Development of retinal blood vessel segmentation methodology using wavelet transforms for assessment of diabetic retinopathy

D. J. Cornforth [1], H. J. Jelinek [1], J. J. G. Leandro [2], J. V. B. Soares [2], R. M .Cesar, Jr.[2], M. J. Cree [3], P. Mitchell [4], and T. Bossomaier [1].

[1] *Charles Sturt University, Australia*
Email: {dcornforth, hjelinek, tbossomaier }@csu.edu.au

[2] *Computer Science, University of Sao Paulo, Brazil*
Email: {jleandro, joao, cesar}@vision.ime.usp.br

[3] *Physics and Electronic Engineering, University of Waikato, New Zealand*
Email: cree@phys.waikato.ac.nz

[4] *University of Sydney Department of Ophthalmology, Westmead Hospital, Australia*
Email: paul_mitchell@wmi.usyd.edu.au

Abstract

Automated image processing has the potential to assist in the early detection of diabetes, by detecting changes in blood vessel diameter and patterns in the retina. This paper describes the development of segmentation methodology in the processing of retinal blood vessel images obtained using non-mydriatic colour photography. The methods used include wavelet analysis, supervised classifier probabilities and adaptive threshold procedures, as well as morphology-based techniques. We show highly accurate identification of blood vessels for the purpose of studying changes in the vessel network that can be utilized for detecting blood vessel diameter changes associated with the pathophysiology of diabetes. In conjunction with suitable feature extraction and automated classification methods, our segmentation method could form the basis of a quick and accurate test for diabetic retinopathy, which would have huge benefits in terms of improved access to screening people for risk or presence of diabetes.

1. Introduction

Diabetes affects almost one million Australians, and has associated complications such as vision loss, heart failure and stroke. Any improvement in early diagnosis would therefore represent a significant gain with respect to reducing the morbidity and mortality of the Australian population. This work describes one step along the path to providing automated diagnostic support tools for the early detection of diabetes associated retinal pathology. We focus on the blood vessel network in the retina that is affected by diabetes. Currently, the correct assessment of these images using fluorescein requires a specialist, which presents difficulties in remote or rural areas and is also a relatively invasive and dangerous procedure.

The alternative is to use non-mydriatic colour images that although more difficult to interpret, are more desirable for use in remote or rural areas as they can be obtained by trained rural health professionals.

Non-mydriatic camera images were analysed in an attempt to fine tune and improve the quality of retinal blood vessel detection in medical imaging. For this purpose we used the Morlet wavelet as a tool for segmenting retinal blood vessels in combination with adaptive thresholding, supervised classification and supervised classifier probabilities combined with adaptive thresholding. The advantage of wavelet analysis is its multiscale analysing capability in tuning to specific frequencies, allowing noise filtering and blood vessel enhancement in a single step. Twenty non-mydriatic camera images of the retinal fundus were analysed, from the Stare database (http://www.parl.clemson.edu/stare). The results were compared against the corresponding gold-standard images indicating the true location of vessels as determined by two ophthalmologists and presented in the manner of free-response receiver operator characteristic (FROC) curves.

2. Image processing for medical diagnosis

Identification of vascular anomalies associated with diabetes health care represents a large portion of the assessment carried out by ophthalmologists, which is time consuming and in many cases does not show any anomalies at the initial visit. Utilizing non-specialist health workers in identifying diabetic eye disease is an alternative but trials have shown that correct identification of retinal pathology may be poor (i.e. only 50% of the cases). This success rate decreases for early proliferative retinopathy stages. Telemedicine is an attractive approach. However, this procedure is not time effective and does not lessen the burden on a comparatively small number of ophthalmologists in rural areas that need to assess the images. In addition significant technical problems lessen the availability of telemedicine (Yogesan et al., 2000).

Automated assessment of blood vessel patterns that can be used by rural health professionals is now being extended from fluorescein-labelled to non-mydriatic camera images (Cesar & Jelinek 2003; McQuellin et al., 2002). This work presents the evolution of retinal blood vessel segmentation as a function of the effectiveness of the wavelet transform. We outline the use of the wavelet transform combined with mathematical morphology, supervised training algorithms and adaptive thresholding.

Several methods for segmenting blood vessels using either rule-based or supervised methods have recently been reported for both fluorescein and non-mydriatic colour retinal images (Leandro et al., 2003; Sinthanayothin et al., 1999; Staal et al., 2004). Mathematical morphology, which is a rule-based method, has previously revealed itself as a very useful digital image processing technique for detecting and counting microaneurysms in fluorescein and non-mydriatic camera images (Cree et al., 2004; Spencer et al., 1996). Wavelet transform theory has grown rapidly since the seminal work by Morlet and Grossman, finding applications in several realms (Goupillaud et al., 1984). The wavelets space-scale analysis capability can be used to “decompose” vessel structures into differently scaled Morlet wavelets, so as to segment them from the retinal fundus. Supervised classification, which requires manually labelled images for training vessel pixel candidates, can then be added to the pattern recognition process (Gardner et al., 1996; Leandro et al., 2003; Sinthanayothin et al., 1999; Staal et al., 2004).

3. Methods

The purpose of these experiments was to assess the relative merits of several techniques for segmentation of blood vessels from retinal images. Twenty digital images were used from the Stare database. This database also includes the opinions of two experts who had indicated the position of the vessels from colour images to establish two "gold standards" as separate images.

Our strategy was to use three methods for segmenting retinal blood vessels from directly digitized colour retinal images. The experimental procedure followed was to pre-process the images first to optimise the use of the wavelet transforms. The methods tested were:

1. Wavelet transform plus adaptive thresholding,
2. Wavelet transform plus supervised classifiers,
3. Wavelet transform plus pixel probabilities combined with adaptive thresholding.

In addition, we compared two training techniques: training on one or more complete images, then classifying the remaining images, and training on a window of the image then classifying the remainder of the same image.

Initially the methods were compared qualitatively, but the best of these methods were selected and compared numerically by plotting on a graph of true positive against false positive results from the classification. This graph resembles a free receiver operator curve (FROC) to aid the reader in its interpretation. True positives occur when the classifier labels a pixel as belonging to a vessel and the "gold standard" segmentation also labels the pixel as vessel.

3.1 Pre-processing

In order to reduce the noise effects associated with the processing, the input image was pre-processed by a mean filter of size 5x5 pixels. Due to the circular shape of the non-mydriatic image boundary, neither the pixels outside the region-of-interest nor its boundary were considered, in order to avoid boundary effects. For our wavelet analysis we used the green channel of the RGB components of the colour image as it displayed the best vessels/background contrast.

3.2 Continuous wavelet transform plus adaptive thresholding

Applying the continuous wavelet transform approach provides several benefits but resulted in some loss of detail as the scale parameter was fixed. We therefore adopted a pixel thresholding approach that represented each pixel by a feature vector including colour information, measurements at different scales taken from the continuous wavelet (Morlet) transform and the Gaussian Gradient, as well as from mean filtering applied to the green channel. The resulting feature space was used to provide an adaptive local threshold to assign each pixel as either a vessel-pixel or a non-vessel pixel.

The real plane $R \times R$ is denoted as R^2, and vectors are represented as bold letters, e.g. $\mathbf{x}, \mathbf{b} \in R^2$. Let $f \in L^2$ be an image represented as a square integrable (i.e. finite energy) function defined over R^2 (Arnéodo et al., 2000). The *continuous wavelet transform* (**CWT**) is defined as:

$$T_{\psi}(\mathbf{b},\theta,a)(\mathbf{x}) = C_{\psi}^{-1/2}\frac{1}{a}\int \psi^{*}\left(a^{-1}r_{-\theta}(\mathbf{x}-\mathbf{b})\right)f(\mathbf{x})d^{2}\mathbf{x} \qquad (1)$$

where C_{ψ}, ψ, $\mathbf{b}$, $r_{-\theta}$, θ and a denote the normalizing constant, the analysing wavelet, the displacement vector, the rotation operator, the rotation angle and the dilation parameter, respectively (ψ^{*} denotes the complex conjugate). The double integral is taken over R^{2} with respect to vector variable **x,** being denoted by $d^{2}\mathbf{x}$. The Morlet wavelet is directional (in the sense of being effective in selecting orientations) and capable of fine tuning specific frequencies. These latter capabilities are especially important in filtering out the background noise, and comprise the advantages of the Morlet wavelet with respect to other standard filters such as the Gaussian and its derivatives. The 2D Morlet wavelet is defined as:

$$\psi_{M}(\mathbf{x}) = \exp(j\,\mathbf{k}_{0}\cdot\mathbf{x})\exp\left(-\frac{1}{2}|A\mathbf{x}|^{2}\right) \qquad (2)$$

where $j=\sqrt{-1}$ and $A = diag\left[\varepsilon^{-1/2}, 1\right]$, $\varepsilon \geq 1$ is a 2×2 array that defines the anisotropy of the filter, i.e. its elongation in some direction. In the Morlet equation (2), which is a complex exponential multiplying a 2D Gaussian, $\mathbf{k}_{0}$ is a vector that defines the frequency of the complex exponential. Using the Morlet transform to segment the blood vessels, the scale parameter is held constant and the transform is calculated for a set of orientations θ = 0, 10, 20, 30, ..., 180. The ε parameter has been set as 4 in order to make the filter elongated and $\mathbf{k}_{0} = [0, 2]$, i.e. a low frequency complex exponential with few significant oscillations. The transform maximum response (in modulus) from all orientations for each position, **b**, is then taken, emphasizing the blood vessels and filtering out most of the noise. The blood vessels can then be detected from this representation.

3.3 Feature extraction

The pixel feature space was formed by Morlet wavelet responses (taken at different scales and elongations), Gaussian Gradient responses (taken at different scales) and colour information, which determine each pixel's colour. This resulted in a computationally demanding high dimensional feature space. At the same time, Morlet responses taken at close scales are highly correlated, as are the Gaussian Gradient responses for similar scales. Therefore we used a feature extraction approach to obtain a lower dimensional feature space, while trying to preserve structure important for discrimination. Feature extraction was performed by a linear mapping provided by nonparametric discriminant analysis (Fukunaga, 1990). Nonparametric discriminant analysis consists of building two matrices. The first is a nonparametric between-class scatter matrix, constructed using k-nearest neighbour techniques, which defines the directions of class separability. The second is the within-class scatter matrix, which shows the scatter of samples around their mean class vectors. These matrices were built based on the labelled training samples. The two matrices are then used to find a projection (given by a linear mapping) that maximizes class separability while minimizing the within-class scatter in the projected feature space.

During the adaptive thresholding process, the dimensional nature of the features forming the feature space might give rise to errors. Since the feature space elements may be considered as random variables, we applied a normal transformation in order to obtain a new

relative random variable, redefined in a dimensionless manner. The normal transformation is defined as:

$$\hat{X}_j = \frac{X_j - \mu_j}{\sigma_j} \qquad (3)$$

where X_j is the *jth* feature assumed by each pixel, μ_j is the average value of the *jth* feature and σ_j is the associated standard deviation.

3.4 Supervised classification

In methods 2 and 3, supervised classification was applied to obtain the final segmentation, with the pixel classes defined as C1 = vessel-pixels and C2 = non-vessel pixels using the Bayesian classifier consisting of a mixture of Gaussians (Theodoridis, 1999). In order to obtain the training set, retinal fundus images have been manually segmented, thus allowing the creation of a labelled training set into 2 classes C1 and C2 (i.e. vessels and non-vessels). In this work, the hand-drawn vascular tree provided by the ophthalmologist has been used – our training pattern – so that we obtained its feature space. Two different strategies for deriving the training set were applied:

1. Some images were completely segmented by an expert and a random sub-set of their pixels was used to train the classifier.
2. Only a small portion (window) of a sample image was manually segmented. The labelled pixels are then used to train the classifier, which is applied to the same image in order to conclude its segmentation.

This 2nd strategy has been devised so that a semi-automated fundus segmentation software can be developed, in which the operator only has to draw a small portion of the vessels over the input image or simply click on several pixels associated with the vessels. The remaining image is then segmented based on this partial training set without the need of tuning any additional parameters. This approach requires a small effort from the operator, which is compensated for by the fact that image peculiarities (e.g. due to camera model and settings) are directly incorporated by the classifier. Note that this method should be repeated for every new image.

3.5 Post-processing

The output produced by the classifier leads to a binary image where each pixel is labelled as vessel or non-vessel. Some misclassified pixels appeared as undesirable noise in the classified image. In addition, for some vessels, only their boundaries were classified, so that it was necessary to perform post-processing by using morphological tools to obtain the final desired segmentation. Finally, to optimize the vessel contours, morphological operations have been applied, beginning by area open to eliminate small noisy components. The vessels were completely filled by morphological dilation and area close (Cesar & Jelinek, 2003).

3.6 Comparison

The results were obtained in the form of segmented images and compared by the experts. Here we present the results in the manner of FROC curves to assess the trade off between true-positive fraction and false-positive fraction of our methods in detecting the pixels associated with the vessel patterns.

4. Results

Before we present the results of the experiments comparing the different applications of the segmentation procedures, we provide for comparison an example of the application of the wavelet transform to fluorescein images (Leandro et al., 2001). Figure 1(a) shows a typical image of the retinal fundus with the optic disk on the right hand side and the protruding blood vessels that course throughout the image. Figure 1 (b) shows the result of image segmentation using the Morlet wavelet transform. The latter shows the difficulty in obtaining a clear segmentation. Background noise and variable grey levels across the image introduce artefacts. In particular, this method did not remove the optic disc located on the right hand side of the image and was very susceptible to hue variation that resulted in the large grey area to the left of the image.

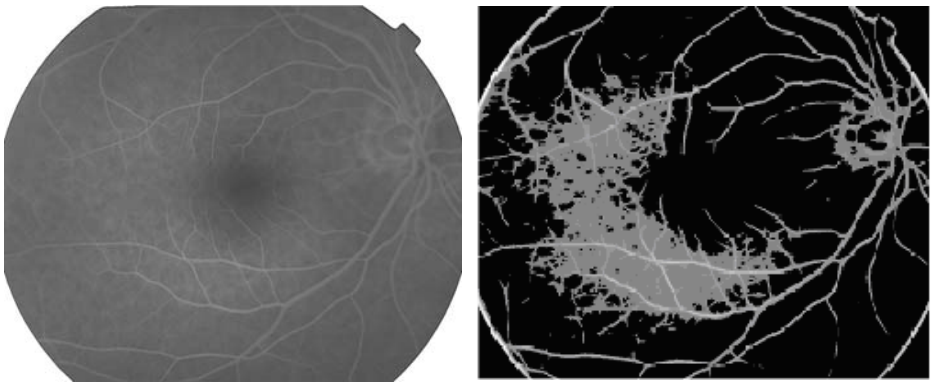

(a) *(b)*

Figure 1. Wavelet transform in blood vessel segmentation: a) original fluorescein retinal image; b) and example of segmentation of blood vessels from a fluorescein image using the Morlet wavelet transform only.

In method 1 we applied the wavelet transform plus adaptive thresholding to colour non-mydriatic camera images. Figure 2 (a) shows a typical grey scale representation of a colour image obtained from the digital camera. The optic disk is noticeable as a light grey area on the left hand side with blood vessels emanating from it. Notice the variable brightness across the image, and especially the presence of the optic disc, which can introduce artefacts during the image processing. Figure 2 (b) shows the same image after application of the Morlet wavelet transform and thresholding. This is much more successful than using the fluorescein monochrome image (Figure 1). The optic disc has been successfully removed, but artefacts remain. In particular, notice the extra “vessels” apparent at the bottom of Figure 2 (b) at approximately 5 o’clock. Many disconnected segments also remain, and some smaller vessels clearly visible in (a) have not been detected in (b).

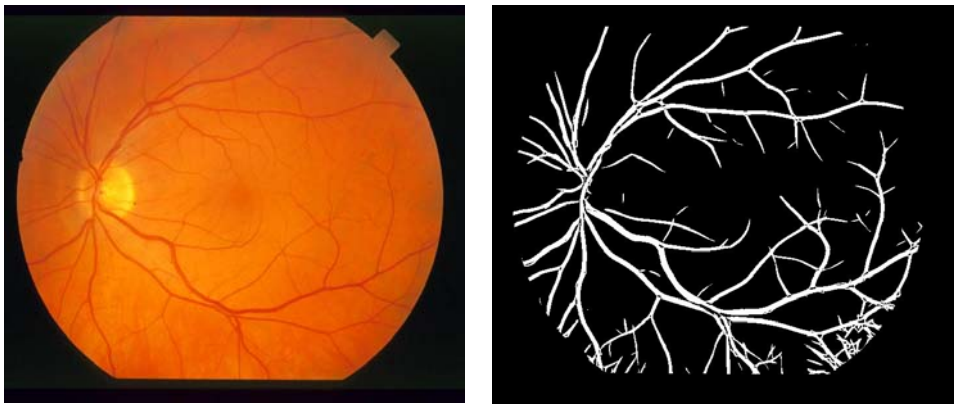

(a) *(b)*

Figure 2. Segmentation of non-mydriatic colour images for method 1: a) grey-scale image of original retinal fundus; b) segmentation of retinal blood vessels using wavelet transform and adaptive thresholding.

For a more sophisticated approach to dealing with the image variations in hue of background and blood vessels, we applied a supervised learning algorithm. The classifier was first trained using all pixels from entire images. All pixels were labelled by the experts, as shown in Figure 3(a). The trained classifier was then used to segment other images. In Figure 3 (b) we show the result of supervised classification, where the classifier has been trained on four other images, and then used to segment the image of Figure 2 (a). Comparing this with Figure 2 (b) the improvement is obvious. Many of the artefacts at the bottom (5 o'clock) of that image have now disappeared. However, many of the smaller vessels towards the centre of the image have not been detected, and there are still many disconnected vessel segments.

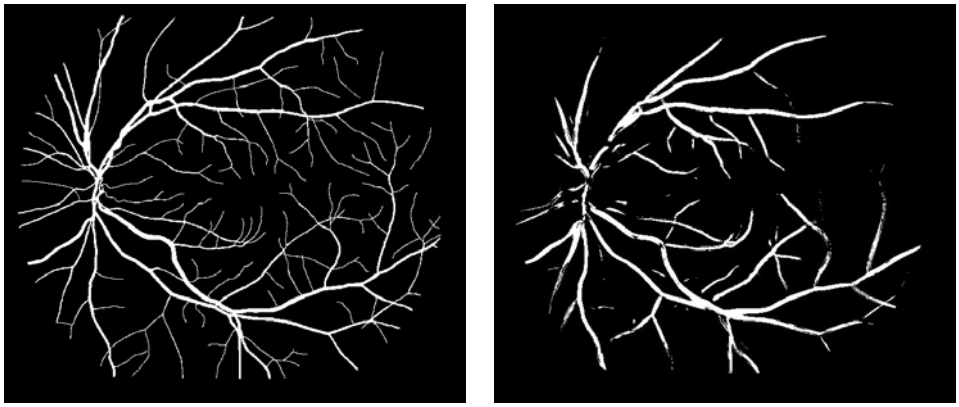

(a) *(b)*

Figure 3. Results from method 2, wavelet transform plus pixel classification: (a) an example of the training set provided by the experts; (b) an example of a segmented image obtained using the total vessel pattern as a training set.

For method 3, we combined the wavelet transform with the supervised classification and mixed adaptive thresholding. In this case, instead of using the simplified approach of Leandro *et al.* the thresholding procedure was applied to the pixel probability of being vessel as estimated by the supervised classifier approach (Leandro et al., 2001). This led to the results shown in Figure 4. Here many of the smaller vessels are now visible, and the number of disconnected vessel segments is much less.

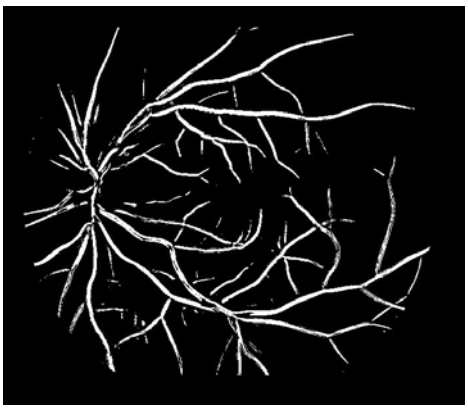

Figure 4. The same image after adaptive thresholding on the probability of each pixel being part of a vessel.

A variation of the pixel classification is to train the classifier with a window of the image, then use it to segment the remainder of the image. This should provide more accurate classification, as it corrects for different image parameters. Figure 5 (a) shows the window containing the training data for method 2. This represents a portion of the retinal vessels as identified by the expert. Figure 5 (b) shows the result of the segmentation when only using a part of the figure as a training set. The number of small vessels detected has increased, and the segmentation is of superior quality. Compare this with figure 3.

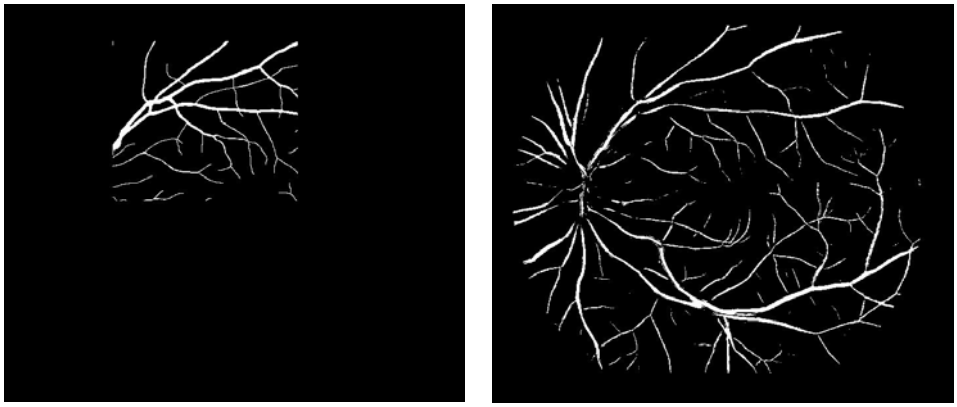

(a) *(b)*

Figure 5. Segmented image (b) obtained using only a portion of the vessel pattern as a training set (a).

Finally, we applied the adaptive thresholding (method 3) to the vessel probability of each pixel of the window based classification. A typical result is shown in Figure 6. This represents the best result obtained so far, where most of the smaller vessels have been detected. The main problem with this approach is that it does not take the probability of being background into account. We hope to address this problem in our future work.

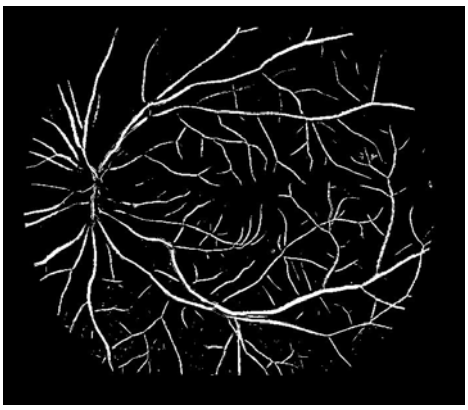

Figure 6. Typical result of using the window method to train the classifier, followed by an adaptive thresholding process.

It is clear from these results that methods 2 and 3, each using the supervised classifier approach, provide the best results. We now present quantitative results from these two methods in Figure 7. For method 2, (wavelet transform and supervised classifier) each source image resulted in a single binary output image with pixels either marked as 'true' (a vessel pixel) or marked 'false' (not a vessel pixel). Each output image produced the single point on the graph in figure 7. The average from the 20 images processed is shown as a circle marked "Adaptive threshold" in the legend.

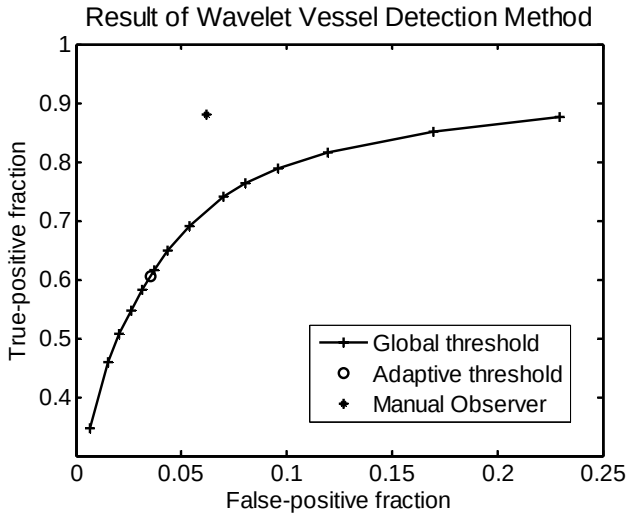

Figure 7: Numerical results from methods 2 and 3, shown in the same form as an ROC graph. Method 2 uses an adaptive threshold, so a single point was obtained, being the average of the values obtained from the 20 images. Method 3 uses a global threshold, so many points were generated as the threshold was varied. Each point is the average of the results from the 20 images. The grey cross indicates the evaluation of one ophthalmologist.

Method 3 (wavelet transform and adaptive threshold) resulted in 20 grey-scale images, where the brighter the pixel the more likely it landed in the vessel class. A global threshold was applied to each image to generate a point on the graph. The threshold was varied from high (poor sensitivity) to low (good sensitivity but too many false positives). The average values taken from the 20 images produced a number of points tracing out the curve (appearing in the legend as "Global threshold".

Discussion

We have demonstrated some new techniques for the automated processing of non-mydriatic images in the study of diabetic retinopathy that can certainly be extended to other contexts in pattern recognition. The results we have obtained so far suggest that pixel classification, in conjunction with wavelet transform and adaptive thresholding, can provide noise-robust vessel segmentation. The approach reported here improved on previous results by reducing the level of interaction required with the segmentation program, providing a useful tool for non-specialists such as community health workers in assessing fundus complications associated with diabetes (Antoine et al., 1997; Cesar & Jelinek, 2003; da Costa, 2001; Grossmann, 1988; Zana & Klein, 2000).

Wavelets are especially suitable for detecting singularities (e.g. edges) in signals, extracting instantaneous frequencies, and performing fractal and multifractal analysis (Antoine et al., 1997; Grossmann, 1988). Applying the wavelet transform allows noise filtering and blood vessel enhancement in a single step. Our results indicate that for the same false-positive fraction, the supervised learning with adaptive thresholding obtained a greater than 75% sensitivity compared to the ophthalmologist with approximately 90%. Future research is now focused on fine tuning our algorithms on a larger data set.

Although these methods are targeted at segmentation in retinal blood vessels, there is no reason why they may not be applied in other areas, especially in medical imaging, where it is necessary to extract intricate branching patterns from images with a noisy background.

Acknowledgements

RMC and JS are grateful to FAPESP (Research Support Foundation of the State of São Paulo, Brazil) and to CNPq (Brazil's National Council for Scientific and Technological Development). HJ was in receipt of grants from CSU and Australian Diabetes Association. The authors also wish to acknowledge the contribution of Alan Luckie and Tien Wong for their expert advice on diabetic retinopathy and arteriolar narrowing.

References

Antoine J. P., Barache D., Cesar Jr. R. M. & da Costa L. (1997), Shape characterization with the wavelet transform, Signal Processing 62(3): 265-290.

Arnéodo, A., Decoster, N. & Roux S. G. (2000), A wavelet-based method for multifractal image analysis. I. Methodology and test applications on isotropic and anisotropic random rough surfaces, The European Physical Journal B 15: 567-600.

Cesar Jr. R. M. & Jelinek H. F. (2003), Segmentation of retinal fundus vasculature in nonmydriatic camera images using wavelets, In Suri J. S. & Laxminarayan S., (Eds.), *Angiography and plaque imaging*, CRC Press, London: 193-224.

Cree M., Luckie M., Jelinek H. F., Cesar R., Leandro J., McQuellin C. & Mitchell P. (2004), Identification and follow-up of diabetic retinopathy in rural health in Australia: an automated screening model, In AVRO, Fort Lauderdale, USA 5245/B5569.

da Costa L. F. (2001), On neural shape and function, In World Congress on Neuroinformatics: ARGESIM / ASIM - Verlag Vienna: 397-411.

Fukunaga K. (1990), Introduction to statistical pattern recognition, 2nd ed., Academic Press, Boston.

Gardner G. G., Keating D, Williamson T. H., & Elliot A.T. (1996), Automatic detection of diabetic retinopathy using an artificial neural network: a screening tool, British Journal of Ophthalmology 80: 940-944.

Goupillaud P., Grossmann, A. & Morlet J. (1984), Cycle-octave and related transform in seismic signal analysis, Geoexploration 23: 85-102.

Grossmann A. (1988), Wavelet Transforms and Edge Detection, In Albeverio, S., et al., (Eds.), Stochastic Processes in Physics and Engineering, Reidel Publishing Company, Dordrecht.

Leandro J. J. G., Cesar R. M. Jr., Jelinek H. F. (2001), Blood vessels segmentation in retina: preliminary assessment of the mathematical morphology and of the wavelet transform techniques, In SIBGRAPI-01, Florianópolis - SC, IEEE Computer Society Press: 84-90.

Leandro J. J. G., Soares J. V. B., Cesar R. M. Jr., & Jelinek H. F. (2003), Blood vessel segmentation of non-mydriatic images using wavelets and statistical classifiers, In Proceedings of the Brazilian Conference on Computer Graphics, Image Processing and Vision (Sibgrapi03), Sao Paulo, Brazil, IEEE Computer Society Press: 262-269.

McQuellin C. P., Jelinek H. F. & Joss G. (2002), Characterisation of fluorescein angiograms of retinal fundus using mathematical morphology: a pilot study, In 5th International Conference on Ophthalmic Photography, Adelaide: 83.

Sinthanayothin C., Boyce J. & Williamson C. T. (1999), Automated localisation of the optic disc, fovea and retinal blood vessels from digital colour fundus images, British Journal of Ophthalmology 83(8): 902-912.

Spencer T., Olson J. A., McHardy K., Sharp P. F. & Forrester J. V. (1996), An Image-Processing Strategy for the Segmentation and Quantification of Microaneurysms in Fluorescein Angiograms of the Ocular Fundus, Comput. Biomed. Res. 29: 284-302.

Staal J. J., Abramoff M. D., Niemeijer M., Viergever M. A. & van Ginneken B. (2004), Ridge-based vessel segmentation in color images of the retina, IEEE Trans Med Imag 23(4): 501-509.

Theodoridis S. (1999), Pattern Recognition, Academic Press, Baltimore.

Yogesan K., Constable I. J., Barry C. J., Eikelboom R. H. & Tay-Kearney M. L. (2000), Telemedicine screening of diabetic retinopathy using a hand-held fundus camera, Telemedicine Journal 6(2): 219-223.

Zana F. & Klein J. C. (2000), Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation, IEEE Transactions on Image Processing 10(7): 1010-1019.