Enhancement and Segmentation of Retinal Images

Dr. T. Morris, Mrs. Z. Newell

Department of Computation, UMIST,
P.O. Box 88, Manchester, M60 1QD, U.K.
t.morris@co.umist.ac.uk

Abstract

Glaucoma is one of the major causes of preventable blindness in the world. It induces nerve damage to the optic nerve head via increased pressure in the ocular fluid. It is presently detected either by regular inspection of the retina, measurement of the intra ocular pressure (IOP) or by a loss of vision. It has been observed that nerve damage precedes the latter two events and that direct observation of the nerve head could therefore be a better method of detecting glaucoma, if the observations could be made reliably. This paper describes our work in enhancing and isolating the optic nerve head in images of the retina. We describe previous attempts that have been made using simple image processing techniques and the current multiresolution approaches we are taking, we present samples of our initial results. Once the nerve head has been located, its shape will be quantified using measurements that have already been shown to be effective.

1. INTRODUCTION

Glaucoma is the largest cause of preventable blindness in the U.K., accounting for some 13% of the patients registered blind [1]. Nerve damage of the retinal ganglion cells which pass through the optic nerve head (that region of the retina where nerve fibres and blood vessels pass through the eye, the "blind spot") is induced by increased pressure of the aqueous humour which affects the circulation of the associated capillary network. At present it is detected either by regular inspection of the optic nerve head, measurement of the intra ocular pressure (IOP) or by a loss of peripheral visual field. Irreversible loss of sight may have occurred by the time glaucoma is diagnosed. Treatment of glaucoma is monitored by the same tests, with the same problem, i.e. the tests respond indirectly to the disease and further loss of sight may have occured by the time ineffectual treatment is recognised.

It has been observed that changes in the optic nerve head topography precede changes in the IOP and loss of vision [2] and that direct observation of the optic nerve head could therefore be the best method of detecting glaucoma, if the observations could be quantified reliably. Manual assessment of the shape of the nerve head has been shown to be unreliable, automation could improve this [3].

Automatically identifying the neuroretinal rim (the boundary between the retina and the nerve head) is not an easy task. The images are taken under low light conditions and are noisy. (The data used in the present investigation was collected by attaching a video camera to the eyepiece of a Zeiss fundus camera. The retina was illuminated with a photographic flash and the video image grabbed. Cox and Wood [4] describe the instrumentation in more detail.) The structure itself is indistinct and partially obscured by blood vessels. In fact, if a simple edge detector is applied to these images it is the blood vessel boundaries that give the strongest responses, the required features respond very weakly, if at all. Figure 1 shows a normal retina, portions of the required boundaries are clearly visible, other portions are obscured or very indistinct. The original image is 512 by 512 pixels and 8 bit greyscale.

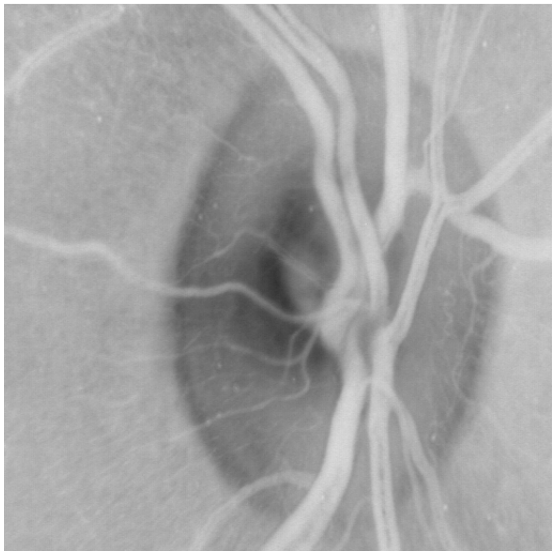

Figure 1.
Retinal image, white light illumination.

A substantial amount of work has been reported on attempts at automatically detecting the neuroretinal rim, though none has met with significant success. Early reported work simply thresholded the image, either globally [5] or adaptively [6], but met with little success due to uneven illumination. Cox and Wood [4] presented a semi-automated method: an observer indicated extremal points on the boundary which were automatically connected by tracing along the boundary. They showed how important it was for the same observer to perform all of the measurements since the inter observer variability was similar to the difference between normal and abnormal classes [3]. Morris, Cox and Wood [7] initially presented a completely automatic method which traced beteween points on the boundary identified automatically by their grey level gradient properties. They have latterly returned to semi automated methods which are proving to be more reliable [8]. Lee and Brady [9] and Donnison and Morris [10] have both investigated using active contour (snake) methods to locate the boundary. Both sets of authors have highlighted the importance of pre-processing the images to emphasise the difference between retinal and optic disk regions of the image before searching for the boundary. Donnison and Morris appear to have had more success in their implementation, possibly due to their formulation of the active contour.

Multiresolution methods have long been seen as an attractive method of segmenting images. From a theoretical viewpoint, they mimic some behaviours of the human visual system, pragmatically they allow us to make early decisions as to the approximate locations of image features and to focus attention on those areas for further processing. A number of approaches have been followed. Image pyramids [e.g. 11] are generated by progressively reducing the size of an image from the base (the original full resolution image) to the top (a single pixel whose value is the average grey value of the image). They seem to have been used most often in edge detection applications. Scale spaces [e.g. 12] may be generated by applying a feature detector at varying resolutions to the original image; a volume is generated: two axes coincide with the axes of the original image, the third represents scale. As the scale is varied the size of the detected feature changes. Marr's edge detector [13] could be used, the scale parameter would equate to the σ of this operator. The common theme in all of these approaches is that information extracted at a lower resolution is used to guide the extraction of information at a higher resolution. In our work we are using the wavelet transform [14, 15] as a means of generating the hierarchical representation of the image data and thus guide its segmentation.

The wavelet transform generates an image which can be divided into four quadrants. Three represent horizontally and vertically oriented and corner features at some scale. The fourth quadrant repeats this basic structure at half of this scale. In our implementation, up to five repetitions can be made. It is our intention to use this structure to guide the interpretation of the original image.

We have taken two approaches to the segmentation. In the first we have simply wavelet transformed the data, filtered and reverse transformed. The nerve head has been enhanced by modifying the wavelet coefficients. In the second approach, we have radially resampled the image: the circular nerve head becomes a linear boundary which is enhanced by a one dimensional transform.

2. IMAGE FILTERING

Setting the coefficients corresponding to small scale features to zero will remove those features and blur the edges of larger scale features. This effectively enhances the gross image features, but it does not allow their boundaries to be accurately delimited. Instead, we choose to suppress the small scale features' coefficients and enhance the larger scale coefficients (cf the non-ideal filters). Figure 2 shows a retinal image which has been transformed using the Daubechies wavelet, the first three stages of the transformed data have been removed, the next two have been suppressed to 30% and 90% of their original values whilst the remaining gross data has been enhanced by 10%; the data has then been inverse transformed.

The filter coefficients were chosen by a process of trial and error, based on a subjective assessment of the resulting image. Smaller blood vessels have been successfully removed. It is not possible to remove the larger ones without also removing the similarly sized nerve head boundary. We are currently developing more rigorous methods of chosing the filter coefficients which take into account the features' spatial extents.

Whilst this technique is currently unable to segment the image it is a reliable method of enhancing the differences between these two classes of pixels, and thus provides a useful

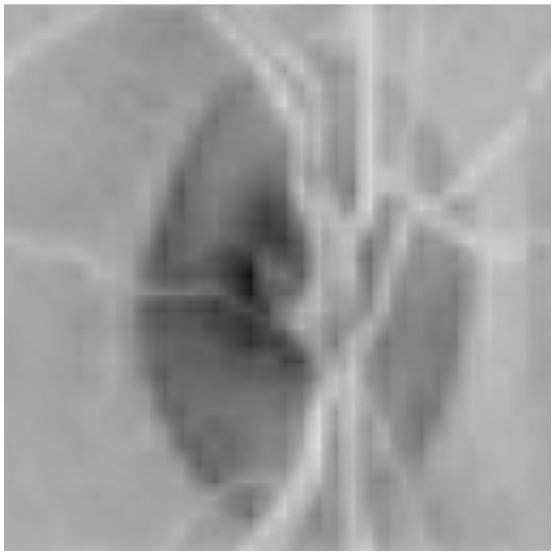

Figure 2.
Figure 1, transformed, filtered and reverse transformed.

preprocessing stage for further analyses. More importantly, this step allows us to suggest regions of the image that are definitely nerve head and regions that are definitely retina, and thus focus attention on the remaining regions.

3. IMAGE RESAMPLING AND FILTERING

When imaged, the optic nerve head appears as an approximately elliptical object and it is usually processed as one. Processing may be simplified by resampling the image in a radial sense: this transforms the nerve head boundary into a locally linear structure.

Resampling is achieved by defining an origin for the resampling vector, a maximum number of samples to be estimated per vector, an initial direction for the vector, an angular increment between vectors and the sense in which the angle increases.

We require the origin to be inside the optic nerve head; the exact location is unimportant. The image's centroid is a point that satisifes this requirement. Image values are interpolated at 256 unit increments along each resampling vector. If necessary the vector is padded with zeros. The vector is initially in the nine o'clock direction, its orientation is incremented in a clockwise sense by an amount sufficient to yield 512 vectors. A typical resampled image is shown in figure 3. The centroid is along the top edge of the image which is 512 vectors long. The first column represents the vector towards the image's left hand side and is 256 samples long.

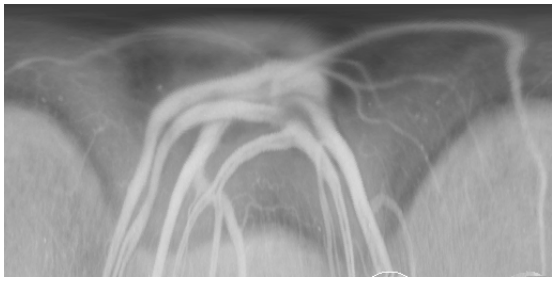

Figure 3.
Radially resampled image.

The columns of the resampled image are passed through a one dimensional wavelet transform, filtered and reverse transformed. A typical result is shown in figure 4. The different stages of the transformed data have been given the same weightings as before (0.0, 0.0, 0.0, 0.3, 0.9, 1.1).

Figure 4.
Processed version of figure 3.

We are currently experimenting to determine the optimum values of the weightings, and also investigating methods of segmenting these images that utilise the structure of the boundary.

4. SEGMENTATION

Segmentation is a trivial problem once the image has been satisfactorily enhanced. Our earlier work [4, 7, 8, 10] has demonstrated that the optic nerve head boundary can be followed reliably provided that points on the boundary can be identified (i.e. a point to initialise the following) and there is sufficient contrast across the boundary. We are confident that images enhanced using these methods may be accurately segmented.

5. CONCLUSIONS

Automated assessment of the optic nerve head may provide information that is useful in the diagnosis and treatment of glaucoma. We have described earlier work in this area and presented the results of our recent work in enhancing the optic nerve head. This involves selective modification of the wavelet transform coefficients.

We are currently investigating methods of reliably, and automatically selecting the weightings. We are also investigating the efficacy of varying the coefficients within a stage of the transform, this will take advantage of the transform's spatial localisation and may improve the discrimination between the larger vessels and the nerve head boundary.

We have also presented early results of enhancing a radially resampled version of the images. We are currently investigating this idea.

We contend that completely segmenting the image is simple provided that the enhancement stage has performed well. The final goal of our research is to characterise the shape of the nerve head, which will be done using features previously shown to be effective [3]. Finally, we shall validate our characterisation using a set of images from patients with known diagnoses.

References

[1] F.G. Brown, R. Fletcher, *Glaucoma In Optometric Practice*, Blackwell Scientific Publishers, Oxford, 1990.

[2] T.G. Zeyen, J. Caprioli, Progression Of Disk And Field Damage In Early Glaucoma, *Arch Ophthalmol* 111:62-65, 1993.

[3] M.J. Cox, I.C.J. Wood, Inter- And Intra-Image Variability In Computer Assissted Optic Nerve Head Assessment. *Opthal. Physiol. Opt.* 11:36-43, 1991.

[4] M.J. Cox, I.C.J. Wood. Computer-Assisted Optic Nerve Head Assessment. *Ophthal. Physiol. Opt.* 11:27-35, 1991.

[5] A.R. Rosenthal, D.G. Falconer, P. Barrett. Digital Measurement Of Pallor-Disc Ratio. *Arch. Ophthalmol.* 98:2027-2031, 1980.

[6] P. Nagin, B. Schwartz, K. Nanba, The Reproducibility Of Computerised Boundary Analysis For Measuring Optic Disc Pallor In The Normal Optic Disc. *Ophthalmology* 92:243-251, 1985.

[7] D.T. Morris, M.J. Cox, I.C.J. Wood. Automated Extraction Of The Optic Nerve Head Rim, *Americal Academy of Optometrists Annual Conference*, Boston, Dec. 1993.

[8] T. Morris, I. Wood, The Automatic Extraction Of The Optic Nerve Head *American Academy of Optometrists Biennial European Conference*, Amsterdam, May 1994.

[9] S. Lee, J.M. Brady, Integrating Stereo And Photometric Stereo To Monitor The Development Of Glaucoma. *Proceeding of the British Machine Vision Conference*, September 1990.

[10] C. Donnison, T. Morris, Identifying The Neuroretinal Rim Using Dynamic Contours, *Image and Vision Computing* (in review).

[11] A. Rosenfeld, Pyramids, Multiresolution Image Analysis. *Proc Third Scandinavian Conference on Image Analysis*, Copenhagen, July 1983.

[12] T. Lindberg. Scale-Space Theory: A Basic Tool For Analysing Structures At Different Scales. *J of Applied Statistics*, 21(2):224-270, 1994.

[13] D. Marr, E. Hildreth. Theory Of Edge Detection. *Proc Roy Soc* B207:187-217, 1980.

[14] S.G. Mallat. A Theory For Multiresolution Signal Decomposition: The Wavelet Representation. *IEEE Pat. Anal. Mach. Intell,* 11:674-693, 1989.

[15] I. Daubechies. *Ten Lectures On Wavelets*. CBMS 61, SIAM Press, Philadelphia PA, 1992.