THE UNIVERSITY OF
WARWICK

Original citation:
Bhalerao, Abhir, Patanaik, Amiya, Anand, Sarabjot Singh and Saravanan, Pounnusamy (2008) *Robust detection of microaneurysms for sight threatening retinopathy screening.* In: 6th Indian Conference on Computer Vision, Graphics and Image Processing, Bhubaneswar, India, Dec 16-19 2008. Published in: Sixth Indian Conference on Computer Vision, Graphics & Image Processing, ICVGIP 2008 pp. 520-527.

Permanent WRAP url:
http://wrap.warwick.ac.uk/17478

warwick**publications**wrap
highlight your research

http://wrap.warwick.ac.uk

Robust Detection of Microaneurysms for
Sight Threatening Retinopathy Screening

Abhir Bhalerao[1], Amiya Patanaik[2], Sarabjot Anand[1], Pounnusamy Saravanan[3]

[1]Department of Computer Science, University of Warwick, UK abhir@dcs.warwick.ac.uk

[2]Indian Institute of Technology, Kharagpur, India amiyan@gmail.com

[3]Department of Diabetes and Endocrinology,
University Hospitals of Coventry and Warwickshire, UK p.saravanan@warwick.ac.uk

Abstract

Diabetic retinopathy is one of the major causes of blindness. However, diabetic retinopathy does not usually cause a loss of sight until it has reached an advanced stage. The earliest sign of the disease are microaneurysms (MA) which appear as small red dots on retinal fundus images. Various screening programmes have been established in the UK and other countries to collect and assess images on a regular basis, especially in the diabetic population. A considerable amount of time and money is spent in manually grading these images, a large percentage of which are normal. By automatically identifying the normal images, the manual workload and costs could be reduced greatly while increasing the effectiveness of the screening programmes. A novel method of microaneurysm detection from digital retinal screening images is proposed. It is based on filtering using complex-valued circular-symmetric filters, and an eigen-image, morphological analysis of the candidate regions to reduce the false-positve rate. We detail the image processing algorithms and present results on a typical set of 89 image from a published database. Our method is shown to have a best operating sensitivity of 82.6% at a specificity of 80.2% which makes it viable for screening. We discuss the results in the context of a model of visual search and the ROC curves that it can predict.

1. Introduction

Diabetic retinopathy is present in 30% of diabetic population and is the leading cause of blindness in the developed world [4]. Blindness from retinopathy can in theory be prevented but requires regular eye checks and aggressive control of blood sugar levels. Retinopathy may also be an early indicator of other health problems such as heart disease. In the UK, community based retinal screening programmes have been set up where patients with diabetes are routinely screened and images stored in a central database.

For each patient, four images are taken every year. These are then manually graded by trained retinal screeners to identify sight threatening retinopathy. Patients are referred to eye specialists for certain treatments (laser) to prevent them developing advanced retinopathy and blindness; but a significant number of patients (approximately 60-70%) do not have retinopathy. This is a time-consuming process and involves a dozen or so screeners for Warwickshire.

When the patient is screened, one or more digital photographs of each eye are acquired then examined by optometrists who grade the images on a standard scale. The aim is to identify so called "RR" or referable retinopathy when the patient is sent for assessment to a consultant ophthalmologist. In most cases, however, the images show no signs of disease and are classed "R0", and the patient will be asked to return typically after 12 months. Manual grading of these images is a skilled yet laborious task and only about 30 or so cases can be screened by one grader per day. The scale of the problem becomes clear when typically, there are tens of thousands of patients (e.g. Warwickshire region have 35,000 patients) with mulitple images per patient (e.g. 100,000 per year) that need to be manually assessed. These are some of the reasons that a robust, computer aided diagnositc (CAD) system is needed. The cost benefits of a system that could reduce the load, even by 50% are clear. But there also other gains to be had: quality assurance improvements, by maintaining a consistently high-quality standard for the screening; better utilization of skilled graders; and in the near future, using results from the system to perform longitudinal analysis of the disease and its treatment, e.g. [8].

Fundus images are acquired non-invasively, thorough the dialated pupil of the eye, using a digital camera and a flash. Typically, the resolution of state-of-the-art system is very good with up to 2500 x 1500 pixel resolution images being digitized. Microaneurysm (MAs) are the earliest signs of retinopathy at the so called pre-proliferative stage. They persist as the disease continues to progress. If the early signs are not immediately treated, the disease quickly reaches the proliferative stage when new vessels are created

(neovascularization) threaten eye sight and may cause the retina to become detached. The grading process decides on referable cases, "RR", when there is one or more MAs identified on the fundus photographs. These need not be in the critical foveal region; a single optic-disc radius around the fovea is sometimes used to signal more critical cases. MAs are caused by high-sugar levels in the blood which causes the walls of tiny blood vessels (capillaries) to distend and they appear as roughly circular dark spots of fairly consistent size on the red-free images (green channel). They have similar intensity values to the blood vessels and indeed small, spot shaped aretefacts can sometimes appear in and around the vessels. These are true-negative lesions but a trained grader can quickly dismiss these. Latter stages of the disease, blot-aneurysms, cannot normally be confused with MAs as they are diffuse in appearance. The main factors that make the manual and automatic assessment of the disease difficult are:

- fundus images can have significant contrast variation related to the field of view (FoV) and insufficient dialation of the pupil,

- variation in pigmentation because of ethnicity and other types of disease lesions such as those related to age,

- spot artefacts near vessels that have similar shapes and hue

Approaches to MA detection have principally focussed on their intensity characteristics, size and shape. Most of the presented methods have used mathematical morphology operators but all methods perform a series of basic steps: (1) *contrast normalization* to remove shading artefacts (see figure 4); (2) detection of *candidate* lesions (these are dark spots); (3) reduction of the candidate set by eliminating false-positive detections. False-positive detection of regions which are on or near the vessels is particularly problematic. Some reported methods have performed explicit vessel detection to exclude such regions [5, 9, 2]. A number or vessel detection methods exist and can be applied to mask and exclude such detections. Others have extended the basic mathematical morphology framework to locally measure linear structures and suppress candidates at these points.

One of the better performing systems has been reported by Fleming et al. [5] . Their system has a sensitivity of 85.4% with a specificity of 83.1%. Recently, they reported that at 50% sensitivity, an area-under-roc curve (AUR) of 0.911 for MAs. This implies that for grading, their CAD system would be able to reject 50% of the normal cases with high sensitivity (95% or more). Much of the problem however of insufficiently high sensitivity has been because of the false-positives not being filtered correctly. Fleming's method consists of a number of stages: contrast normalization based on a median filter; candidate detection using tophat morphological operations; region growing after thresholding; candidate evaluation based on fitting a 2D Gaussian; background region growing using a watershed segmentation; vessel exclusion using linear structuring elements; further classification of the remaining candidates using a kNN

classifier. The number of steps and associated parameters (thresholds) makes it hard to reproduce their technique and to readily analyse its shortfalls and advantages. In this work, we have consciously kept the number of stages and parameters to a minimum.

Streeter and Cree [10] reported a method which begins by suppressing vessel like structures by first running 8 directional linear structuring elements of size 25 pixels using a tophat step. Then a pair of circular-symmetric matched filters of size 17 and 13 are applied to the output. The output is thresholded and produces an inital set of 250 false positive lesions per image. Regions growing from these seed points with two further intensity thresholds then leaves a set of candidate points. Eighty features each are derived from these candidates and after forward-backward feature selection reduced to 16 discriminating with an LDA classifier. The final classification was tested and shown to produce about 6 false-positives per image at a sensitivity of 56%. Here we also uses matched filters and a feature selection/classifier, but based on an eigen-image analysis of the candidate regions. Our approach only has two user defined thresholds, one of which is set using a supervised analysis of MA and non-MA neighbourhoods and the second which can be taken as the optimal operating point from an ROC analysis.

The main contributions of this paper are: a new fast and robust MA detector based on a complex-valued, circular-symmetric operator; a filtering method on an eigen-image analysis of their shapes; and the use of a model of visual search recently proposed by Chakraborty [3]. The presented method is both simple and verifiable. The eigen-image classifier is highly effective and our results on a standard retinopathy database [11] confirm the visual results for sensitivity and specificity. After a description of the method, presentation of results, we discuss the ROC analysis with reference to the model of visual search and make predictions for the algorithm performance on larger data.

2. Method

The key idea of our approach is that the microaneurysms have a particular shape profile which is very different from other features present in the fundus image like vessels. This is exploited by employing an orientation matched filter. But before this can be done, the images are preprocessed to enhance features and make it suitable for further processing. Thresholding on the output of orientation matched filter obtains a set of potential candidates. This technique also picks up certain noise artifacts which resemble the shape profile of a microaneurysm, to remove these false-positives, eigen image analysis is applied on the potential candidate regions and a second threshold applied on the eigen-space projection of the candidate regions. The main stages of the methods are:

1. contrast normalization using a median filter

2. *blob* detection using a Laplacian of Gaussians (LoG)

3. shape filtering using a circular-symmetry operator on an orientation map of the data; combining the outputs with the LoG produces a candidate set

4. small windows around candidates are then used as features into an eigen-space learnt from ground-truth TPs and FPs.

Only 3 parameters are required: the size of the LoG blob detector; the threshold on the combined shape/orientation field analysis; and the classifier threshold on the eigen-space projection. We show that in fact, that for a given image resolution, only the second threshold requires tuning by finding the best performing value against ground truth data (ROC analysis).

2.1. Speckle reduction and Contrast normalization

The green channel, I, of a colour image of size $W \times H$ pixels is taken and filtered using 3×3 median filter to remove any salt and pepper noise. The median filtered output M is multiplied by a local-mean filtered image, L_{25}, obtained by calculating a moving average mean in windows of size 25×25. At the same time, a confidence map C is produced by thresholding the output against a low intensity value, $I_{min} = 20$:

$$M(x,y) = [M_{3\times 3} * I(x,y)] \times L_{25}(x,y) \quad (1)$$

$$C(x,y) = \begin{cases} 1 & \text{if } L_{25}(x,y) \geq I_{min} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

This is effective at eliminating regions with low contrast and also estimating a quality value for the image, $Q_I = \sum C(x,y)/(W \times H)$.

Due to shape of the eye and improper illumination, fundus images are corrupted by a shading artifact. A simple model for the shading artefact across an image is a multiplicative model $I = \hat{I} \times S_{\times}$, where $S_{\times}$ is the shading bias. To estimate the contrast normalized image, $\hat{I}$, a number of methods have been reported which all assume that the shading bias is low-frequency and the unbiased image is piecewise constant. In this work, a large window median filter is used to obtain an estimate of $S_{\times}$.

$$\hat{S}_{\times}(x,y) = M_{53\times 53} * M(x,y). \quad (3)$$

Then, the contrast normalised image is $\hat{I}(x,y) = M(x,y)/(\hat{S}_{\times}(x,y)+1)$. The confidence map, C, is applied to the output contrast is rescaled to one standard deviation either side of the mean of $\hat{I}$ (figure 4).

2.2. Blob detection

After preprocessing the input, a Laplacian of Gaussians filter is used to enhance small, spot features (figure 1).

$$L_{\nabla^2}(x,y) = \nabla^2(\sigma) * \hat{I}(x,y), \quad (4)$$

$$\nabla^2(\sigma) = -\frac{1}{\pi\sigma^4}\left(1 - \frac{x^2+y^2}{2\sigma^2}\right)e^{-\frac{x^2+y^2}{2\sigma^2}}.$$

This results in strong positive responses for dark blobs of extent σ and strong negative responses for bright blobs of similar size. However, this also enhances features other than MAs like certain parts of vessels and noise artefacts (figure 4).

Figure 1. Top: Laplacian of Gaussians and Difference of Gaussians profiles. LoG used as blob detector is controlled by σ; DoG is the radial magnitude, $f(r; \sigma_2)$, of circular-symmetry filter, $\mathbf{f}$. Bottom: LoG and orientation filtering operations on a single MA region shown in $\hat{I}$.

2.3. Circular-Symmetry Operator

The complex valued, circular symmetry operator [6]

$$\mathbf{f}(x,y) = f(r)e^{i2(\beta\phi+\alpha)} \quad (5)$$

where $f(r)$ is a radial magnitude function, is correlated with a orientation map of the contrast normalised image,

$$\mathbf{R}(x,y) = m(x,y)e^{i\Theta(x,y)} * \mathbf{f}^*(x,y) \quad (6)$$

and the magnitude and phase of the orientation map is obtained by running a pair of directional filters (Sobel operators in s_x and s_y directions)

$$m(x,y) = \|\hat{I}(x,y) * (s_x(x,y), s_y(x,y))^T\|, \quad (7)$$

$$\Theta(x,y) = arg[\hat{I}(x,y) * (s_x(x,y), s_y(x,y))^T].$$

Here the superscript, $*$, is a complex conjugation, and ^{T} the transpose.

The circular-symmetry operators, $\mathbf{f}$, belongs to a class of *double-angle* filters which can be varied by changing the harmonic number β and the phase offset α. We have used the simplest form, $\beta = \frac{1}{2}, \alpha = 0$, which defines *single-angle* circular patterns of orientation (see figure 4) and matches well the appearance of the orientation field around a microaneurysm. The output of the circular-symmetry operator filtering, $\mathbf{R}$, is also complex valued and has positive phases for dark spots, such as MAs, and negative phase for

white spots such as hard-exudates. This allows us to easily distinguish MAs in the presence of proliferative lesions in advanced retinopathy. The radial profile, $f(r)$, of the operator is also matched to the expected amplitude profile of the edges-map around a MA as the difference of Gaussians

$$f(r; \sigma_2) = e^{-\frac{r^2}{2\sigma_2^2}} - e^{-\frac{r^2}{4\sigma_2^2}}. \tag{8}$$

We have tuned σ_2 to correspond with σ used for the LoG filtering (see figure 1).

2.4. Candidate selection from matched filtering

The outputs of LoG and Circular-symmetry operators are combined by taking the product, $L_{\nabla^2} \times \|\mathbf{R}\|$. This map is thresholded at t, to locate a set of candidate lesion regions. The centroids of these regions are stored for further processing by the next stage. In the evaluation presented below, we explain how the value of t is chosen by finding the best operating-point on a receiver operator characteristic (ROC) curve.
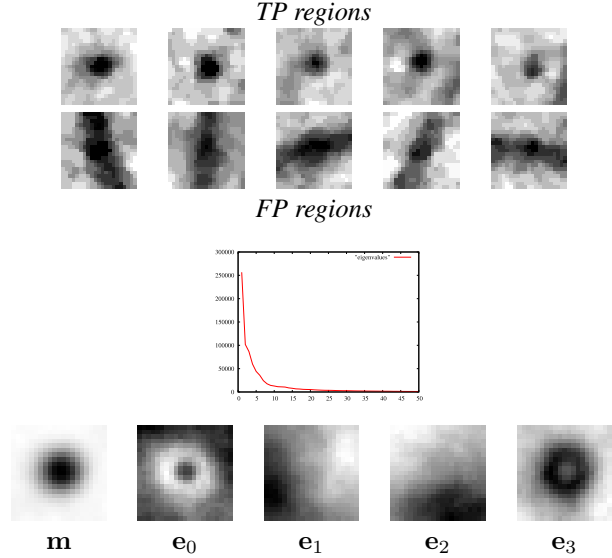
2.5. Eigen-image analysis of candidate regions

At this point, the candidate regions are likely to include non-MA lesions as well as MA regions. We can class these as false-positive detections, FPs, and true-positive detections, TPs. Figure 2 shows a few examples of FP and TP regions. These are obtained by taking a small, square window of size B around each candidate centroid.

To further discriminate between TPs and FPs, an eigen-image analysis is employed (see for example [7]). Each candidate region is projected into an eigen-space obtained from a principal component analysis (PCA) of a set of training sample regions - like those shown in figure 2. Each square region can be represented by a B^2 dimensional vector, $\mathbf{v}_i$, then the PCA performed by calculating the mean and co-variance of the set of N training samples $\mathbf{v}_i$. The projected region is obtained by the linear transformation

$$\mathbf{w}_i = \mathbf{W}^T \mathbf{v}_i \tag{9}$$

where the projection matrix $\mathbf{W}$ is the M principal eigen-vectors, $\mathbf{e}_j$, of the covariance of the data: $\sum (\mathbf{v}_i - \mathbf{m})(\mathbf{v}_i - \mathbf{m})^T$, $\mathbf{m} = \frac{1}{N} \sum \mathbf{v}_i$. Now, the dimensionality of $\mathbf{w}$ is M and significantly less than that of the data, $\mathbf{v} \in \mathbb{R}^{B \times B}$. In the evaluations we used a window of size, $B = 21$ and the variance analysis of the PCA (figure 2) shows that only the first 10 or so modes are significant. Hence, $M = 10$ is reasonable. In fact, if we visualize the training regions using only the first three principal dimensions, it appears that the first mode, $\mathbf{e}_0$, is sufficient to discriminate the TP and FP regions (figure 3). Hence, we can use a simple threshold of approximately 1000 on the projection $\mathbf{e}_0^T \mathbf{v}_i$ to discriminate any remaining false-positive regions.

TP regions

FP regions

m   **e**$_0$ **e**$_1$ **e**$_2$ **e**$_3$

Figure 2. Top: Selection of TP and FP candidate regions used to perform PCA. Middle: Plot of eigen values from PCA of TP and FP training regions. Bottom: the mean eigen-image, $\mathbf{m}$, and the first 5 eigen vectors of the PCA. Variation is circular-symmetric in 1st and 4th modes, while the remainder is a result of local contrast variation.


Figure 3. Scatter plot showing coordinates of TP and FP candidate regions in principal coordinates of eigen-space. A single mode, e_0, is strongly discriminating at value 1000.

3. Evaluation

3.1. Materials

We have used a set of 89 images from the DIARETDB1 database published by the University of Kuopio, Finland [11]. This database is a set of high-resolution RGB fundus images of size 1500 x 1152 pixels (50 degree field of view) together with labelled ground-truth by 3 experts of various lesions: MAs, hemorrhages, hard-exudates, soft-exudages and neovascularization. In this work, we are interested in only the MA labelled images.
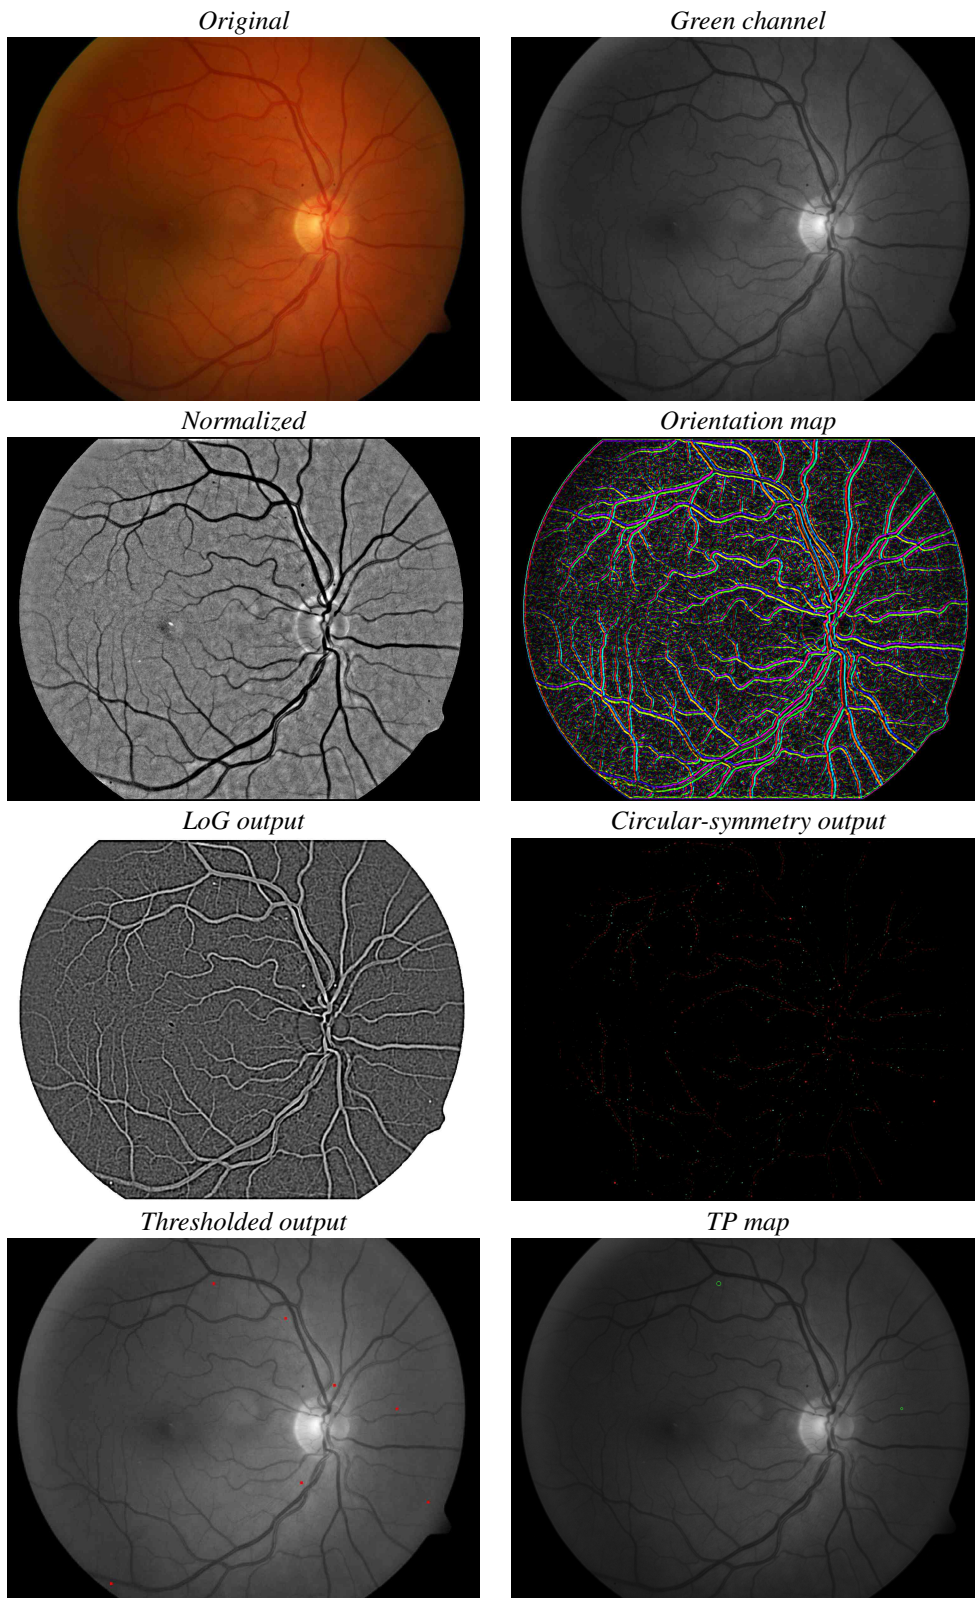
Figure 4. Stages of MA detection method on image078 of DIARETDB1 database. MA detections shown in red on green channel (bottom left). True-positive detection from ground-image.

There were 205 separate MA lesions across the entire set, with 9 images containing no MAs. Figure 4 illustrates the results on one image of the test data (image078). This image had two true-positive lesions marked on the ground-truth, plus two true-negatives. The automatic detector has found two microaneurysms not picked up by expert reader and also one of the two true-negatives which are the result of dust on the camera lens (these persist on other images of the DIARETDB1 dataset).

3.2. ROC Analysis

Receiver Operator Characteristics (ROC) analysis was applied to validate the MA detection method against the example images from the DIARETDB1 database. All the images were processed by the automatic method with the LoG and circular-symmetric operators tuned to with $\sigma = 3$ and $\sigma_2 = 1$. The eigen-image classfier threshold was set to 1000. The threshold t on the product of the matched filters $L_{\nabla^2} \times \|\mathbf{R}\|$ was varied between the interval $t = [160, 240]$, to produce a set of MA detections for each input image. These pixels from these detection were then tested with the hand labelled ground truth to calculate sensitivity (true-positive rate), and specificity (true-negative rates):

$$Sen = \frac{TP}{TP + FN} \qquad Spec = \frac{TN}{TN + FP},\qquad (10)$$

where TP is true-positive, FP is false-positive, TN is true-negative and FN is false-negative. The sensitivity is the proportion of actual MA regions that are found and the higher this number the more accurately the algorithm is detecting the disease regions. The specificity is the proportion of non-MA regions correctly detected. In the ROC plots presented below, $1 - Spec$ is plotted, which is the false-positive rate, and this has to be small indicating that algorithm does not produce too many false-alarms. For our application, reducing the $1 - Spec$ rate will correspondingly allow a greater proportion of the images to be discarded without referral to an expert. However, since the DIARETDB1 images do not show the TN lesions, we have to use a statistical argument to determine which pixels could be classed as such.

Blurred ground truth *Confidence intervals*

Figure 5. Confidence intervals $TN = [0.3, 0.75]$ (green), $TP = [0.75, 1.0]$ (red) on ground-truth data used to generate ROC. Example shown for image003 of DIARETDB1 database.

To acertain the number of true-negatives, we took confidence intervals of the labelled ground-truth (GT) data. The GT images are labelled in an additive way; i.e. each grader's markup is summed, and no restriction was placed on how the labelling should be done [11]. It is then possible to blur this markup and take confidence intervals as shown in figure 5. For example, to find those pixels that have the highest agreement between the graders, a confidence level of 0.75 can be taken. Similarly, we can take an intermediate interval, e.g. [0.3, 0.75], to locate pixels which are potential true-negative regions because at least one grader believes they could be lesions. Figure 6 is a plot of the ROC analysis for all the DIARETDB1 images. The best operating point on this curve is at a sensitivity of 82.6% at a specificity of 80.2%.

Figure 6. ROC curve for threshold t on matched filter product on DIARETDB1 database. The best operating point is 82.6% at 80.2%.

3.3. A Model for Visual Search

One of the problems with using ROC curves is that although they can be an effective way of finding the best operating point for a given algorithm, and assess its performance independently, it can be hard to compare between ROC curves from other algorithms when the data is different. In particular, it would be informative to predict the algorithm performance as if the data were worse or better than that given. Or, indeed, to hypothesize how it might perform on different data knowing only, for example, the relative incidence of disease in the poplulation. In [3], Chakraborty proposed a model for visual search which can be used to predict ROC curves for a free-response paradigm such as an observer searching for lesions in an image. The model is parameterized by perceived lesion-to-noise ratio and how good the observer is at detecting the lesions and non-lesions. Here, we fit the model to the performance curves of our algorithm and use it to predict the performance as the model parameters are varied.

The visual search model proposed in [3] makes a number of statistical assumptions. It assumes that the observer is presented with an image containing a number of lesion sites, some of which are signal sites and some noise sites. The observer then makes decisions for each potential site marking it as a TP or not. The parameters of the model therein are

1. The number of noise sites, n, are drawn from a Poisson distribution with rate parameter λ.

2. The number of signal sites, u, are drawn from a Binomial distribution with trial size s, and success probability ν.

3. The decision at each site is based on a random variable, ζ, drawn independently from a unit variance, Gaussian distribution, $N(m, 1)$, where $m = 0$ for a noise site, and $m = \mu$ for a signal site.

Translating these parameters to our MA detection problem, we can deduce that the trial size of the number of signal sites, s, is the same as the estimate of the expected value for the number of MAs per image. In our data, this was $205/80$, as 9 images had no MAs and a total of 205 were present, hence $s = 2.5$. We can relate the success probability, ν, to some large value as we expect that our method is able to find all potential MA sites. It is designed to over-segment the data and the expectation is that all the signal sites will be located every time. We can nominally set $\nu = 1.0$. The parameter, μ, which separates the decision distributions indicates the discriminating power of the algorithm: the larger the μ, fewer errors are made and the better are the ROC curves which result.

To plot an ROC curve, in [3], the sensitivity and specificity under the above model are analytically derived as

$$Spec(\zeta|n) = e^{-\frac{\lambda}{2} + \frac{1}{2}\lambda\, erf(\frac{\zeta}{\sqrt{2}})}, \tag{11}$$

$$Sen(\zeta|\mu, \lambda, \nu, s) = 1 - \left(1 - \frac{\nu}{2} + \frac{\nu}{2}\, erf(\frac{\zeta - \mu}{\sqrt{2}})\right)^s. \tag{12}$$

$$e^{-\frac{\lambda}{2} + \frac{1}{2}\lambda erf(\frac{\zeta}{\sqrt{2}})}.$$

For our experiments, since our threshold t takes arbitrary values, we do not know the spread of the decision variable ζ, so we add additional parameters μ_n and ς to model the spread of the Gaussian normal distributions from which it is drawn. This can be simply achieved by substituting $\zeta \to \frac{\zeta - \mu_n}{\varsigma}$ into the error functions. We used non-linear regression fitting to estimate the unknown parameters: λ, $\{\mu_n, \varsigma\}$ and μ for our experimental data and plotted the resulting curve in figure 7.

Finally, we illustrate the predictive abilities of the ROC model by altering the estimated parameters. In figure 8, the separation of the noise and signal normal distributions, μ, is increased and decreased which correspondingly improves or worsens the performance of the algorithm. Similarly, changing the expected rate of MAs per image, controlled by s, worsens the performance significantly if $s = 1.5$.

4. Conclusions

Development of computer aided methods for the screening for diabetic retinopathy is an important and active area of research in medical image analysis. Although a number of algorithms have been proposed, none have shown the
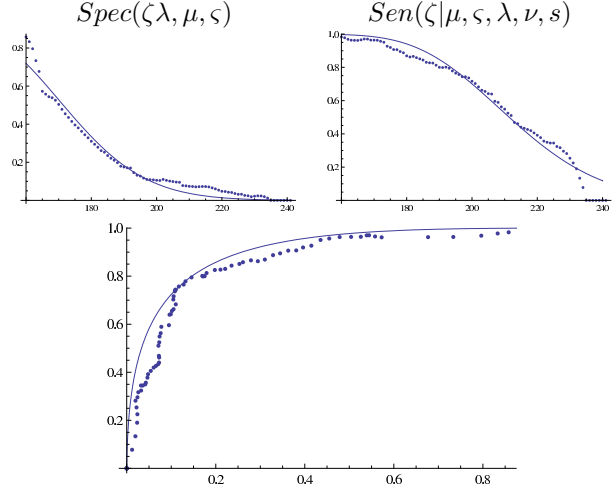
Figure 7. Fitting to ROC data using model of visual search. $s = 2.5, \lambda = 7.6, \nu = 1.0, \mu = 58.5, \varsigma = 31$.

Figure 8. Predictions from ROC analysis. Top: Varying the separation power of detector: $\mu = 68.5, \mu = 38.5$. Increasing μ improved the performance. Bottom: Varying the incidence of MAs per image $s = 3.5, s = 1.5$. Decreasing s, worsens the performance.

promise of performance which can easily be scaled. Many of the methods presented elsewhere use sets of arbitrary image processing steps and perhaps do not fully exploit the intensity and shape characteristics of the microaneurysm lesions. This make it hard to reproduce from explainations in the literature and harder still to tune them to their best performance. We have presented a model based approach to the problem that employs fairly standard linear filtering and eigen-image analysis. We employ both intensity and shape filters, using a circular-symmetric operator on an ori-

entation map of the input. The output of these allows us to quickly identify a set of potential MA regions. These are then classified using a PCA and a linear classifier.

The method only requires the tuning of two parameters that are tied to image resolution and leaves a threshold that can be adjusted using a simple ROC analysis. The resultant method has been shown to perform comparably well on a set of 89 standard retinal fundus images for which ground truth was made available. Our ROC analysis of the detector performance have shown it to produce a sensitivity of 82.6% at a specificity of 80.2%. At a specificity of 67% the sensitivity is 89% which compares well with the results published elsewhere (see for example [9]). However, given that we have only run the detector on 89 images, it remains to be seen if the performance can be maintained on larger data where the numbers of true-positive images will be much smaller. We shown how a model for visual search can be used to fit to the ROC data and enable predictions on the performance of the method on enhanced or degraded data to be made. We believe that this sort of extrapolation could be useful to gauge both the power and robustness of a given method.

A number refinements are still be made to our system. We have not fully explored the eigen-image analysis and it is apparent that there is some residual contrast variation that the initial normalization step is not removing fully. We are exploring parametric normalization schemes which may cure this problem [1]. In the near future, we will be acquiring a much larger set of example images (1000 or more), to test further the performance. Our current results are encouraging but we acknowledge that the lower incidence rates of MAs in larger cohorts may be problematic.

References

[1] A. Bhalerao and S. Anand and P. Saravanan. Retinal Fundus Image Contrast Normalization using Mixture of Gaussians. In *Proceedings of 42nd Asilomar Conference on Signals, Systems and Computers*. IEEE Signal Processing Society, 2008.

[2] D. Abràmoff, M. Niemeijer, M. S. A. Suttorp-Schulten, M. A. Viergever, S. R. Russell, and B. van Ginneken. Evaluation of a System for Automatic Detection of Diabetic Retinopathy from Color Fundus Photographs in a Large Population of Patients With Diabetes. *Diabetes Care*, 31:193–198, 2008.

[3] D. P. Chakraborty. ROC curves predicted by a model of visual search. *Physics in Medicine and Biology*, 51(14):3463–3482, 2006.

[4] J. Evans, C. Rooney, S. Ashgood, N. Dattan, and R. Worlmald. Blindness and partial sight in England and Wales April 1900-March 1991. *Health Trends*, 28:5–12, 1996.

[5] A. D. Fleming, S. Philip, K. A. Goatman, J. A. Olson, and P. F. Sharp. Automated Microaneurysm Detection Using Local Contrast Normalization and Local Vessel Detection. *IEEE Transactions on Medical Imaging*, 25(9):1123–1232, 2006.

[6] H. Knutsson and G. H. Granlund. Apparatus for determining the degree of variation of a feature in a region of an image that is divided into discrete picture elements. *Swedish Patent 8502569-0 1986/US Patent 4.747.151, 1988*, 1996.

[7] X. Mo and R. Wilson. Discriminant Eigengels: A Statistical Approach to 2-D Gel Analysis. In *Workshop on Genomic Signal Processing and Statistics*, Raleigh, North Carolina, USA, October 2002.

[8] H. Narashimha-Iyer, A. Can, B. Roysam, C. V. Stewart, H. L. Tenenbaum, A. Majerovics, and H. Singh. Robust Detection and Classification of Longitudinal Changes in Color Retinal Fundus Images for Monitoring Diabetic Retinopathy. *IEEE Transactions on Biomedical Engineering*, 53(6):1084–1098, 2006.

[9] S. Philip, A. D. Fleming, K. A. Goatman, S. Foneseca, P. Mcnamee, G. S. Scotland, G. J. Prescott, P. F. Sharpe, and J. A. Olson. The efficacy of automated "disease/non disease" grading for diabetic retinopathy in a systematic screening programme. *British Journal of Ophthalmology*, 91:1512–1517, 2007.

[10] L. Streeter and M. J. Cree. Microaneurysm Detection in Colour Fundus Images. In *Proceedings of Image and Vision Computing 2003, New Zeland*, pages 280–285, 2003.

[11] T. Kauppi and V. Kalesnykiene and J.-K. Kamaraniene and L. Lensu and I. Sorri and A. Raninen and R. Voutilainen and J. Pietilä. The DIARETDB1 diabetic retinopathy database and evaluation protocol. In *Proceedings of British Machine Vision Conference, 2007*, pages 252–261, 2007.