Analysis of Retinal Images Using Mathematical Morphology

Piotr Jasiobedzki, Chris J. Taylor
Wolfson Image Analysis Unit, Dept. of Medical Biophysics,
University of Manchester, Manchester M13 9PT, UK
Email: pcj@wiau.mb.man.ac.uk

Abstract

An automatic analysis of angiographic retinal images is discussed. A detection method based on texture differences between perfused and non-perfused regions is described. First the image is tessellated into a large number of primary regions using the positions of vessels and other structures in the image. The image is represented as a region adjacency graph. Arcs in the graph are classified as vessels using features obtained from grey level profiles in the original image. For each region the degree of non-perfusion is estimated by measuring textural properties using mathematical morphology. Adjacent regions are subsequently merged using information on the degree of perfusion.

Keywords: mathematical morphology, texture, retinal images, blood vessels detection, microcapillary perfusion.

1. Introduction and Medical Background

The optical system of the eye permits direct observation of retinal blood vessels including some capillaries. Diabetes Mellitus, which affects 2% of the population, is one of the most important diseases which causes abnormalities in the retinal capillaries. Such Diabetic Retinopathy is the commonest cause of blindness in people in developed countries.

Changes in the retinal vascular network in diabetes include occlusion of capillaries with formation of non-perfused regions. This process is accompanied by local dilation of adjacent vessels and formation of microaneurysms, small (25–100 μm in diameter) circular blood filled sacks, which tend to surround areas of non-perfusion. Microaneurysm counting has been used as an indirect method of grading the severity of Diabetic Retinopathy [Kohner et al. 86]. After an initial increase in the microaneurysm number there is a subsequent decrease as the retinopathy progresses by extensive capillary closure. Therefore a microaneurysm count can only be regarded as a reliable measure of the severity of early retinopathy.

Fluorescein angiography is an established clinical technique for observing the retinal circulation. Fluorescent dye is injected into a vein in the arm and a sequence of photographs of the illuminated retina is taken through the pupil. Standard images taken for clinical tests cover a region about 10 mm in diameter around the fovea and the optic disc (figure 1). Vessels filled with dye are observed as bright

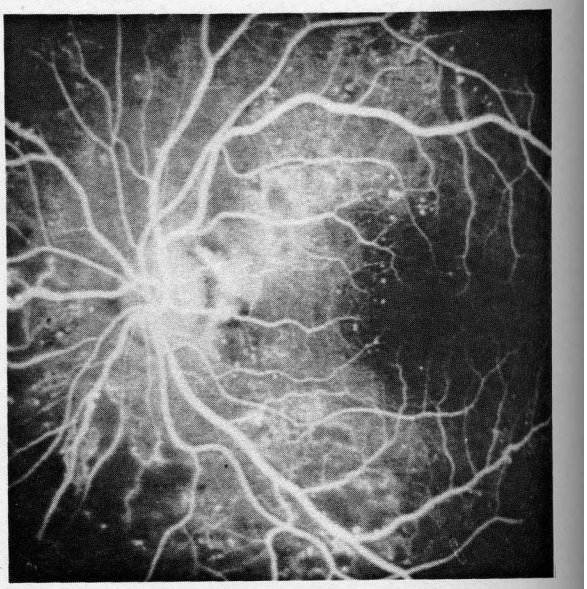

Figure 1. Angiographic image of the retina.

ribbons on a non-uniform background. The vessels are roughly cylindrical and assuming uniform flow of the dye and low attenuation of light by the dye and blood, the grey level profile taken across the vessel can be approximated by an ellipse. The maximum

intensity level for each vessel depends on its diameter. Major vessels can be easily detected, but contrast decreases with decreasing vessel diameter. Separate microcapillaries can be observed in the region around the fovea where there is only a thin layer of microcapillary net and the contrast is

good. Further from the fovea where the microcapillary net is much thicker, there is less contrast with the background and detection of individual capillaries is not feasible. We are attempting to measure degree of perfusion (density of microcapillary net) directly in the retinal image.

2. Previous Work

Previous work on analysing retinal images has concentrated mainly on detecting vessels and microaneurysms. Detection of vessels was usually performed in two steps:

(i) detection of vessel segments,

(ii) identification of vessel tree by joining segments and resolving connectivity at bifurcations and crossing points.

Vessel segments have been detected using various methods. Tanaka [Tanaka 80] used adaptive thresholding followed by binary thinning to detect major vessels. Katz [Katz et al. 88] segmented and thinned output from the Sobel edge detector to outline vessels. Only major vessels were detected since smaller ones could not be distinguished from other structures. Chaudri [Chaudri et al. 89] described a more robust approach based on detecting local vessel direction by convolving the image with a set of masks. The masks were obtained by rotating a linear feature detector of a gaussian cross-section which roughly matched the grey level profile across vessels. This method allowed detection of much smaller vessels than found previously [Katz et al. 88]. The technique was only applied at a single scale but application of a range of kernel sizes would have allowed a much wider range of vessel widths to be detected. Application of a different linear feature detector ("road detector") has been also investigated [Akita 82].

The structure of the vessel tree has been obtained either directly from detected ribbons [Katz et al.88] or from branching and crossing points which were identified in the binary image [Tanaka 80]. Vessel segments joining these points were then extracted from the image. Akita [Akita 82] described the use of relaxation labelling to provide a connectivity interpretation for the vessel segments and to label vessels as arteries or veins. Caliber of vessels at arterio-venous crossings is used for ranking hypertension.

Lay detected microaneurysms using a morphological "top hat" transformation of the positive image [Lay, Badouin 84]. The top hat transformation was performed by subtracting the image opened using linear structuring elements from the original. The resulting image contained objects of circular shape and of diameter less than the structuring element length.

The problem of detecting non-perfused zones has received little attention. Goldberg [Goldberg et al. 89] described results of clustering regions of similar grey level. The image was smoothed and segmented, with a fixed threshold, into perfused and non-perfused regions. This approach did not take into account non-uniform background and did not try to relate detected regions to any constant objects (major vessels, optic disc, etc.) or to monitor increase in number and extent of non-perfused regions over a period of time. We are not aware of any other published work in this area.

3. Outline of the Proposed Method

The requirements for the analysis of retinal images for the purpose of detecting and monitoring the early stages of Diabetic Retinopathy can be summarised as follows:

(i) to detect and measure non-perfused regions (without a visible microcapillary net),

(ii) to relate the position of non-perfused regions to landmarks in the image (optic disc, fovea, position of major vessels),

(iii) to detect and compare the same regions in later images.

The approach we have taken follows the method used in the manual detection of non-perfused regions by an ophthalmologist. Regions which contain a fine and regular structure of capillaries are regarded as well perfused. Areas where this structure is not visible are regarded as non-perfused. Larger vessels do not directly contribute to perfusion. Boundaries of non-perfused regions are usually aligned with the position of vessels or are created by joining adjacent groups of microaneurysms. The automatic processing of each image is performed in the following steps:

(i) detecting the network of vessels,

(ii) defining primary regions,

(iii) measuring texture in primary regions,

(iv) merging similar adjacent regions.

Detection of the network of vessels is based on detecting linear structures in the image and identifying vessels among them.

Defining primary regions is based on partitioning the image into irregularly shaped areas using the positions of vessels and chains of microaneurysms as boundaries.

As mentioned in the introduction, it is possible under good conditions to detect single capillaries in a small area around the fovea. In other regions of the observed image the background has less contrast and capillaries form a more complicated and thick layer of entangled vessels. Our approach is based on measuring texture in order to grade regions. The texture is measured for each region separately using mathematical morphology.

Adjacent regions are merged together if their textural properties are similar and their boundaries are not major vessels. It is essential for the purpose of this work to relate the position of each detected region to a constant set of local landmarks such as the fovea, optic disc and major vessels. Our approach aligns the regions to the position of vessels and allows us to map detected regions to these elements of the retina which do not change even over a long period of time. One of the problems in analysing textured images is reliable identification of the boundaries between regions of different texture. We solve this problem by using non-textural information to define the boundaries.

4. Detecting the Network of Vessels

Detection of the network of vessels is performed in two stages. First linear structures corresponding to vessels, lines of adjacent microaneurysms and other structures present in the image are detected, and represented as a network of arcs and nodes. In the next step each of the arcs is processed separately by referring to the original image and labelled as vessel or other structure.

The original image (figure 2) is smoothed by morphological opening and resampled giving a low resolution version of the image. The resulting image is thinned to give one pixel wide lines. Thinning is performed by conditional erosion in a 3 x 3 neighbourhood using twelve directional masks. This is based on the method of [Goetcherian 80] which uses only eight masks. Four additional masks are used for pruning the image of lines with one unconnected end.

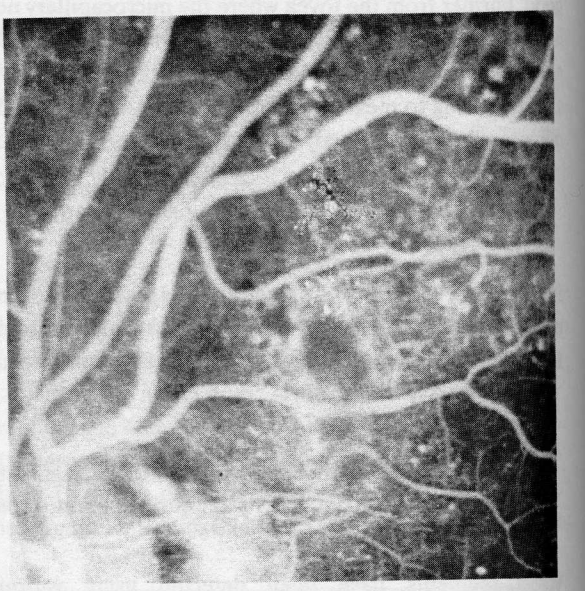

Figure 2. Fragment of the image selected from figure 1.

For each mask two neighbourhoods are defined: L and H. If the value of the central pixel g (m, n) is higher than all pixels in the L neighbourhood and lower or equal to pixels in the H neighbourhood this pixel is set to the maximum value of the pixel in the L neighbourhood. If not it is left unchanged:

$$\left(g(m,n) > \max_{(i,j)\in L}(g(i,j))\right) \wedge \left(g(m,n) \leq \min_{(i,j)\in H}(g(i,j))\right)$$

$$\Rightarrow g(m,n) = \max_{(i,j)\in L}(g(i,j))$$

Each mask is applied in turn.The processing is continued until the image is unchanged by subsequent thinning passes by the set of masks. Thinning creates plateaus of uniform intensity which are surrounded by ridges. The ridges are one pixel wide with the exception of junctions, where ridges of different intensities meet.The connectivity of all connections of ridges is retained (figure 3).

The thinned image is segmented by detecting inner points on the plateaus i.e. points which have all 8 - connected neighbours of equal value. The remaining points are binary thinned and give truly one pixel wide arcs (figure 4). The arcs form a network with all end points connected.

The arcs in the binary image correspond to different structures in the original image. Some of them may represent vessels, lines of microaneurysms or boundaries between re-

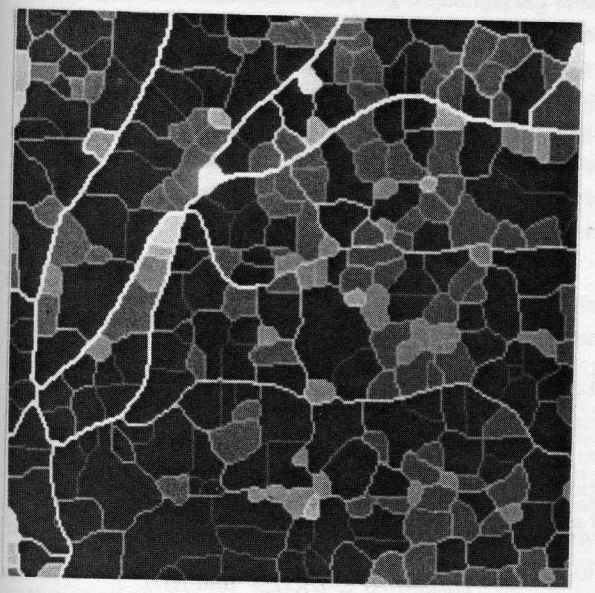

Figure 3. Smoothed, resampled and thinned version of image from figure 2.

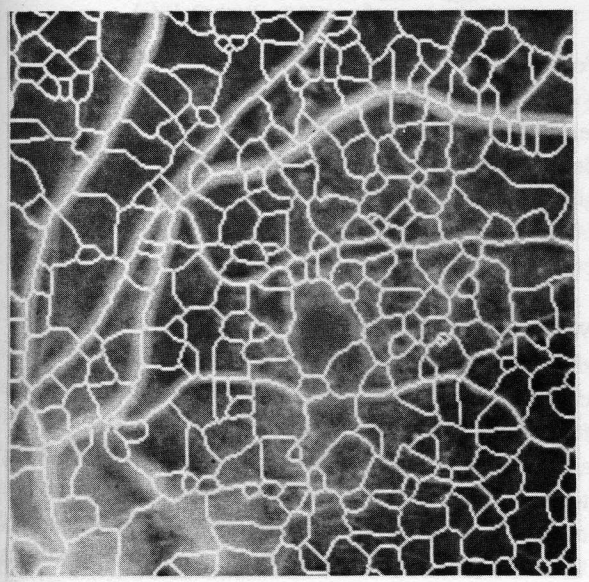

Figure 4. Network of arcs detected in the image from figure 3.

gions of different intensity or texture. The arcs are stored in a region adjacency graph and are used for three purposes:

(i) to delineate primary regions,

(ii) to merge primary regions into larger patches,

(iii) to enable identification of vessel tree.

For the purpose of region delineation information on the type of object (vessel, line of microaneurysms, artifacts) an arc corresponds to is not essential. However as experiments on extracting textural features show it helps to reliably calculate textural features if the region is reduced so as not to include parts of larger vessels. In order to merge adjacent regions into large patches or to identify the vessel tree it is necessary to detect vessels among the arcs in the graph.

To identify the actual type of object that each arc represents, it is necessary to refer to the original image and analyse the grey level distribution around the arc. This is performed by generating a sequence of lines perpendicular to the arc and sampling the original image along the lines. The length of lines is larger than the maximum vessel width. Each line is placed in such way that it is perpendicular to the arc and the middle point on the line coincides with a point on the arc. Grey level profiles are obtained by sampling the image along these lines. Each profile is convolved with a derivative mask to give a gradient profile. The positions of two extreme values of the gradient, to the left and right of the centre are detected. The fragment of the original profile between the detected points is selected as the most likely location for a vessel. An ellipse is fitted to this section of the profile using a least square error criterion. Figure 5 shows the grey level profile taken across a vessel and it's approximation by an ellipse. Width, height and residual error of the fitted ellipse are measured.

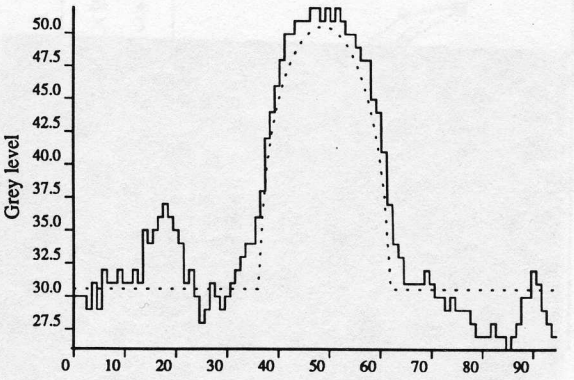

Figure 5. Profile taken across a vessel and it's approximation by an ellipse.

For each arc several profiles are sampled and analysed, and following parameters are derived:

(i) mean and standard deviation of the width of the fitted ellipses,

(ii) mean and standard deviation of the aspect ratio (ratio of height and width) of the fitted ellipses,

(iii) mean and standard deviation of the residual error of the fitted ellipses.

These six features are used for classification of arcs into vessels or other structures currently using a box classifier. Limiting values for each feature have been obtained by training on images with a wide range of vessel widths.

Classification works well for isolated vessels. At some bifurcations or crossings of vessels the grey level profile taken in direction perpendicular to the arc is no longer of elliptical shape, and the arc is not properly classified. This does not affect detection of primary regions because of the method used for drawing detected arcs (see below).

5. Detecting Primary Regions

Having detected vessels in the image it is possible to delineate primary regions in images in such way that their boundaries are aligned with the edges of vessels. Boundaries of regions are obtained by drawing the arcs which have been classified as vessels and dilating them using a circular structuring element of diameter equal to their width. Other arcs are drawn as lines of unit width. Figure 6 shows the primary regions detected after labelling arcs in the network in figure 4.

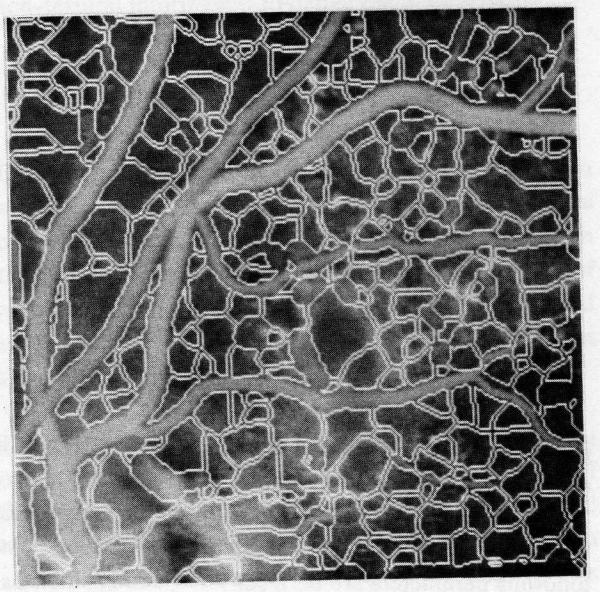

Figure 6. Primary regions detected in the image in figure 2.

The image partitioned into regions in this way, is represented as a region adjacency graph. Information regarding topological relations between regions, connectivity of arcs forming boundaries and properties of regions and boundaries is stored in the graph.

6. Measuring Texture in Primary Regions

Mathematical morphology provides a way of describing sizes and shapes of objects and their relative positions in binary or grey level images. The basic morphological operations for a function $f(y)$ using a binary structuring element B are defined as [Serra 82]:

$$erosion : (f \ominus B)(x) = \min_{y \in B_x}(f(y))$$

$$dilation : (f \oplus B)(x) = \max_{y \in B_x}(f(y))$$

$$opening : f_B = ((f \ominus B)(x) \oplus B)(x)$$

$$closing : f^B = ((f \oplus B)(x) \ominus B)(x)$$

Following detection, each of the primary regions is analysed separately in order to measure its texture. We are interested in differentiating between regions which contain a dense and regular structure of microcapillary net and regions where this structure is severely reduced. The number of capillaries and areas between them reflect the degree of perfusion of a region.

Werman [Werman, Peleg 85] calculates morphological erosions and dilations using linear structuring elements of increasing lengths, oriented at four directions. This yields a feature matrix for each texture which is then used for classification purposes. Serra [Serra 82] uses random set theory to model texture. A texture is described by a probabilistic model. Covariance estimated for the model and measured for the image are used in classification. Maragos [Maragos 89] defines pattern spectrum as a function of structuring element size. The value of the pattern spectrum for a given size corresponds to the difference between images opened (closed) with consecutive sizes of structuring elements. The entropy of the pattern spectrum is used a texture measure.

A box measure of a fractal surface as described in [Mandelbrot 82] can be implemented using morphological dilations and erosions with a three dimensional structuring element.

This method was described for the first time in [Peleg et al. 84] without actually calling it morphological. An algorithm corresponding to dilation with a pyramid shaped structuring element was used. Rigaut [Rigaut 88] describes segmentation of biological images by measuring fractal dimension in windows. Textured regions do not need to have fractal properties over a large range of scale to be classified properly.

The main problem in using morphological operations on images to describe texture is that it is possible either to use a combination of many measures as in [Werman, Peleg 85] or to try to model the process which generates the texture, estimate its parameters and limit interest to a class of applications. In our application we are interested in differentiating between regions which contain dense structure of microcapillaries and regions where this structure is severely reduced. Major vessels do not contribute to the degree of perfusion of the region. The size of small areas between microcapillaries indicates the degree of perfusion.

The one dimensional case is illustrated in figure 7. Dia-

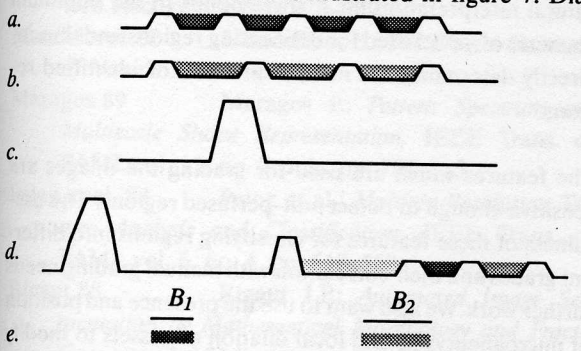

Figure 7. The effect of closing on 1D profile: (a) profile across well perfused region, (b) profile across less perfused region, (c) profile across non–perfused region, (d) profile across a "real" region, (e) two structuring elements of different sizes.

grams *a* to *d* represent grey level profiles obtained by sampling the image along lines in regions with different degrees of perfusion. Profile *a* corresponds to a region with a relatively dense and regular network of capillaries. In profile *b* this density is lower, while in *c* only one large vessel is present. Profile *d* is a more realistic illustration of the expected appearance. Closing with a structuring element B_1 fills the valleys in the profile *a* increasing the average grey level but does not change the profiles *b* and *c*. Increasing the structuring element size to B_2 fills the valleys in profile *b* but *c* remains unchanged. For a periodic texture one may expect a significant increase in average grey level intensity for a precisely defined size of the structuring element. In the real case we obtain profiles more similar to *d*. Then for a region containing a microcapillary network we obtain responses for a wide range of structuring element sizes. For non–perfused regions we obtain some response due to the presence of noise and other visible structures unrelated to the capillaries.

In the two dimensional case it is possible to select the shape and size of the structuring element. Larger sizes of structuring elements can be obtained by multiple application of the structuring element of unit size:

$$f^{(n+1)B}(x) = (((f \oplus nB) \oplus B)(x) \ominus (n+1)B)(x)$$

$$n = 0, 1, 2, \ldots$$

$n = 0$ *corresponds to the original image*

On a square grid a circular structuring element is usually approximated by an octagon obtained by alternating plus and square shaped structuring elements. In our case we are interested in texture with no particular orientation so we use a pseudo–circular structuring element.

The diagram in figure 8 shows changes of average grey level

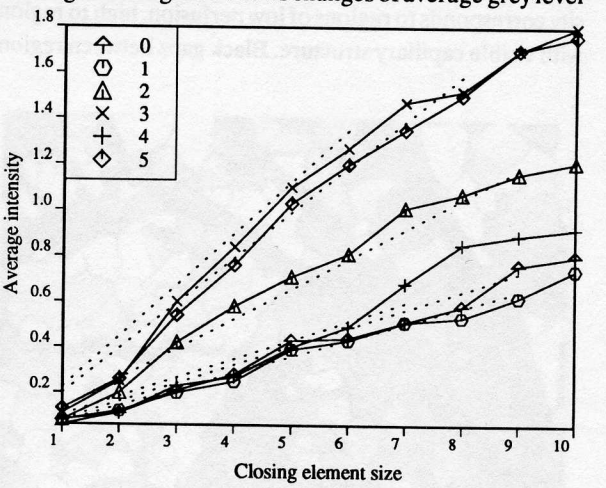

Figure 8. The effect of closing the image on average intensity for selected regions from figure 6: regions 0, 1, 4 are non–perfused, regions 2, 3, 5 are perfused.

as a function of the size of structuring element for number of perfused and non–perfused regions. The points on each of the curves correspond to the difference between the closed and the original image. The difference is calculated for the whole region and normalised by its size. For non–perfused regions the slope is low (curves *0, 1, 4*), for per-

fused ones (curves *2, 3, 5*) the slope is high. We calculate the least square error (LSE) straight line for the initial part of each curve. The number of points for each fit is selected so as to limit the error to a predefined constant value. We use the slope of this line as a feature for classification purposes.

In addition we use the opening of the image in a similar way in order to measure the size and number of the capillaries in the region. We measure the normalised difference between original and opened images as a function of the structuring element size. The slope of the LSE line fitted to the initial fragment of each of the curves is used as a second feature to describe regions.

Thus we currently use two features. One of which depends on the size and number of gaps (density of capillaries) in the region, whilst the other depends on the size and number of capillaries. In practice it is unlikely that these features are truly independent but they encode essentially different properties of the texture.

The image in figure 9 illustrates feature values obtained for closing which have been linearly mapped into the corresponding regions of the image shown in figure 6. Low intensity corresponds to regions of low perfusion, high to regions with visible capillary structure. Black gaps between regions correspond to boundaries between regions where the texture was not measured.

Figure 9. Feature map calculated by closing the image and linear mapping of the value.

Comparing results of manual grading of regions by an ophthalmologist with obtained feature maps we have found a good correspondence between large, non–perfused regions. The correspondence decreases for small regions due to the fact that texture is a property which can be measured more reliably in large regions. This means that textural features calculated for a small region have to be treated with lower confidence than for a large one. Small regions could be merged into bigger ones before calculating their textural features.

7. Conclusions and Further Research

There is a good conformation between the results of our method of tessellating images and the boundaries defined by ophthalmologists. There is one major parameter which controls the process of detecting regions: the size of structuring element used for opening. This parameter has a natural interpretation for it corresponds to the minimum diameter of vessel used for delineating regions and thus indirectly determines the number and size of identified regions.

The features which are used for grading the images are sensitive enough to detect non–perfused regions. The usefulness of these features for classifying regions into different grades and their correlation with manual grading needs further work. We also want to use the presence and position of microaneurysms and local dilation of vessels to modify the classification of regions derived from their texture since both these features accompany non–perfused regions.

8. Acknowledgements

This research has been supported by a grant from the Kershaw Trust. We are grateful to Prof D. McLeod from Dept. of Ophthalmology, University of Manchester for proving us retinal images and for helpful advice.

9. References

Akita 82 Akita K., Kuga H.: *Pattern Recognition of Blood Vessel Networks in Ocular Fundus Images.* Proc. SPIE vol. 375, pp. 436–441.

Chaudri et al. 89 Chaudri S. et al.: *Detection of Blood Vessels in Retinal Images Using Two–Dimensional Matched Filters.* IEEE Trans on Medical Imaging, vol. 8, no. 3, pp. 263–269.

Goetcherian 80 Gotecherian V.: *From Binary to Grey Tone Image Processing Using Fuzzy Logic Concepts.* Pattern Recognition, vol. 12, pp. 7–15.

Goldberg et al. 89 Goldberg E.G. et al.: *Quantification of Progressive Diabetic Macular Non–perfusion.* Ophthalmic Surgery, vol. 20, no. 1, pp. 42–45.

Jagoe et al. 90 Jagoe J.R. et al: *Quantification of Retinal Damage During Cardiopulmonary Bypass: Comparison of Computer and Human Assessment.* IEE Proc., vol. 137, pt. I, no. 3, June, pp. 170–175.

Katz et al. 88 Katz N., Goldbaum M., Nelson M., Chaudri S.: *An Image Processing System for Automatic Retina Diagnosis.* Proc. of SPIE vol. 902, pp. 131–137.

Kohner 86 Kohner E.M., Sleightholm M.: *Does Microanerysm Count Reflect Severity of Early Diabetic Retinopathy.* Ophthalmology, vol. 93, May, pp. 586–589.

Lay, Badouin 84 Lay J., Badouin B.: *Computer Analysis of Angioflourograms.* Proc. of VII Int. Conf. on Pattern Recognition, pp. 927–929.

Mandelbrot 82 Mandelbrot B.: *The Fractal Geometry of Nature.* Freeman.

Maragos 89 Maragos P.: *Pattern Spectrum and Multiscale Shape Representation.* IEEE Trans. on PAMI, vol. 11, no. 7, July, pp. 701–715.

Peleg et al. 84 Peleg et al.: *Multiple Resolution Texture Analysis and Classification.* IEEE Trans. on PAMI, vol. 6, no. 4, pp. 518–523.

Rigaut 88 Rigaut J.P.: *Automated Image Segmentation by Mathematical Morphology and Fractal Geometry.* J. of Microscopy, vol. 150, pp. 21–30.

Serra 82 Serra J.: *Image Analysis and Mathematical Morphology.* Academic Press.

Takahashi 84 Takahashi N., et al.: *Distribution of Capillary Nonperfusion in Early–stage Diabetic Retinopathy.* Ophthalmology, vol. 91, December, pp. 1431–1438.

Tanaka 80 Tanaka M., Tanaka K.: *An Automatic Technique for Fundus Photograph Mosaic and Vascular Net Reconstruction.* Proc. of MEDINFO 80, Tokyo, pp. 116–120.

Werman, Peleg 85 Werman M., Peleg S.: *Min–Max Operations in Texture Analysis.* IEEE Trans. on PAMI, vol. 7, no. 6, pp. 730–733.