

Characterization of changes in blood vessel width and tortuosity in retinopathy of prematurity using image analysis

Conor Heneghan^{a,*}, John Flynn^{a,b}, Michael O’Keefe^c, Mark Cahill^{c,d}

^a*Department of Electronic and Electrical Engineering, University College Dublin, Belfield, Dublin 4, Ireland*

^b*Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA*

^c*Department of Ophthalmology, The Children’s Hospital, Temple Street, Dublin 1, Ireland*

^d*Beetham Eye Institute, Joslin Diabetes Center, 1 Joslin Place, Boston MA 02215, USA*

Received 23 March 2001; received in revised form 1 August 2001; accepted 22 October 2001

Abstract

Many retinal diseases are characterised by changes to retinal vessels. For example, a common condition associated with retinopathy of prematurity (ROP) is so-called plus disease, characterised by increased vascular dilation and tortuosity. This paper presents a general technique for segmenting out vascular structures in retinal images, and characterising the segmented blood vessels. The segmentation technique consists of several steps. Morphological preprocessing is used to emphasise linear structures such as vessels. A second derivative operator is used to further emphasise thin vascular structures, and is followed by a final morphological filtering stage. Thresholding of this image is used to provide a segmented vascular mask. Skeletonisation of this mask allows identification of points in the image where vessels cross (bifurcations and crossing points) and allows the width and tortuosity of vessel segments to be calculated. The accuracy of the segmentation stage is quite dependent on the parameters used, particularly at the thresholding stage. However, reliable measurements of vessel width and tortuosity were shown using test images. Using these tools, a set of images drawn from 23 subjects being screened for the presence of threshold ROP disease is considered. Of these subjects, 11 subsequently required treatment for ROP, 9 had no evidence of ROP, and 3 had spontaneously regressed ROP. The average vessel width and tortuosity for the treated subjects was 96.8 μm and 1.125. The corresponding figures for the non-treated cohort were 86.4 μm and 1.097. These differences were statistically significant at the 99% and 95% significance level, respectively. Subjects who progressed to threshold disease during the course of screening showed an average increase in vessel width of 9.6 μm and in tortuosity of +0.008. Only the change in width was statistically significant. Applying a simple retrospective screening paradigm based solely on vessel width and tortuosity yields a screening test with a sensitivity and specificity of 82% and 75%. Factors confounding a more accurate test include poor image quality, inaccuracies in vessel segmentation, inaccuracies in measurement of vessel width and tortuosity, and limitations inherent in screening based solely on examination of the posterior pole.

© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Morphological processing; Retinopathy of prematurity; Segmentation; Screening

1. Introduction

Pathological changes of the retinal vasculature are a feature of many diseases. For example, diabetic re-

tinopathy is often characterised by the presence of new blood vessels, venous beading, microaneurysms, and intraretinal macular abnormalities. Another example is the presence of plus disease, concurrent with retinopathy of prematurity, which is characterised by an increase in vessel width and tortuosity. Currently, these systematic changes in vessel characteristics are determined qualitatively by direct inspection (ophthalmoscopy) or through examination

*Corresponding author. Tel.: +353-1-706-1925; fax: +353-1-283-0921.

E-mail address: conor.heneghan@ucd.ie (C. Heneghan).

of photographic records of the retina. Relatively little work has been carried out to date on automated quantitative analysis of digital retinal imagery. This is probably for two reasons: (a) until recently, digital images of the retina could only be obtained through scanning of conventional film photography or slides, which made image processing relatively cumbersome, and (b) the tasks required of automated analysis are quite demanding from an image processing point of view.

In this paper, we present a technique for automatically segmenting out the retinal vasculature from an image of the retina, and for using this segmentation to provide measurements of vessel widths and tortuosity. The clinical motivation for undertaking this study was to quantitatively assess a set of clinical imagery obtained from subjects who were being screened for retinopathy of prematurity (ROP). In ROP, normal retinal vessel development is halted and replaced by abnormal blood vessels. It is a serious complication of premature birth with an incidence of between 16 and 56% in infants weighing 1500 g or less or born at a gestational age of 30 weeks or less (McNamara and Connolly, 1999; Palmer et al., 1991), and can lead to severe long-term vision loss or blindness if not treated correctly.

An international classification of ROP has been developed to standardize evaluation and assist in clinical research into the disease (Committee for the Classification of Retinopathy, 1984). The classification divides the eye into a set of zones (Fig. 1) and the severity of ROP is determined according to three parameters: (a) the zone in which new vessels are located, (b) how much of the retina is involved (which is determined by dividing the retinal area into clock hours) and (c) how much fibrosis is associated with the vessels. A fourth independent parameter is the presence or absence of 'plus disease'. Plus disease is when normal blood vessels located near the optic disc (which is near the posterior pole of the retina) become

dilated and tortuous. All the parameters are cumulative and if they are all present in sufficient severity an eye is said to have threshold disease. Other inflammatory changes accompany these vasculature changes and threshold disease is associated with the imminent development of severe, sight-threatening ROP. Fortunately, the majority of cases of ROP regress spontaneously, but once threshold disease has been detected treatment is indicated. The mainstay of treatment for ROP is ablation of the immature retina that is producing the abnormal growth factors that stimulate the new blood vessels, using either cryotherapy (Cryotherapy for Retinopathy of Prematurity Co-operative Group, 1988, 1990a,b) or laser treatment (McNamara et al., 1992).

Screening premature infants for threshold disease makes sense from a public health perspective. ROP is a significant health problem, its natural history is known, the at-risk population can be identified, the screening test is sensitive and specific, and an economic treatment that is effective in altering the course of the disease is available. International guidelines have been established for screening. A key issue is to avoid screening before the disease is normally present, which would result in the false impression that there is no ROP, but not delaying too long relative to the natural history of the disease, in case any ROP may be too advanced to respond to treatment. In general, the first examination takes place between 4 and 6 weeks of chronological age or between 31 and 33 weeks post-conceptual age. Subsequent examinations depend on the presence and severity of any ROP detected. Subjects with no ROP may have only one further examination 6 weeks later, whereas subjects with significant disease may require weekly or even daily examinations, to assess if they have reached threshold disease. Typically these examinations are carried out with an indirect ophthalmoscope, a device that allows the ophthalmologist to get a stereoscopic wide-angle view of the retina, but which is quite challenging to use correctly.

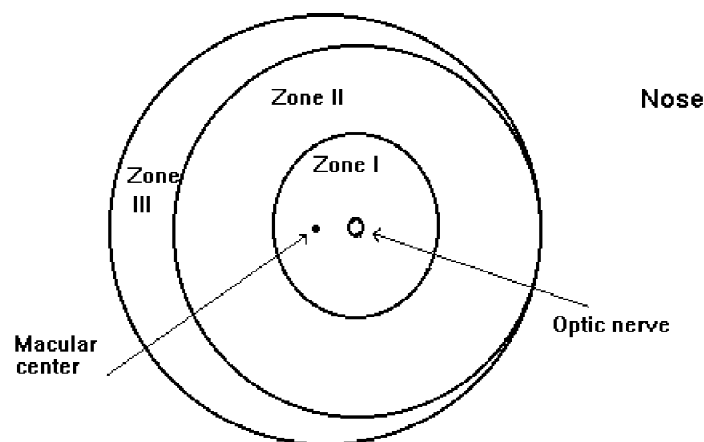


Fig. 1. Guideline diagram used in the International Classification of Retinopathy of Prematurity.

The main goal of this screening process is to reliably identify subjects who have progressed to threshold disease, so they can be promptly treated. At present this screening process is carried out by ophthalmologists skilled in the examination of infants' eyes. Any system which can assist ophthalmologists in increasing the accuracy of their screening, or which could allow less highly trained individuals to carry out the screening (e.g., ophthalmic nurses) may be of clinical benefit. A possibility of providing some automated assistance in this screening process lies in accurate computer measurement of vessel width and tortuosity near the posterior pole (back) of the retina. The clinical justification for this is that it has been recently demonstrated that the absence of dilated and tortuous vessels in the posterior pole is a reliable marker for the absence of threshold ROP (Saunders et al., 2000; Wallace et al., 2000), lessening the need for indirect ophthalmoscopy of the peripheral retina. The posterior pole can be more easily visualised using a direct ophthalmoscope (which is considerably easier to use), or a fundus camera. Therefore analysis of the region near the posterior pole can be used as a screening test in its own right, since only subjects with changes in this region will exhibit threshold disease. This opens the possibility of screening for ROP by non-ophthalmologists using a direct ophthalmoscope (Saunders et al., 2000), or by automated techniques which provide quantitative measurements of vessel width and tortuosity in the posterior pole.

This paper proposes a technique for measuring vessel width and tortuosity in the posterior pole of the retina, and shows how it can be used to demonstrate statistically significant changes in these parameters for ROP subjects. The technique firstly implements segmentation of the retinal vasculature. This is carried out using morphological pre-processing based on linear structuring elements, followed by enhancement using a smoothed second-derivative operator. A final stage of morphological post-processing is then used, prior to thresholding. This results in a binary mask representing detected vessels from which vessel widths can be calculated. This binary mask is then skeletonised. The tortuosity of segments of this skeleton (corresponding to specific blood vessels) can be calculated. By considering the average vessel width and tortuosity at the posterior pole, it can be shown that (a) subjects with threshold disease have wider and more tortuous vessels, though only the changes in vessel width can be considered to be statistically significant based on our measurement technique, (b) that subjects who progress to threshold disease subsequent to their first examination experience a statistically significant increase in average vessel width, and (c) that treatment leads to a rapid reduction in average vessel width and tortuosity. A rudimentary automated screening system for threshold disease based solely on measurement of vessel width and tortuosity is proposed, and the sensitivity and specificity of this test are shown to be 82% and 75%, respectively, for the test data set.

2. Mathematical preliminaries

2.1. Notation

Much of the processing carried out in this paper is based on morphological image processing. Excellent reviews of this topic can be found in the books by Serra or Soille (Serra, 1982; Soille, 1999). For completeness, we briefly define the notation used in this paper for morphological processes.

A binary image f is a mapping of a subset D_f of Z^n , called the definition domain of f , into the couple $\{0, 1\}$ (for 2D images, n is 2),

$$f: D_f \subset Z^n \rightarrow \{0, 1\}. \quad (2.1)$$

A greyscale image f is a mapping of a subset D_f of Z^n , called the definition domain of f , into a bounded set of nonnegative integers N_0 ,

$$f: D_f \subset Z^n \rightarrow \{0, 1, \dots, t_{\max}\}, \quad (2.2)$$

where t_{\max} is the maximum value of the data type used, (e.g., 255 in a greyscale 8-bit image). The *intersection* of two greyscale images is defined as

$$f \wedge g = \min[f(x), g(x)], \quad (2.3)$$

where $f \wedge g$ denotes the greyscale intersection. Similarly, $f \vee g$, the greyscale union is defined,

$$f \vee g = \max[f(x), g(x)]. \quad (2.4)$$

The (greyscale) erosion of a set X by a structuring element (SE) B is denoted by $\varepsilon_B(X)$ and can be defined as the minimum value of the image in the neighbourhood defined by the structuring element when its origin is at x ,

$$[\varepsilon_B(f)](x) = \min_{b \in B} f(x + b). \quad (2.5)$$

The (greyscale) dilation of a set X by a structuring element B is denoted by $\delta_B(X)$ and can be defined as the maximum value of the image in the neighbourhood defined by the structuring element when its origin is at x ,

$$[\delta_B(f)](x) = \max_{b \in B} f(x + b). \quad (2.6)$$

The opening $\gamma_B(f)$ of an image f by a structuring element B is defined as the erosion of f by B followed by the dilation with the transposed SE \hat{B} .

$$\gamma_B(f) = \delta_{\hat{B}}[\varepsilon_B(f)]. \quad (2.7)$$

The closing $\phi_B(f)$ of an image f by a structuring element B is defined as the dilation of f by B followed by the erosion with the transposed SE \hat{B} .

$$\phi_B(f) = \varepsilon_{\hat{B}}[\delta_B(f)]. \quad (2.8)$$

We also make use of geodesic dilation and erosions in this work. Letting f denote the marker image and g the

mask image, with both definition domains equal and f less than or equal to g at every point, the geodesic dilation of size 1 is defined as

$$\delta_g^{(1)}(f) = \delta(f) \wedge g, \quad (2.9)$$

where the dilation is with respect to the 3×3 unit mask. The geodesic dilation of arbitrary size n is defined recursively as

$$\delta_g^{(n)}(f) = \delta_g^{(1)}[\delta_g^{(n-1)}(f)]. \quad (2.10)$$

Similarly, the *geodesic erosion* of size 1 is defined as

$$\varepsilon_g^{(1)}(f) = \varepsilon(f) \vee g, \quad (2.11)$$

where the erosion is with respect to the 3×3 unit mask, and in this case f is greater than or equal to g at every point. The geodesic erosion of arbitrary size n can be defined recursively as

$$\varepsilon_g^{(n)}(f) = \varepsilon_g^{(1)}[\varepsilon_g^{(n-1)}(f)]. \quad (2.12)$$

The *reconstruction by dilation* of a mask image g from a marker image f is defined as the geodesic dilation of f with respect to g until stability and is denoted by $R_g(f)$,

$$R_g(f) = \delta_g^{(i)}(f), \quad (2.13)$$

where i is such that $\delta_g^{(i)}(f) = \delta_g^{(i+1)}(f)$. The *reconstruction by erosion* is defined similarly,

$$R_g^*(f) = \varepsilon_g^{(i)}(f), \quad (2.14)$$

where i is such that $\varepsilon_g^{(i)}(f) = \varepsilon_g^{(i+1)}(f)$.

The hit-or-miss transformation, *HMT*, of a binary image X by a composite structuring element $B = (B_f, B_b)$ is the set of points x , such that when the origin of B coincides with x , B_f fits X and B_b fits X^c ,

$$HMT_B(X) = \{x | (B_f)_x \subseteq X, (B_b)_x \subseteq X^c\}. \quad (2.15)$$

The thin transformation, *THIN*, of a binary image X by a composite structuring element B is defined as the set difference between X and the hit-or-miss transform of X by B ,

$$THIN_B(X) = X \setminus HMT_B(X). \quad (2.16)$$

Hit-or-miss transforms and thinnings can selectively preserve or remove pixels that have certain neighbourhood configurations.

The *skeleton by thinning* of an image X with Golay letter L is defined as the repeated thinning of X with all rotations of the letter L until idempotence,

$$SKEL_L(X) = THIN_{L=1 \dots n}(X)^\infty. \quad (2.17)$$

Here the ∞ symbol denotes iteration until idempotence, and n denotes n rotations of L ($n = 8$ for the 8-connected grid).

3. Extraction and characterization of vessel widths and tortuosity

The work presented here focuses on quantifying the thickness and tortuosity of the retinal vasculature being screened for sight-threatening ROP. A first stage is to segment the vasculature from other anatomical structures in the image such as the optic disk, by creating a binary mask image of the complete vasculature. This output mask image marks pixels in the original picture that are part of vessels as white (binary 1) with all other pixels as black (binary 0). Accurate vasculature segmentation is fundamentally important, as further analysis of vessel properties depends on the accuracy of this segmentation.

In developing this algorithm, we initially considered two types of retinal images – black and white images obtained using fluorescein angiography, and colour images obtained using a wide-angle retinal camera under white light, though our later results will focus on colour retinal images only. In fluorescein angiography, a contrast enhancement dye (fluorescein) is injected into the blood stream, which fluoresces under ultraviolet light. Fluorescein angiograms typically have higher contrast than colour fundus images, and segmentation of vessels in fluorescein angiograms using our system is more successful than in the colour images we present. However, fluorescein angiography is not used in screening for ROP, so in the rest of this paper we restrict our comments to colour fundus images with relatively poor contrast.

3.1. Characteristics of the vasculature

The vasculature has a number of characteristics that can be used in developing a segmentation technique:

- The vessel cross-sectional grey level profile approximates a Gaussian shape
- The vasculature is piecewise linear, and can be represented by many connected line segments
- The direction and grey level of a vessel do not change abruptly, they are continuous
- The vasculature is tree-like; all vessels are connected to all other vessels and they all originate from within a single area, the optic disk

Some factors that hinder vascular segmentation are:

- Vessels are obviously not all the same size, shape or colour
- The contrast can sometimes be quite low; and the vessel colour can be close to that of the background
- Some background features (e.g., underlying choroidal structures, or the nerve fiber layer) have similar attributes to vessels
- Vessel crossings and bifurcations may confuse some techniques
- The edge of the optic disk can be wrongly segmented as a vessel

3.2. Existing work in vasculature segmentation

Previous research has considered the problem of extracting vasculature from retinal and angiographic images. A brief survey of methods used is presented here.

3.2.1. Matched filtering

Matched filtering has been used to emphasise vessels in a fundus image (Zhou et al., 1994; Chaudhuri et al., 1989; Kochner et al., 1998). Since vessels are Gaussian in cross section, a convolution kernel with a Gaussian cross section is usually used. This method is quite effective to a point, but rarely useful alone, and is slow as the convolution kernel used may be quite large, and needs to be applied at several rotations. In addition, the kernel responds optimally to vessels that have the same σ , the standard deviation of the underlying Gaussian function, as the kernel, and thus may not respond to thin vessels as well as wide vessels, etc. In addition, the kernel, since it has a minimum length, may give weak responses for very tortuous vessels. The proposed automated algorithm of Chapman et al. (2001) fits into the category of matched filtering, but only identifies single vessels at a time with manual intervention.

3.2.2. Morphological methods

Morphological image processing is particularly effective at extracting image features whose shape is known a priori. The vasculature, known to be piecewise-linear, is just such an image feature, and is thus suited to morphological processing. Morphology has been applied to vasculature segmentation (Zana and Klein, 1997, 1999; Matsopoulos et al., 1999) with considerable success and can also be used in micro-aneurysm extraction. Morphology has the advantage of speed and noise resistance. A potential disadvantage is that morphological methods do not exploit the known vessel cross-sectional shape.

3.2.3. Artificial neural networks

Artificial neural networks have been applied with promising results to ophthalmic image analysis (Gardner et al., 1996). They are useful in that they can be trained to recognise patterns, such as vessels, in the images, and can thus extract the vasculature. They can also extract other features such as micro-aneurysms. Image processing using neural networks requires, and indeed provides, little insight into the underlying phenomena in the image. This can be both a disadvantage and an advantage to neural networks, in that they make classifications, healthy, unhealthy, etc., based on statistical probabilities, and not objective reasoning. Another disadvantage is that they require exposure to training data or a 'gold standard'. This gold standard consists of a number of images whose vascular structure must be precisely marked by an ophthalmologist.

3.2.4. Vessel tracking algorithms

Vessel tracking algorithms (Sun, 1989) segment a vessel

between two points. They represent a departure from the methods outlined above, in that they work at the level of a single vessel rather than the entire vasculature. Usually, the algorithm works by stepping along the vessel, measuring its width along the route. Various properties of the vessel including average width and tortuosity can also be measured en route. The main advantage of vessel tracking methods is that they provide highly accurate vessel widths, and can provide information about individual vessels that is usually unavailable using other methods. Unfortunately, they require the starting point, and usually the end, of a vessel to be defined by a user and are thus, without additional techniques, of limited use in fully automated analysis.

3.2.5. Manual techniques

Several researchers have considered the fractal topology of the retinal vasculature (Daxer, 1993; Mainster, 1990). In order to make accurate measurements of fractal dimension, a highly reliable segmentation algorithm was required, so these authors used manual tracing of the digital image using a digitising tablet. An advantage of this method is the ability to use clinical knowledge to differentially segment arteries and veins. However, this method is obviously not suited to automated analysis.

3.3. Proposed approach to segmentation

The method proposed in this paper uses a morphological filter to emphasise linear structures such as vessels. The differential properties of the vessels provide additional emphasis. An additional morphological filter, and hysteresis thresholding is used to obtain the final binary image which represents the mask of the vessels. The algorithm is based on the work of Zana and Klein (1997, 1999).

To make the discussion more concrete, we will provide specific values of the parameters used to analyze 640×480 pixel colour images, taken with a 120° RetCam® fundus camera. In practice, the software we have developed automatically sets these parameters based on the size of the image opened, and we also allow the facility to easily change the default values used.

3.3.1. Initial morphological filtering

Morphological image processing is particularly good at extracting image features whose shape is known a priori. In particular, morphology provides an attractive solution for detecting quasi-linear shapes such as vessels in an image. Since the vasculature is known to be piecewise linear, an algorithm that robustly extracts linear shapes is useful in vasculature extraction.

The aim of the initial morphological filtering operation is to emphasise the vasculature, preserving vessel crossings and bifurcations. The approach used in this project uses the difference between the supremum and infimum of the

openings of the original image $I_0(x, y)$ with two linear structuring elements of different lengths.

Consider an image $I_0(x, y)$ containing many bright linear shapes (the vessels) as well as flat homogeneous areas such as the background. A linear shape is defined as a bright part of an image with a minimum length L and a maximum width W (where usually $W \ll L$). The vasculature in a retinal image is composed of many such connected linear shapes, and the aim of our initial processing is to preserve image structures which satisfy the criteria of being at least L pixels long, and no more than W pixels wide. Morphological opening with a structuring element of a given shape preserves image structures that can contain the structuring element and removes those that cannot. Thus, opening the image with a linear structuring element, B , of length L_1 and width 1 preserves linear shapes when the structuring element and the shape are approximately parallel (or more accurately, when $L_1 \sin(\alpha) < W$, where α is the angle between the single-pixel structuring element and the linear shape of interest). If many such structuring elements are used, with different angular rotations, then all linear shapes with length greater than or equal to L_1 should

be preserved by at least one rotation. Noise and other non-vessel structures that cannot contain the structuring element at any rotation will not be preserved by such an operation. Thus, a cleaner version I_c of the image I_0 , can be obtained by taking the supremum of the openings of the image with linear structuring elements with many different rotations (Fig. 2). Mathematically this can be expressed as

$$I_c = \sup_{i=1, \dots, 12} \{\gamma_{B_i}(I_0)\}, \quad (3.1)$$

where B_i represents structuring element B at rotation i , and twelve rotations of the single pixel structuring elements each 15° apart are used. In our particular case, the value of L_1 used was 17 pixels. The length of the structuring element should be chosen so as to preserve vessels but remove noise and non-vessel structures. While the length of the optimal structuring element for a given retinal image depends on the tortuosity of the vessels in the image and the amount of noise present, the overall algorithm is relatively insensitive to small changes in the value of L_1 . A more detailed discussion on the sensitivity of the results to parameter values is presented in Section 4.1.

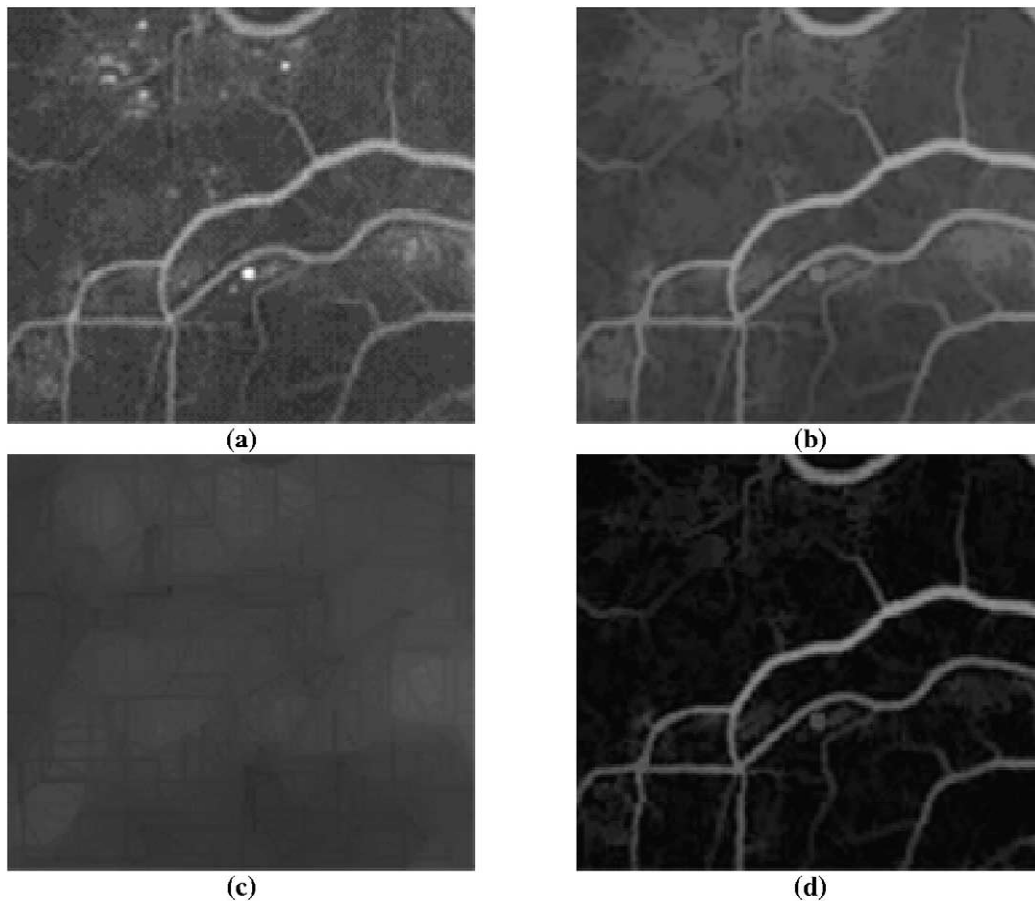


Fig. 2. Illustration of the effects of the initial morphological filtering step. (a) A section (160×160 pixels) of the original fundus image I_0 . This image is a cropped section of a fluorescein angiogram. (b) The corresponding section of I_c obtained by application of Eq. (3.2). Note that bright spots have been removed. (c) The corresponding section of I_b obtained by application of Eq. (3.3). This can be thought of as the fundus background image. (d) The corresponding section of I_v obtained by application of Eq. (3.4).

Much of the small detail (small or tortuous vessels, etc.) that is lost in the operation of Eq. (3.1) can be recovered if morphological reconstruction is used. Morphological reconstruction extracts the peaks in a mask (conditioning) image that are touched by a marker image. If the image produced by Eq. (3.1) is used as the marker image and the original image is used as the mask image then an improved version of Eq. (3.1), incorporating reconstruction, is

$$I_c = R_{I_0}(\sup_{i=1, \dots, 12} \{\gamma_{B_i}(I_0)\}). \quad (3.2)$$

This operation provides a significantly cleaner version of the image but is not sufficient to allow final segmentation of the vasculature. The operation of Eq. (3.2) preserves large homogeneous regions whose size exceeds L_1 pixels in at least one dimension, and is thus of little use alone. This is because such regions can clearly contain the structuring element, (in at least one direction) and are thus preserved by Eq. (3.2). A method is required that preserves only linear shapes which do not exceed a desired width.

Opening the image with a single-pixel linear structuring element, B , of length L_1 , preserves linear shapes. A corollary to this is that opening the image with a linear structuring element, B , of length L_2 and width 1 removes linear shapes when the structuring element and the shape are approximately orthogonal (or more strictly when $L_2 \sin(\alpha) > W$ where α is the angle between the single-pixel structuring element and the linear shape of interest, and W is the width of that linear shape). Removing the linear shapes from an image corresponds to replacing them by their (local) background. Thus for a linear shape of width less than L_2 there is at least one direction in which opening with a linear structuring element B of length L_2 will remove the shape. Thus, the local background image can be obtained by taking the infimum of openings with a linear structuring element B of length L_2 taken in many directions. Mathematically,

$$I_B = \inf_{i=1, \dots, 12} \{\gamma_{B_i}(I_0)\}, \quad (3.3)$$

where twelve structuring elements at angles from 0° to 165° are used. The value of L_2 used in our processing is 17 pixels. This operation preserves large homogenous areas of colour but also removes small linear shapes. There are now two images; I_c , that contains both the linear shapes and the large homogenous areas and another image, I_B , that contains only the large homogenous areas. Subtracting these images will yield a third image, I_v , which contains only the linear shapes.

$$I_v = I_c - I_B. \quad (3.4)$$

To summarise, this initial morphological filtering stage yields an image I_v that contains only linear shapes. Fig. 2 gives a typical example of the results of this processing when applied to a fundus image. When the original image is particularly clean, the vessel mask may be obtained by a

simple threshold. However, for most images further processing is required.

3.3.2. Second derivative properties of the vasculature

As noted previously, the vessel cross-sectional greyscale profile has an approximately Gaussian shape. The vessel width can be defined as the distance between $\pm\sigma$ of the underlying Gaussian function. The second derivative of a Gaussian function is negative for $|x| < \sigma$, then positive, before diminishing to zero. [This is under the assumption of vessels which are brighter than their background; sign changes are required for images with an inverted Gaussian profile.] Thus, taking the negative of the second derivative in the cross sectional direction at a point in a vessel will yield a positive value between $-\sigma$ and σ (either side of the vessel centre line) and a negative value just outside the vessel. Hence if the supremum of the negative second derivatives of the image taken in a number of different cross-sectional directions is computed, then the true vessel cross-sectional direction should give the strongest response. Strong responses should be obtained inside the vessels and weak or negative responses just outside. The vessels should be brightened by this operation, while other structures with zero or positive second derivatives in their cross-sectional direction will be darkened. Thinner vessels will be especially brightened, because the magnitude of the second derivative of a Gaussian is proportional to $1/\sigma^4$, and thinner vessels have smaller σ . This is desirable, as thinner vessels usually appear darker than thicker vessels in the original image.

Differential operators are particularly sensitive to noise, and care must be taken in calculating the second derivative. A robust method of computing the second derivative of a 1-D signal is by convolving a Gaussian filter with the signal to smooth it and then calculating the second derivative of this smoothed signal. Using the properties of convolution these two operations can be combined into one convolution; convolution by the second derivative of the Gaussian. This method of finding the negative second derivative works well at emphasising vessels, however it was found that the vessel boundaries become ‘jagged’ unless the method is augmented with additional smoothing. The reason for this ‘jaggedness’ is that the second derivative calculation only takes into account those pixels along the cross-sectional direction. If pixels along the vessel direction are also used in the second derivative calculation, then the jaggedness is considerably reduced.

One method of achieving this is by smoothing along the cross-sectional direction *and* the vessel centre line direction prior to finding the negative second derivative. Since it is assumed that vessels are piecewise linear, smoothing along a vessel’s centre line should not lead to any loss of detail in the image, but should reduce the jaggedness. Mathematically, this step can be summarised,

$$I_{\text{diff}} = \sup_{\theta=0 \dots 180} (I_v * G_\theta * (G''_{\theta+90})), \quad (3.5)$$

where G_θ denotes the one dimensional Gaussian convolution kernel at an angle θ and $G''_{\theta+90}$ denotes the one dimensional second derivative of the Gaussian convolution kernel (the Laplacian-of-Gaussian) at an angle of $(90 + \theta)^\circ$ (the two kernels are orthogonal). Twelve evenly-spaced angles between 0° and 165° were used in evaluating Eq. (3.5). The one-dimensional Gaussian function is given by

$$G_\theta(r) = \frac{1}{\sqrt{2\pi\rho^2}} \exp\left(\frac{-r^2}{2\rho^2}\right), \quad (3.6)$$

and was implemented using a 7×1 pixel convolution mask with $\rho = 1.75$, i.e.,

$$G_\theta(r) = [0.2301 \quad 0.5205 \quad 0.8494 \quad 1.0000 \quad 0.8494 \\ 0.5205 \quad 0.2301].$$

The one-dimensional Laplacian-of-Gaussian operator took the form

$$G''_\theta(r) = \frac{K(r^2 - \rho^2)}{\rho^4} \exp\left(\frac{-r^2}{2\rho^2}\right) \quad (3.7)$$

and was implemented using the following 7×1 pixel convolution mask with $\rho = 1.75$,

$$G''_\theta(r) = [0.6686 \quad 0.3378 \quad -0.5060 \quad -1.0000 \\ -0.5060 \quad 0.3378 \quad 0.6686].$$

The scaling factor K was set equal to 10. The output of Eq. (3.5) results in positive and negative values. All negative values are set to zero, prior to proceeding to the final stage of processing.

3.3.3. Final morphological filtering

The result of the operation described by Eq. (3.5) will be that vessels are brightened and other non-vessel parts are darkened. However, there are also responses at the edge of flat regions, due to rapid brightness changes. In addition, the centre of wide vessels appear dark, compared to vessel edges. This is because the Gaussian approximation is less valid for wider vessels, as their centres tend to be flat, and thus have lower second derivatives and appear dark in the derivative image. However, the bright non-vessel parts tend to be non-linear (not made up of line-segments) and can thus be removed by a morphological filter similar to that used in the initial morphological filtering stage. The background, as well as flat areas, has already been removed by the two previous stages, thus a morphological filter defined by Eq. (3.2) suffices to filter linear structures, and removes the bright non-linear parts of the image. Thus step one of the final morphological filtering operation is

$$I_l = R_{I_{\text{diff}}}(\sup_{i=1, \dots, 12} \{\gamma_{B_i}(I_{\text{diff}})\}). \quad (3.8)$$

The dark areas within vessels are also usually not particularly linear in shape. To remove the dark areas from

within the vessels we apply the dual filter to the one above. The dual filter uses closings instead of openings and takes the infimum rather than the supremum. Also the reconstruction is the *reconstruction by erosion* (denoted by R^*),

$$I_f = R^*_{I_l}(\inf_{i=1, \dots, 12} \{\phi_{B_i}(I_l)\}). \quad (3.9)$$

This filter ‘fills-in’ holes in the background that cannot fit the linear structuring element. This gives the final image, I_f , which is thresholded to give the vascular segmentation. Figs. 3(c) and 3(d) illustrate the effect of the differential filter combined with the final stage of morphological filtering.

3.3.4. Thresholding

The final step to produce the binary output mask of the vasculature is to threshold the image I_f . As stated before, the vasculature is tree-like: every vessel is connected to every other vessel. Some thresholding method that exploits this connectedness is appropriate. *Hysteresis thresholding* is just such a method.

An ordinary threshold transformation sets all pixels in the input image above a certain grey scale level to the value 1 and the remaining ones to the value 0. More generally, in mathematical terms, the threshold operator, T , sets all pixels x of the input image f whose values lie in the range $[t_i, t_j]$ to 1 and all others to 0,

$$[T_{[t_i, t_j]}(f)](x) = \begin{cases} 1, & \text{if } t_i \leq f(x) \leq t_j, \\ 0, & \text{otherwise.} \end{cases} \quad (3.10)$$

Hysteresis thresholding uses two levels, t_{low} and a higher level t_{high} . Any pixel with greyscale value above t_{high} is set to 1, as with ordinary thresholding. However, pixels that have a greyscale value above t_{low} and are connected (either eight-connected or four-connected depending upon application – we use eight-connectedness) to pixels with greyscale values above t_{high} are also set to 1. Isolated pixels above t_{low} are set to 0 (as are all pixels below t_{low}). Thus, thin dark vessels above the t_{low} level that are connected to bright vessels above t_{high} are included in the vasculature mask.

Hysteresis thresholding is easily implemented using morphological processing. Two binary images I_{low} and I_{high} are created; I_{low} by thresholding with the lower threshold value t_{low} , and I_{high} by thresholding with the higher threshold value t_{high} . I_{high} is then reconstructed into I_{low} (I_{high} is the marker image and I_{low} the mask image). Mathematically,

$$HYST_{(t_{\text{low}}, t_{\text{high}})}(I) = R_{T_{[t_{\text{low}}, t_{\text{max}}]}(I)}[T_{[t_{\text{high}}, t_{\text{max}}]}(I)], \quad (3.11)$$

where $T_{[t_1, t_2]}$ is the threshold transformation defined by Eq. (3.8). Thus, the final mask image containing the vasculature, I_m , can be created by hysteresis thresholding I_f with suitable values for t_{low} and t_{high} .

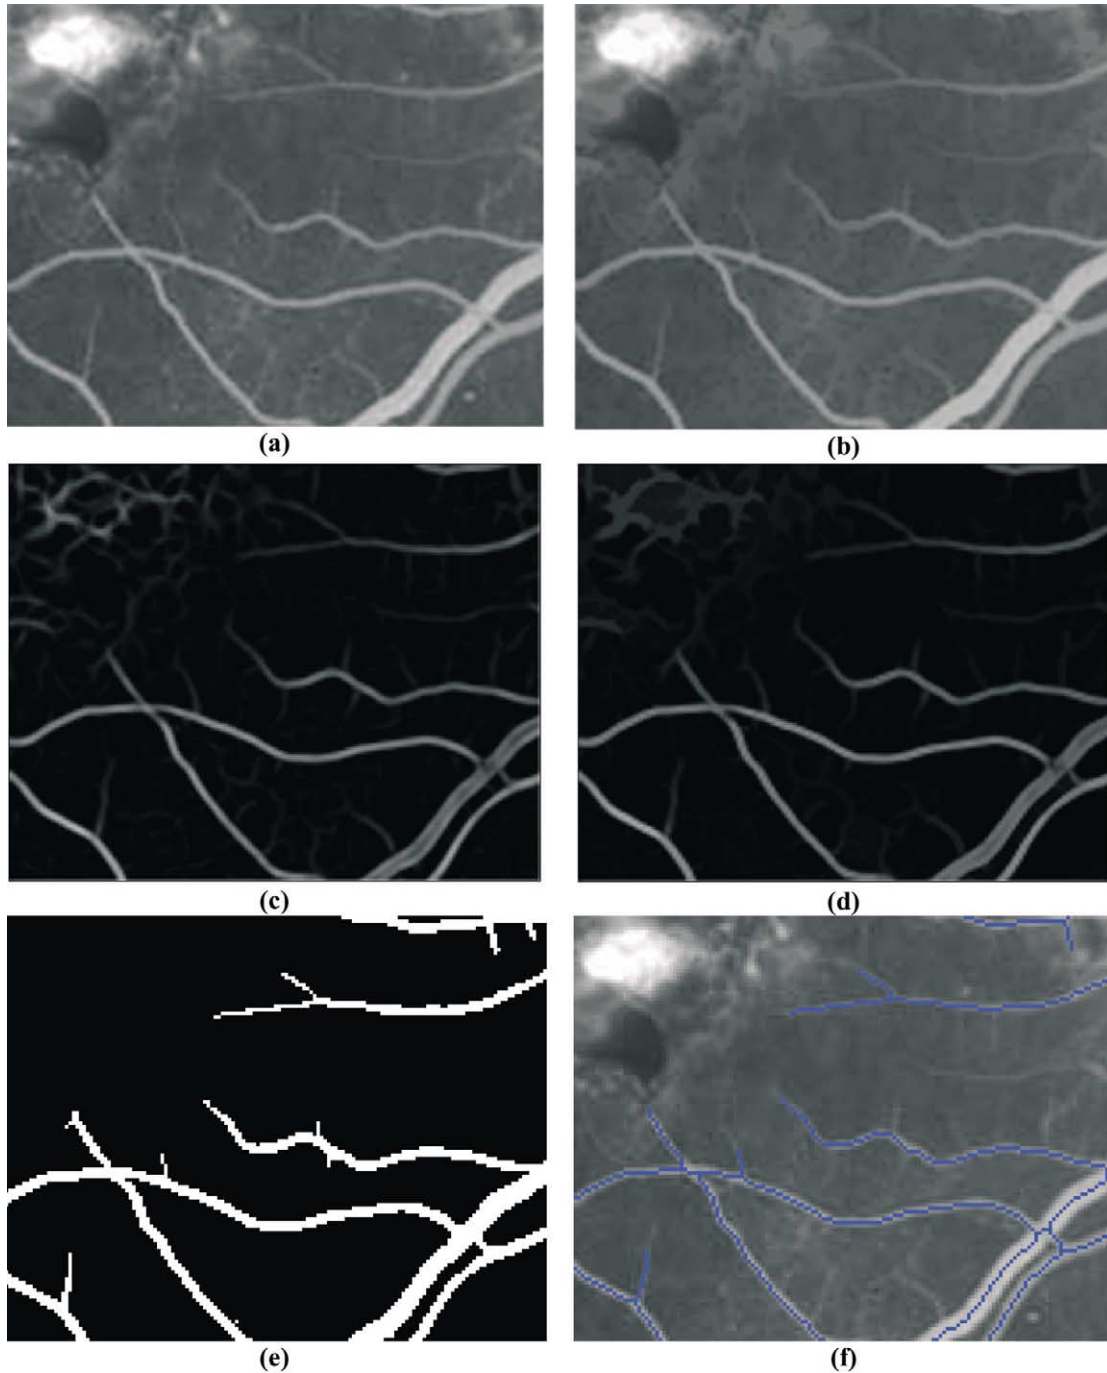


Fig. 3. Illustration of the effects of the differential filter, the final stage of morphological filtering, thresholding, and skeletonization. (a) A section (160×120 pixels) of fundus image. This image is a cropped section of a fluorescein angiogram. (b) The same section after the morphological preprocessing of Fig. 2. (c) Same section of image after application of the differential filtering in Eq. (3.5). (d) Same section of image after application of the morphological processing in Eqs. (3.8) and (3.9). (e) Final binary mask image after application of Eq. (3.12). (f) Skeleton of binary mask image overlaid on the original image.

$$I_m = \text{HYST}_{(t_{\text{low}}, t_{\text{high}})}(I_f). \quad (3.12)$$

Typical values of threshold used range from 25 to 40 (on a scale of 0 to 255). In practice, this thresholding stage has the most significant impact on the accuracy of the segmentation, and at present is carried out interactively. In

Section 4.1, we further consider the sensitivity of our scheme to choice of threshold value. The accuracy of detection can be further improved by removing those connected pixel regions whose area is below a certain value, the minimum grain size (chosen to be 170 pixels in our case). This is implemented by a simple labelling

scheme which searches for labelled clusters with less than 170 pixels, and removes them from the image. The effect of the thresholding stage is pictured in Fig. 3(e).

3.4. Further processing

At this stage, a mask containing the vasculature has been created. This mask can be processed directly to calculate vessel widths directly. However, by reducing the vasculature mask to a symbolic skeleton, processing is simplified and generalised. The skeleton of the vasculature corresponds to the centre lines of the vessels. This skeletal image can be used to detect vessel crossing and bifurcation points as well as to yield information about the vessel widths and tortuosities, which are the immediate focus of this work.

3.4.1. The skeletonisation and skeletal filtering algorithm

To create the skeletal image the skeleton-by-thinnings transformation defined by Eq. (2.17) is used,

$$I_{\text{SKEL}} = \text{SKEL}_L(I_{\text{mask}}), \quad (3.13)$$

where SKEL_L denotes the skeleton by homotopic thinnings using the Golay letter ‘L’ (Sonka et al., 1998).

This skeleton image contains many small branches, corresponding to noise in the vessel widths, etc. Skeletal filtering helps by providing a more regular skeleton structure. We used a *skeletal pruning* operation to achieve this goal. The pruning operation is implemented by thinning the skeletal image with the m rotations of the Golay letter ‘E’,

$$\text{PRUNE}(X) = \text{THIN}_{E=1 \dots m}(X). \quad (3.14)$$

This transformation removes one pixel from the end of each skeletal branch. A pruning of size n removes n pixels from each skeletal branch and is defined by,

$$\text{PRUNE}^{(n)}(X) = (\text{THIN}_{E=1 \dots m}(X))^n. \quad (3.15)$$

While the transformation defined in Eq. (3.15) will in effect completely remove small branches from a skeleton, it will also shorten all (preserved) skeletal branches by n pixels. These branches can be completely recovered by performing some reconstruction. The branches that are removed in the pruning operation are stored in an image I_{rem} .

$$I_{\text{rem}} = I_{\text{SKEL}} \setminus \text{PRUNE}^{(n)}(I_{\text{SKEL}}). \quad (3.16)$$

The endpoints of the pruned skeleton are also found,

$$I_{\text{ends}} = \text{ENDPOINTS}(\text{PRUNE}^{(n)}(I_{\text{SKEL}})), \quad (3.17)$$

where $\text{ENDPOINTS}(I)$ denotes the endpoints transformation, the transformation that preserves only the endpoints of a skeletal image. The $\text{ENDPOINTS}(I)$ transformation is trivial to implement using hit or miss transformations. It should be noted that smaller branches will have no

endpoints in the pruned skeleton, as they will have been completely removed.

Thus reconstructing I_{ends} into $I_{\text{rem}} \cup I_{\text{ends}}$ only recovers those branch parts that came from branches that were not completely removed by the pruning transformation,

$$I_{\text{recv}} = R_{I_{\text{rem}}}(I_{\text{ends}} \cup I_{\text{rem}}). \quad (3.18)$$

If these branch ends are added to the pruned image then the resulting image contains the filtered skeleton,

$$I_{\text{filt}} = \text{PRUNE}^{(n)}(I_{\text{SKEL}}) \cup I_{\text{recv}}. \quad (3.19)$$

I_{filt} should contain the filtered skeleton of I_{SKEL} , where all branches shorter than n pixels are removed completely while longer branches are left unchanged. The value of n used in this work is equal to 10.

This well-known skeletal filtering algorithm works well, but is not perfect. One problem is that if two branches, one long and one short, both meet (join) at a distance n from the long branch’s endpoint, then both branches will be recovered by the reconstruction operation described by Eq. (3.16). This problem was overcome by repeatedly filtering the skeleton with increasing minimal branch size, from 1 to n .

If we denote the skeletal filtering transformation of size n described by Eqs. (3.14) to (3.19) by $\text{FILT}^{(n)}(I_{\text{SKEL}})$ then the repeating filter operation of size n can be described by,

$$\begin{aligned} \text{REP_FILT}^{(n)} \\ = \text{FILT}^{(n)}(\text{FILT}^{(n-1)}(\dots \text{FILT}^{(2)}(\text{FILT}^{(1)}(I_{\text{SKEL}})))) \end{aligned} \quad (3.20)$$

This avoids the problems associated with ordinary skeletal filtering. For the case where a small branch and a large branch meet as described above, the small branch will have been removed completely before the larger branch has been pruned back to the point where the two originally met, and so will not be reconstructed. Fig. 3(f) shows a skeleton calculated in this manner superimposed on the original image.

3.4.2. Identification of fiducials

We define *fiducial* points in the image as corresponding to either bifurcation points, where one vessel splits into two vessels, or crossing points, where two vessels cross. Fiducial points in the skeletal image can be found by using the hit-or-miss transform using all rotations of the composite structuring elements shown in Fig. 4 (there are eight for F_1 and four for F_2). These structuring elements correspond to pixel configurations of fiducial points.

Since the structuring elements F_1 have no background points, they detect both bifurcations (where three skeletal branches meet), and vessel crossings (where four branches meet). Mathematically, the fiducial point detection operation is described,



Fig. 4. The composite structuring elements used to detect fiducial points. Black denotes the foreground set and white the background set. Grey denotes “don’t care” squares.

$$I_{\text{mult}} = FID(I) = HMT_{F_1}(I) \cup HMT_{F_2}(I), \quad (3.21)$$

where F_1 and F_2 are as in Fig. 4 above.

The resulting detected fiducial points can be dilated by the standard square structuring element if required for display. The problem with the resulting set of points is that there are many pairs of points that in reality only correspond to a single vessel crossing point, this is because at a vessel crossing the two vessels are rarely normal to each other and thus in the skeletal representation will appear to enter at one point and leave at another point further along the vessel. To prevent multiple detections of a single crossing point the program checks the distance between fiducial points and merges fiducial points that are near each other to form a single fiducial point. In this example (640×480 pixel images), ‘near’ is defined as being within 10 pixels of each other. The co-ordinates of the merged fiducial points are recorded in an array and this array is written to an output text file where it can be used for further processing, for example, in image registration. Fig. 5 illustrates the process of merging two nearby points to form a single true fiducial point.

In order to extract useful information from the skeleton a portion of the skeleton must first be selected. We will refer to such a portion as a *vessel segment*. The next step is thus to isolate a vessel segment, defined by two points. Because the skeleton is interconnected there are usually many paths between any two points; a method is required to find the true vessel path. The algorithm finds the shortest

path between the two points on the skeleton using standard recursive search techniques. This path will usually correspond to the correct path but may sometimes give an incorrect path. To overcome this problem, a selection of the shortest paths is chosen, and the user can choose the true path, if required. In general there may be many slightly different paths that are in reality only slight variants on the same path. A metric is used to compare paths and only stores those paths that significantly differ from previously stored paths. Once the correct path has been chosen the constituent points on the skeleton are stored in a data structure. This converts the vessel segment from a binary bitmap to a set of vectors.

3.4.3. Finding the vessel width

The next step is to find the width at every point in the vessel segment. For this project, the width at a point p in a binary image is defined as the largest line segment passing through the point p that can be contained within the foreground of the binary image, at all possible rotations. An initial attempt at finding the widths used morphological openings with line segments of increasing length. Morphological opening answers the question “Can the image contain the structuring element?”, and in light of the definition of the width given above, an algorithm using openings would be intuitively correct. A method using morphological openings was developed, but was found to be computationally inefficient, as many openings with

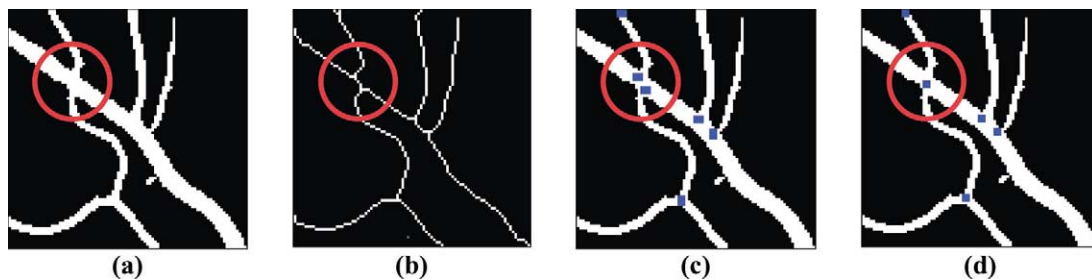


Fig. 5. A single vessel crossing may lead to two crossings in the skeletal image, one of which is spurious. (a) The original mask image. (b) The skeletal image. (c) Detected fiducial points. (d) Detected fiducial points after merging nearby points.




Fig. 6. Finding the vessel width. The line segment is extended from both sides of the pixel until a black pixel is encountered. The distance between the last two white pixels encountered is taken as the width for that rotation. The smallest distance over all rotations is taken as the vessel width.

linear structuring elements of increasing length were needed.

A faster algorithm, that shares the same ideas as the one described above, was developed (Fig. 6). The algorithm works by extending a line segment at an angle θ from both sides of a pixel in the mask image simultaneously until a background ('0') pixel is met. When the last foreground pixel (such that the next one is a background pixel) has been passed in both directions, the distance can be measured for that angle θ , as the length of the line segment. The minimum of the distance obtained using all rotations of the line segment yields the vessel width at that pixel. The width at every point in the vessel segment was computed using this method. An average vessel width figure was computed for the segment by averaging over all widths within the segment.

3.4.4. Calculation of vessel tortuosity

The tortuosity is trivial to calculate once the constituent points of the vessel segment have been recorded in their own data structure. The distance traversed by the vessel is calculated by summing the distance between consecutive points in the segment. Mathematically,

$$d_{\text{curve}} = \sum_{i=1}^{N-1} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2}, \quad (3.20)$$

where (x_i, y_i) are the co-ordinates of the i th pixel in the vessel segment, and the vessel segment has N constituent points. The straight distance is calculated as the distance between the first and last points of the vessel,

$$d_{\text{straight}} = \sqrt{(x_N - x_1)^2 + (y_N - y_1)^2}. \quad (3.21)$$

The tortuosity of a vessel segment is defined as the ratio of the curved distance and the straight distance,

$$\text{Tortuosity} = \frac{d_{\text{curve}}}{d_{\text{straight}}}. \quad (3.22)$$

3.5. Implementation

The image processing routines described above were implemented as a complete package called 'Vessel Finder'.

The programme has a convenient graphical user interface, which provides for opening and displaying of fundus image files (limited to 8- and 24-bit bitmaps). The interface is easy to use and allows fast processing of images, even for someone with little or no computer experience. The image program parameters can be changed quickly and the effects of such changes are (almost) instant. The results of the various image transformations are displayed onscreen as the algorithm proceeds. Individual blood vessel segments can be selected easily using the mouse and information about selected segments is shown onscreen, in real-time. Interested readers can obtain the programme executable by contacting the authors, or by accessing the web-pages at <http://ee.ucd.ie>.

Fast methods exist to implement erosion and dilation using linear structuring elements (Van Herk, 1993; Gil and Werman, 1993; Soille et al., 1996). These methods were used in early versions of the program but, because of certain undesirable properties, were not used in the final version. Fast morphological reconstruction was implemented as per Vincent (1993).

4. Results

4.1. Accuracy of segmentation and parameter sensitivity

In Section 3, we detailed an algorithmic procedure for vessel segmentation and analysis. This algorithm was originally developed for fluorescein angiograms, where the contrast between blood vessels and the background tends to be high, but subsequent efforts were focused on the analysis of colour fundus images taken with a wide-field digital retinal camera. The camera used is called a RetCam[®] 120 (Massie Research Laboratories, Dublin, California). The RetCam[®] is a wide-field digital retinal camera developed specifically for pediatric ophthalmology, which stores images as 24-bit colour bitmaps with a resolution of 640×480 pixels. Prior to analysis the images are converted to grayscale by simple averaging of the red and green channels. The camera gives a 120° field of view. Because of the wide field of view, the image is not perfectly in focus everywhere. Contrast also tends to be relatively poor, and many images taken with this camera suffer from motion artefact as the infants' eyes move. A further complication is that the camera is handheld, so observer motion will also contribute to image degradation. Other researchers (Roth et al., 2001) have compared expert interpretation of RetCam[®] images with direct ophthalmic examination, and have found that the camera images are somewhat inferior to direct examination, particularly in more peripheral regions. Hence, it is reasonable to say that accurate vessel segmentation in this set of images is quite challenging when compared to the same task in fluorescein angiograms where the higher contrast and narrower field of view make for far better images.

Since we wished to conduct quantitative studies of changes in vessel widths and tortuosities measured from these images, we first carried out an analysis of the accuracy of our vessel segmentation algorithm, and the effect of various algorithmic parameters on this accuracy. Adjustable parameters in our algorithm include: (a) the lengths L_1 and L_2 of the structuring elements used in Eqs. (3.1) and (3.2), (b) the values of ρ used in Eq. (3.5), (c) the scaling parameter K used in Eq. (3.7), and (d) the values of the threshold parameters used in Eq. (3.12). Of these, the most significant parameters turned out to be at the thresholding stage.

To provide some quantitative assessment of this sensitivity, the following experiment was carried out. Segmentation was performed on a 260×260 pixel section of a typical image using the algorithm described above. The thresholding scheme was simplified by setting $t_{\text{low}} = t_{\text{high}} = Th$. This diminishes slightly the possible performance, but greatly reduces the parameter space to be explored. The other parameters were set as follows: $L_1 = 17$, $L_2 = 17$, $K = 10$, $\rho = 1.75$ and $n = 10$. The value of the threshold Th was then swept through the range from 20 to 65. To quantify the algorithmic performance, the resulting segmentations were compared to a ‘gold-standard’ image. The gold-standard image was obtained by manually creating a vessel mask in which all vessel pixels were set to one, and all non-vessel pixels set to zero. In practice, even this gold-standard image is also subjective since a skilled human observer will incorrectly annotate certain pixels. Using this gold-standard image, the accuracy of the automated vessel segmentation can be assessed. The gold-standard image is used to decide whether a vessel is truly present or absent, and the results of the segmentation algorithm decide whether a vessel has been detected or not. For each value of the threshold Th , the following table is created:

	Vessel present		Vessel absent	
Vessel detected	True (TP)	positive	False (FP)	positive
Vessel not detected	True (TN)	negative	False (FN)	negative

The sensitivity (Se) and specificity (Sp) of the segmentation are then defined as

$$Se = \frac{TP}{TP + FN}, \quad Sp = \frac{TN}{TN + FP}. \quad (4.1)$$

Fig. 7(a) shows how these quantities vary as Th changes. A perfect segmentation algorithm will yield sensitivity and specificity values equal to unity; any real algorithm falls short of the ideal. The behaviour of the sensitivity and specificity curves as Th varies is as expected. As Th increases, the sensitivity decreases (i.e., more vessels are

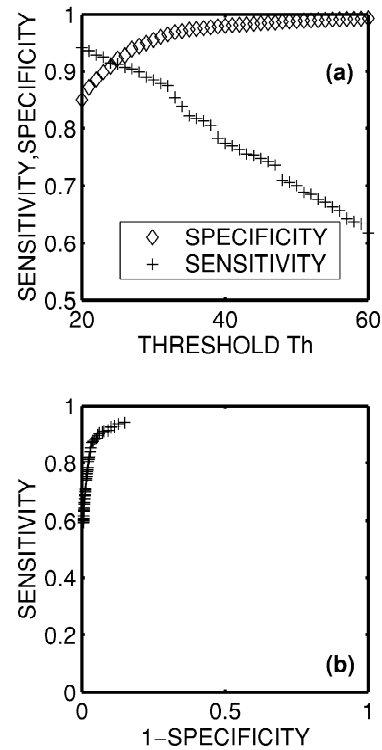


Fig. 7. The sensitivity and specificity of the segmentation algorithm as a function of the chosen threshold value. (a) Sensitivity (+) decreases as the threshold parameter increases, whereas specificity (◇) increases. (b) A receiver operating characteristic curve for the segmentation algorithm. Ideal segmentation would yield the value of SENSITIVITY equal to unity for all values of (1-SPECIFICITY) except zero.

missed which are really present), but the specificity increases (i.e., the number of spurious vessel detections reduces). An alternative way of presenting this information is as a receiver-operating characteristic (ROC) curve [Fig. 7(b)]. This is a plot of (1-SPECIFICITY) versus SENSITIVITY, and ideally should equal a right angle plot with SENSITIVITY=1 for all values of (1-SPECIFICITY) except 0. The closer the ROC curve is to the ideal, the better the overall performance of the algorithm. An optimum choice of threshold yields a sensitivity of approximately 92% and a specificity of approximately 92%. These figures are nearly identical with those quoted by Gardner et al. (1996) for the performance of their segmentation algorithm for vessel identification based on artificial neural networks. Their image set consisted of red-free colour fundus photographs. It is an interesting footnote that two such distinct methods can yield highly similar results.

The most important aspect of Fig. 7(a), however, is that the sensitivity and specificity vary quite rapidly with the choice of threshold (e.g., sensitivity varies by about 10% for a change in threshold from 20 to 30). Unfortunately, this means that the choice of threshold has a critical effect on the accuracy of the segmentation. Therefore, in our current implementation of the algorithm, the thresholding

stage incorporates an interactive choice of threshold, so that the user can control the trade-off between sensitivity and specificity. In practice, we use relatively low values of Th (e.g., 25–35) since at a later stage we can choose which vessels to analyse so that spurious vessels can be neglected at that stage. In other words, it is better to use a low threshold which gives good sensitivity and regain higher specificity in our later choice of vessels to analyze.

Other algorithmic parameters also affect the accuracy of segmentation, though in general not as significantly as the threshold parameters. Using the same 260×260 pixel image, with a fixed threshold value of 30, the influence of other parameters on segmentation performance was considered. Fig. 8(a) shows how the sensitivity and specificity vary as a function of the length L_1 of the structuring element used in Eq. (3.1). As L_1 is increased, the sensitivity decreases and the specificity increases. This is expected; as the structuring element is made longer it only selects long straight sections of vessels, and starts to miss shorter vessels. In particular, too large a value of L_1 will cause tortuous vessels to be missed. However, it is less likely to emphasise small isolated artefacts, hence the specificity is

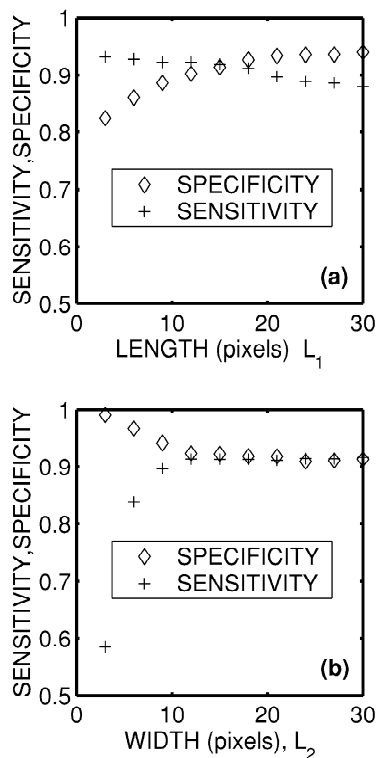


Fig. 8. The sensitivity and specificity of the segmentation algorithm as a function of the values L_1 and L_2 used in the morphological processing. (a) Sensitivity (+) decreases as L_1 increases, whereas specificity (◇) increases. This is because a longer structuring element will be more selective since only straight sections of vessels will be retained by a longer structuring element. (b) Sensitivity (+) increases as L_2 increases whereas specificity (◇) decreases. This is because too short a structuring element will reject true vessels which are wide. Beyond a certain point, however, the parameter L_2 will have negligible impact on segmentation performance.

increased. However, if L_1 is chosen somewhere in the range 10–20 pixels, performance is reasonably constant. Fig. 8(b) shows the changes in sensitivity and specificity as the length L_2 of the structuring element in Eq. (3.2) is varied. For very small values of L_2 , only very narrow vessels are detected, so the sensitivity is poor. However, for values of L_2 in the range 10–30 pixels, performance is virtually constant.

We also considered the influence of the scaling parameter K and found it to be of little influence (provided the threshold parameter is scaled proportionally). Segmentation performance was also relatively insensitive to the value of ρ used in the kernels of Eqs. (3.5).

4.2. Accuracy of vessel width and tortuosity measurement

A goal of this research is to accurately assess changes in vessel width and tortuosity. To verify the accuracy of vessel measurements, a test image [Fig. 9(a)] was synthesized which consists of ‘vessels’ of known width and tortuosity. The straight vessels on the left of this figure range in width from 1 to 6 pixels. They are of constant cross-section with a luminance value of 208 against a background of Gaussian random luminance with a value 128 ± 5 . The tortuous vessels on the right hand side are simple sinusoidal curves of the form $A \sin(\theta)$, with a line thickness of 3 pixels, and a luminance of 208. As the value of A is increased, the tortuosity increases. By numerically performing the arc-length integral,

$$C = 2 \int_0^\pi (1 + A^2 \cos^2 \theta) d\theta, \quad (4.2)$$

the theoretical tortuosity of the vessels can be calculated as a function of A . Fig. 9(b) shows the measured width (in pixels) against the theoretical width. Under these ideal circumstances, the vessel width algorithm performs consistently. The only caveat is that the threshold value should be chosen at a sufficiently low level so that a high level of sensitivity is achieved – otherwise the widths will be under-estimated. Fig. 9(c) shows measured tortuosity against the ideal tortuosity value. The tortuosity algorithm consistently over-estimates the tortuosity by approximately 0.07. This overestimate is due to the discrete nature of a digital image, since pixels have a finite size and hence measuring distance from center to center of each pixel will force a theoretically smooth curve to follow a zig-zag path with a longer length. However, the estimate is consistently biased over the range of tortuosities of interest, so that measurements of tortuosity can be considered as reliable in a relative sense.

4.3. Analysis of ROP images

The goal of the current study is to use image analysis to




Fig. 9. (a) A simple test image that was used to test the accuracy of vessel width and tortuosity measurement. The test image consists of vessels of known widths between 1 and 6 pixels, and vessels of known tortuosities between 1.05 and 1.4. (b) A comparison between vessel width measured using the automated technique and actual real vessel width. (c) A comparison between vessel tortuosity measured using the automated technique and theoretical tortuosity. A consistent bias can be seen due to finite geometric effects.

quantify changes in vessel width and tortuosity in subjects being screened for ROP, since such changes may be of diagnostic or prognostic significance. In the following section, we detail the analysis of a database of retinal images obtained from subjects presenting at The National Maternity Hospital, Holles Street, Dublin between

November 1998 and April 2000. Twenty-three subjects are included in this study, as detailed in Table 1. The average gestational age at birth was 27 weeks, and average birth weight was 975 g.

These subjects were examined in accordance with hospital guidelines, which follow international best prac-

Table 1
Summary of database parameters

Finding	Number	Average gestational age at birth (weeks)	Average weight at birth (g)
Retinopathy of prematurity	11	27	900
Spontaneously regressed retinopathy of prematurity	3	26	860
No evidence of retinopathy of prematurity	9	28	1060

tice in screening. First examinations took place between 4 and 6 weeks of chronologic age or between 31 and 33 weeks post-conceptional age. Subsequent examinations depended on the presence and severity of any ROP detected. Subjects with no ROP may have only one further examination 6 weeks later whereas subjects with significant disease may require weekly or even daily examinations. These examinations were carried out with an indirect ophthalmoscope, a device that allows the ophthalmologist to get a stereoscopic wide-angle view of the retina. Simultaneously digital images of the retina were captured using the RetCam[®]. Subjects who were judged to have threshold disease, and who were subsequently treated, were examined between 1 and 6 times prior to treatment. The subjects who had spontaneous regression of ROP were all examined at least three times, while those with no sign of disease were all examined between two and three times.

The goal of the current study is to quantify changes which occurred in these subjects, specifically changes in vascular width and tortuosity near the posterior pole. The protocol for doing this was as follows. Images from each subject visit were stored (though some of these had to be excluded from the analysis due to poor lighting, poor focus, or motion artefact). Segmentation and skeletonization of the vessel structures were carried out using the algorithm described. Fig. 10 gives some typical results of this process. Fig. 10(a) shows the original colour fundus image. Fig. 10(b) shows the image after conversion to greyscale and morphological preprocessing. Fig. 10(c) shows the result after the differential operations and morphological postprocessing. Since the thresholding stage is critical, the user is able to select values and view results before choosing the final threshold values. Fig. 10(d) shows the result of the thresholding stage, while Fig. 10(e) shows the final skeleton of the vasculature. The inaccuracies seen in this image are representative of the errors normally encountered. Specifically, the edge of the optic disk is often incorrectly segmented as a vessel. Some vessels with low contrast relative to the background are missed, and spurious side-spurs are detected off major vessels. Many spurious segments are detected due to underlying choroidal vessels slightly below the center of the image. Fig. 10(f) shows how a region of interest near the posterior pole is defined as an annulus whose inner boundary is the optic disk, and whose outer boundary is a circle of two optic-disk diameters centred about the optic disk. Within this annulus, the user selects vessels for measurement in the region-of-interest using the mouse. Ideally, arteries and veins should be considered separately, but since it is hard to reliably identify the veins and arteries across images, the analysis is carried out using the combined average. Moreover, since it would be difficult to automate the process of differentiating between veins and arteries, an automated computer-aided screening system is more likely to rely upon measurement of both veins and arteries. In most cases, at least four arteries and four veins

can be identified (i.e., at least one vessel feeding and draining each quadrant). The average width of a vessel was calculated by evaluating its width along its length, and averaging. By further averaging, the overall average width of the main vessels within the region of interest is calculated. The tortuosity of a vessel was defined as the ratio of the actual vessel length over the distance between its end points. In order for this definition to be reasonable, vessels of significant length must be used (i.e., longer than 500 μm). The average tortuosity was calculated by averaging over the weighted tortuosity of all chosen vessels (i.e., the tortuosity value of each vessel was weighted by the length of that vessel). Fig. 10(f) shows how the width, tortuosity and length of each selected vessel was displayed to the user, as well as an overall average set of readings. Vessels which were of insignificant length were not included in the analysis, nor were minor side spurs of the main vessels.

Under ideal conditions, it would be desirable to track exactly the same set of vessels across every visit of the subject. In practice, this was not possible. The desire to obtain good quality images has to be weighed against the desire to minimise the time examining the subjects, as premature infants are typically sick and physiologically unstable subjects and the examination process may cause significant transient changes in an infant's general condition. As a result, only a small set of images could be obtained at each visit, and the image quality is highly variable as a result. Accordingly, the figures extracted as 'average vessel width' and 'average vessel tortuosity' from each examination are imprecise, as there are many potential sources of error in their calculation. Any or all of the following factors can affect the measured value: variations in camera focal length, quality of image, choice of threshold parameters, vessels selected for measurement, the size and centre of the region of interest, and the location of the endpoints of selected vessels. Since the final values for average vessel width and tortuosity are the sum of many small measurements it seems reasonable to model uncertainty in the measurements as Gaussian noise. In order to obtain some quantitative estimate of how reliable the measurements are, we repeated the measurement of average vessel width and average tortuosity ten times on the same image, for several different images. The choice of threshold, vessels selected, and the size of the region-of-interest varied slightly from run to run. The standard deviations of the average width and tortuosity varied from image to image, but a reasonable *upper* bound for the standard deviation of average vessel width measurement was $\sigma_w = 12 \mu\text{m}$. The corresponding standard deviation for average vessel tortuosity measurement was $\sigma_t = 0.01$. Based on these measurements, we conclude that the accuracy of average vessel width measurements are on the order of $\pm 12 \mu\text{m}$, and that the average tortuosity measurements are accurate to within ± 0.01 .

By using a single image for each eye from each subject

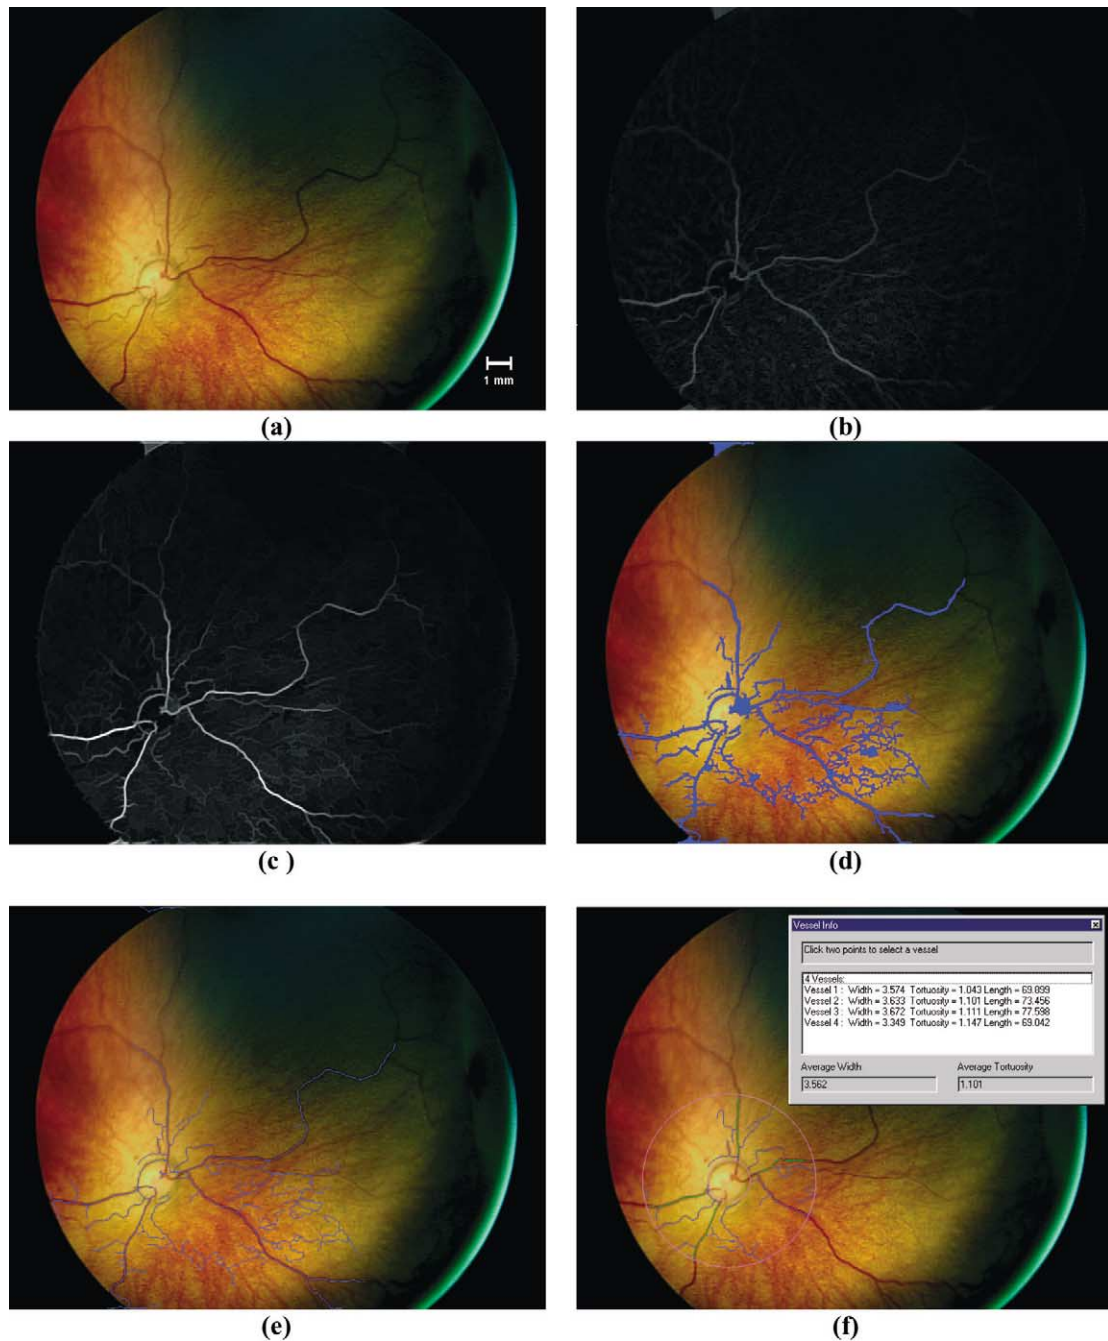


Fig. 10. Panel (a) shows the original RetCam[®] image. Panel (b) shows the image after the initial morphological preprocessing. Panel (c) shows the image after use of the second derivative and morphological post-processing. Panel (d) shows the vessels segmented after hysteresis thresholding with $t_{\text{low}} = 30$ and $t_{\text{high}} = 45$. Panel (e) shows the vessel skeletons after skeletonization. Panel (f) shows a sample screen shot, in which vessels are identified within and annulus constrained by the optic disc and a circle of radius equal to two optic diameters. The optic disc is manually identified. Vessels within the annulus can be chosen by selecting with a mouse. Selected vessels are highlighted in green. The units of width in the analysis window are pixels.

visit, it is possible to draw up a time series of average vessel width and tortuosity. For example, Fig. 11 shows the time course of the vessel width and tortuosity for the left eye of a subject who was judged to have reached threshold disease. The subject was seen three times prior to treatment (the last visit being on the day of treatment itself), and twice after treatment. On first examination the

average vessel width was $82.5 \mu\text{m}$, and the average tortuosity was 1.12. On the day of treatment the average vessel width was $100.8 \mu\text{m}$ and the average tortuosity was 1.16. This figure shows both the rapid increase in vessel width and tortuosity prior to treatment, and how the vessel widths and tortuosity resume their normal values within a few days of treatment. This behaviour is typical of many of




Fig. 11. This indicates how the vessel width and tortuosity vary as a function of time for a single eye in one subject. The subject was examined five times, three times prior to treatment and twice after. The increase and subsequent reduction in vessel width is clearly apparent. A significant increase in vessel tortuosity is also apparent.

the time series we have examined. Fig. 12 shows some of the images from which this time series was extracted. Fig. 12(a) illustrates the image acquired on the first day of examination, while Fig. 12(b) is the image acquired on the day of treatment.

However, the time course of vessel width change in the posterior pole is not always so clear-cut. Fig. 13 shows time series for all subjects of the measured average vessel width. For clarity of presentation, the time series are

presented for each different class (regressed ROP, threshold disease ROP, non-ROP) separately. This plot confirms the subjective observation of clinicians regarding increasing width of the vasculature in subjects with threshold disease ROP, whereas subjects with mild or no ROP experience no significant changes. Fig. 14 shows similar time series plots for all subjects of the average vessel tortuosity. In certain cases, increased tortuosity can be definitely associated with the presence of threshold dis-




Fig. 12. Panel (a) shows the retinal image (000004ce.bmp) acquired at gestational age 236 days in a patient being screened for ROP. Panel (b) shows the retinal image (00000513.bmp) acquired at gestational age 257 days, at which point the patient was judged to have reached threshold disease, and was treated. The increase in vessel width and tortuosity is clearly apparent.

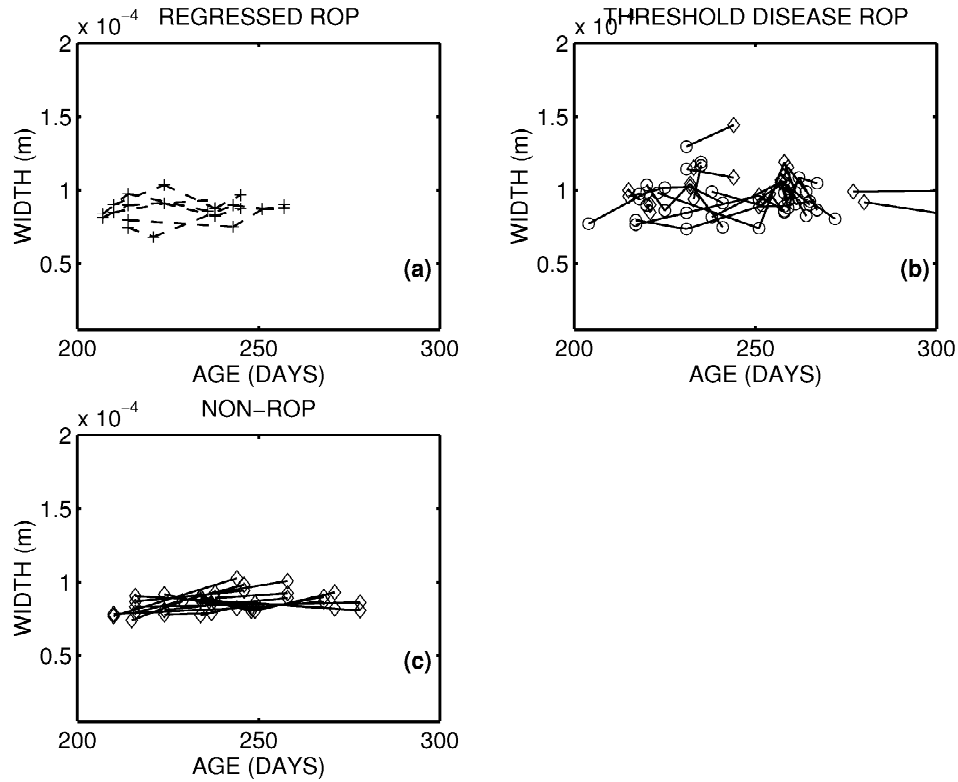


Fig. 13. This indicates how the average vessel width varies as a function of time for all eyes from all subjects in the study. For clarity, the three classes of ROP subjects are shown separately in panels (a), (b) and (c). The lines connect the measurements made on an eye over several different examinations. For the threshold disease ROP subjects, the measurement on the day of treatment is denoted with a \diamond symbol. Two features are apparent in these graphs: (i) in general the threshold disease ROP subjects have higher average vessel widths when they are judged to have reached threshold disease, and (ii) the threshold disease ROP cases tend to experience a rapid increase in vessel width over relatively short periods of time. The subjects not requiring treatment tend to have more consistent average vessel widths.

ease, but this pattern is not consistent. To quantify this numerically, the absolute values of average vessel width and tortuosity was compared for subjects treated for ROP (21 eyes) and those not treated (i.e., regressed and non-ROP classes together comprising 21 eyes in total). The average vessel width and tortuosity for the ROP subjects was $96.8 \mu\text{m}$ and 1.125 respectively, on the day of treatment. Comparative figures for the non-treated class are $86.4 \mu\text{m}$ and 1.097 over all examinations. Using a z -test, the difference in these mean values is significant at the 99% and 95% significance levels, respectively (Devore and Peck, 1996). In other words, ROP subjects do have wider and more tortuous vessels, at least in a statistical sense. However, there is quite a degree of overlap between the measured average vessel widths and tortuosity for threshold disease ROP, non-ROP, and regressed ROP subjects.

Subjectively, it has also been noted that the threshold disease cohort experience a sudden rapid increase in average vessel width at somewhere between 30 and 38 weeks of gestational age. As a statistical indicator of the extent of this increase we calculated the difference in average vessel width and tortuosity between the first day of examination, and the day of treatment. Since some ROP subjects were treated on their first day of examination, only

14 eyes were included in this measure. The average change in average vessel width was $w_d = 9.6 \mu\text{m}$, and the average increase in tortuosity was $t_d = +0.008$. The statistical significance of these numbers was assessed in the following manner. A paired t -test was used in which the null hypothesis was that the observed increase was as follows,

$$H_0^{\text{width}}: w_d \leq \sqrt{\frac{2\sigma_w^2}{n}}, \quad H_0^{\text{tort}}: t_d \leq \sqrt{\frac{2\sigma_t^2}{n}}, \quad (4.3)$$

where n is the number of paired measurements. Qualitatively, we are evaluating the probability of observing an increase in the vessel width and tortuosity purely due to the estimated inaccuracies in the measurements themselves. Our choice of null hypothesis is quite conservative since it assumes a worst case scenario where the errors in measurement all fall in the wrong direction. This null hypothesis for vessel width increase was rejected at the 95% level. However, the null hypothesis for tortuosity increase was accepted. We conclude that the measured vessel width change is statistically significant at the 95% level, but that the measured changes in tortuosity are insufficient for statistical significance. A similar analysis was carried out for the observed changes in average vessel

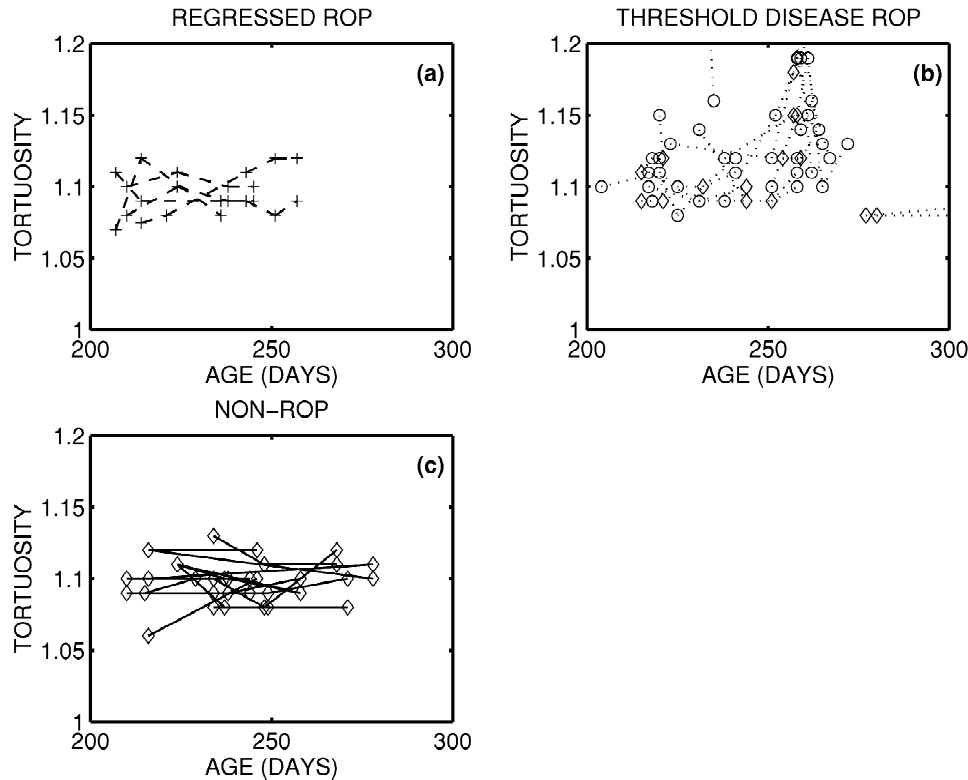


Fig. 14. This indicates how the vessel tortuosity varies as a function of time for all eyes from all subjects in the study. For clarity, the three classes of ROP are shown separately in panels (a), (b) and (c). The lines connect the measurements made on an eye over several different examinations. For the threshold disease ROP subjects, the measurement on the day of treatment is denoted with a \diamond symbol. Some of the ROP cases have a rapid increase in vessel tortuosity, but this pattern is not consistently repeated. Tortuosity is fairly constant in the non-threshold cases.

width and tortuosity for both the non-ROP and regressed ROP subjects. In both cases, the increase was calculated between the first and last day of examination. For the non-ROP subjects the observed changes in vessel width and tortuosity were $+4.6 \mu\text{m}$ and -0.001 , and for the regressed-ROP subjects the corresponding figures were $+7.1 \mu\text{m}$ and $+0.007$. None of these changes are statistically significant at the 95% significance level, using a null hypothesis test similar to Eq. (4.3).

Figs. 13 and 14 also show that a qualitative decrease in vessel width and tortuosity is observed following treatment. For the 17 eyes for which post-treatment measurements are available, the average decrease in vessel width is $10 \mu\text{m}$ and the average decrease in tortuosity was 0.023. Both of these changes are statistically significant at the 95% level if a null hypothesis similar to Eq. (4.3) is used – in other words there is a meaningful decrease in average vessel width and tortuosity following treatment.

4.4. Potential implications for automated screening

In the previous section, we showed (a) that there is a statistically significant difference in absolute average vessel width and tortuosity between ROP subjects requiring treatment and those not, and (b) that there is a statistically significant increase in average vessel width for ROP

subjects. In order to be clinically significant, it is reasonable to devise some classification scheme based on these measured parameters, which can provide some automated aid to clinical judgement.

Since the goal of the screening process is to assess which subjects should receive treatment for ROP, and which should remain untreated, we proposed the following simple test to answer this question,

IF {
 (maximum average vessel width observed in either
 eye > width threshold)
 OR
 (maximum average tortuosity observed in either
 eye > tortuosity threshold)
 }
 THEN
 treat subject.

This is an extremely simplified test since it does not include any of the known risk factors such as low birth weight, birth age, etc., nor does it attempt to normalise the measured average vessel width relative to infant size. It is a bilateral test, in that measurements from both eyes are included in the test statistic, and the outcome is a decision

to treat both eyes. This is unrealistic, as ROP can be unilateral in a small number of cases, and current practice is to evaluate on a per-eye basis rather than on a per-subject basis. However, it does provide a first attempt at devising computer-aided screening based on measurement of width and tortuosity. The sensitivity and specificity of this test change as the threshold values are altered. Fig. 15 shows a receiver-operating-characteristic values for this test over a large range of possible threshold values. The best overall performance is achieved by setting the width threshold equal to $98\ \mu\text{m}$ and tortuosity threshold equal to 1.12. For those parameters, the sensitivity is 82% and the specificity is 75%. Since these figures are obtained retrospectively using a training set, they are probably a slight overestimate of the actual performance on a new test set. To provide some comparison, we include an ROC curve calculated using the published data from Onofrey et al. (2001). Their automated screening technique based on weighting of known risk factors, and overall physician judgement produces a predicted risk estimate for an eye to proceed to threshold disease requiring treatment. By varying the level of risk estimate at which treatment is initiated, it is possible to draw up an ROC curve similar to our results. The technique based on actual vessel measure-

ment does outperform the risk-factor based approach. However, we would agree with Onofrey's statement that "although a software program may be a useful tool for following premature infants with ROP, the clinical examination remains the gold standard", since the sensitivity and specificity values obtained using our system are relatively modest.

5. Conclusions and discussion

We have shown that segmentation of the vascular structure in retinal images is possible by use of a combination of morphological and linear filtering. The quality of the segmentation is dependent on a number of parameters such as image quality, choice of threshold, and choice of structuring elements. Successful segmentation allows a variety of further processing such as:

- Visual highlighting of vessels in the image
- Accurate characterisation of vessel parameters such as thickness and tortuosity
- Location of vessel bifurcation and crossings which can act as intrinsic features for registration schemes.

We have provided some quantitative assessment of the

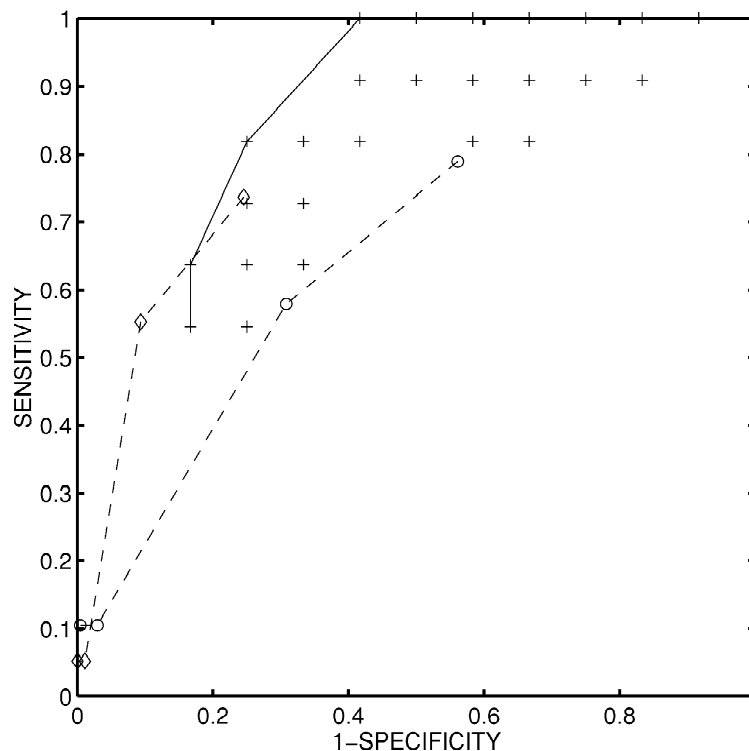


Fig. 15. Receiver-operating-characteristic (ROC) values for the screening test proposed in Section 4.4, for various choices of the threshold values used in the screening test. The ROC values plot SENSITIVITY versus (1 – SPECIFICITY), and are denoted with '+'. Points which are near the value (0.1) are optimal. The solid curve is drawn to indicate the ROC curve corresponding to an optimal choice of threshold values. For example, a suitable choice of screening parameters yields a sensitivity of 82% and specificity of 75%. For comparison, we have drawn up an ROC curve using the data published by Onofrey et al. (2001). The dashed curve marked by '◊' is the ROC curve for their screening test carried out on the first day of examination. The performance is quite good since there are many non-ROC cases which are easily rejected. The dashed curve marked by '○' is the ROC curve for their screening test carried out on a subsequent examination, and specificity is reduced since there are fewer easy decisions to make since the non-ROP cases have been excluded. The performance of their risk-based screening test appears inferior to the screening test we have proposed.

accuracy of our segmentation algorithm to facilitate comparison with other techniques. For a typical image, it is reasonable to expect that approximately 90% of vessels will be detected, and that approximately 10% of non-vessel pixels will be incorrectly classified as vessels. Gardner et al. give the only other quantitative assessment of vascular segmentation. Their artificial neural network classifier gives a sensitivity and specificity of approximately 91% for vessel detection in red-free fundus images (Gardner et al., 1996), which is highly similar to our result. Unfortunately, quantitative assessment of segmentation accuracy is not available for other techniques. The following subjective comments can be made, however, about alternative vessel segmentation strategies. The manual techniques of Mainster and Daxer are definitely highly accurate but not suitable for automation (Mainster, 1990; Daxer, 1993). Matsopoulos et al. (1999) use a combination of morphological processing and matched filtering to obtain their segmentation. The performance of their technique appears highly similar to that proposed here (Fig. 2 of their paper); some vessels are missed, and some spurious vessels are apparent. In particular, their technique is also confounded by the edge of the optic disk. Kochner et al. (1998) estimate that their technique (based on course tracking and steerable filters) detects approximately 70–90% of vessels present in fundus images. Their technique detects few spurious vessels since it tracks vessels from known starting points. The matched filter approach of Chaudhuri et al. (1989) also appears to give similar performance to our technique (Fig. 7 of their paper) with the optic disk edge again being the source of spurious vessel detection. Finally comparison was made with the techniques of Zana and Klein (1997, 1999) which are based on morphological processing and cross-curvature. Subjectively, their technique appears to give the best results of the techniques available in the literature, though again the optic disk edge can be wrongly interpreted as a blood vessel. In addition, their results are mostly presented for fluorescein angiograms which have higher contrast.

We have used our segmentation algorithm to characterise changes in the vasculature of the posterior pole of subjects being screened for retinopathy of prematurity, with the motivation ultimately of using such measurements as an aid in screening for threshold disease. A variety of automated approaches to screening for ROP have previously been suggested. Onofrey et al. (2001) used an automated program to weight the impact of various risk factors in predicting outcomes. Fig. 15 indicates that their predictive test appears to perform worse than a simple screening technique based on vessel width and tortuosity only. Vogele et al. (1999) have also adopted a similar approach.

Other researchers have advocated screening which is directly based on simplified clinical examinations which do not require the use of indirect ophthalmoscopy to view the periphery of the retina. Saunders et al. (2000) have proposed that it may be possible for non-ophthalmologists

to screen for ROP requiring treatment by assessing the posterior pole only, and grading the blood vessels by dilation and tortuosity. In their study, they were able to show that screening by a non-ophthalmologist can achieve a sensitivity of 100%, but that specificity was poor. In a similar vein, Wallace et al. (2000) have suggested that the presence of ‘pre-plus’ disease may act as a useful predictor of eyes which will go on to develop clinically significant ROP. Plus disease, an important prognostic indicator for this condition, is currently clinically classified as absent or present. Plus disease is qualitatively characterised as having increased vascular dilation and tortuosity. They conducted a prospective study in which they *qualitatively* judged for the presence of plus disease by characterising vascular dilation and tortuosity in the posterior pole using five different grades. They concluded that the presence of ‘mild’ dilation and tortuosity increases was a significant indicator for the need for treatment. Since the algorithm proposed here is able to measure changes in both vascular width and tortuosity, it provides a quantitative basis for their approach.

In conclusion, we have shown that it is possible to segment out the retinal vasculature from wide-angle colour images of the retinas of premature infants. This has allowed quantitative measurement of both vessel width and tortuosity in the posterior pole, and has shown statistically significant differences between subjects with threshold ROP disease requiring treatment and those without. It can also be used to show that significant increases occur in vessel width as ROP develops to the point of threshold disease, and that a significant decrease is seen after treatment. Finally, measurement of vessel width and tortuosity can be used to provide a simple automated screening test for detecting subjects with threshold disease, though at present the sensitivity and specificity of this technique is considerably below that needed for clinical practice.

Acknowledgements

The authors are grateful to B. Lanigan for facilitating the collection of the clinical imagery used in this study. The authors also acknowledge the work of the anonymous reviewers for their helpful comments. This work was supported by the Institute of Ophthalmology, the University College Dublin President’s Research Award, and Enterprise Ireland.

References

- Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., Goldbaum, M., 1989. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Trans. Med. Imaging* 8, 263–269.
- Chapman, N., Witt, N., Gao, X., Bharath, A.A., Stanton, A.V., Thom, S.A., Hughes, A.D., 2001. Computer algorithms for the automated

- measurement of retinal arteriolar diameters. *Brit. J. Ophthalmol.* 85, 74–79.
- Committee for the Classification of Retinopathy of Prematurity, 1984. International classification of retinopathy of prematurity. *Arch. Ophthalmol.* 102, 1130–1134.
- Cryotherapy for Retinopathy of Prematurity Cooperative Group, 1988. Multicenter trial of cryotherapy for retinopathy of prematurity: preliminary results. *Arch. Ophthalmol.* 106, 471–479.
- Cryotherapy for Retinopathy of Prematurity Cooperative Group, 1990a. Multicenter trial of cryotherapy for retinopathy of prematurity: three-month outcome. *Arch. Ophthalmol.* 108, 195–204.
- Cryotherapy for Retinopathy of Prematurity Cooperative Group, 1990b. Multicenter trial of cryotherapy for retinopathy of prematurity: one-year outcome – structure and function. *Arch. Ophthalmol.* 108, 1408–1416.
- Daxer, A., 1993. Characterisation of the neovascularisation process in diabetic retinopathy by means of fractal geometry: diagnostic implications. *Graef Arch. Clin. Exp.* 231, 681–686.
- Devore, J., Peck, R., 1996. *Statistics: The Exploration and Analysis of Data*, 3rd Edition. Brooks/Cole, Pacific Grove, CA.
- 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. *Brit. J. Ophth.* 80, 940–944.
- Gil, J., Werman, M., 1993. Computing 2-D min, median and max filters. *IEEE Trans. Patt. Anal. Mach. Intell.* 15, 504–507.
- Kochner, B., Schuhmann, D., Michaelis, M., Mann, G., Englmeier, K.-H., 1998. Course tracking and contour extraction of retinal vessels from color fundus photographs: most efficient use of steerable filters for model based image analysis. In: *Proc. SPIE Medical Imaging*, pp. 22–28.
- Mainster, M.A., 1990. The fractal properties of retinal vessels: embryological and clinical implications. *Eye* 4, 235–241.
- Matsopoulos, G.K., Mouravliansky, N.A., Delibasis, K.K., Nikita, K.S., 1999. Automatic retinal image registration scheme using global optimization techniques. *IEEE Trans. Inf. Technol. Biomed.* 3, 47–60.
- McNamara, J.A., Connolly, B.P., 1999. Retinopathy of prematurity. In: Regillo, C.D., Brown, C.G., Flynn, H.W. (Eds.), *Vitreoretinal Disease: The Essentials*. Thieme Medical Publishers, New York, pp. 177–192.
- McNamara, J.A., Tasman, W., Vander, J.F., Brown, G.C., 1992. Diode laser photocoagulation for retinopathy of prematurity. Preliminary results. *Arch. Ophthalmol.* 110, 1714–1716.
- Onofrey, C.B., Feuer, W.J., Flynn, J.T., 2001. The outcome of retinopathy of prematurity – Screening for retinopathy of prematurity using an outcome predictive program. *Ophthalmology* 108, 27–34.
- Palmer, E.A., Flynn, J.T., Hardy, R.J., Phelps, D.L., Phillips, C.L., Schaffer, D.B., 1991. Incidence and early course of retinopathy of prematurity. *Ophthalmology* 98, 1628–1640.
- Roth, D.B., Morales, D., Feuer, W.J., Hess, D., Johnson, R.A., Flynn, J.T., 2001. Screening for retinopathy of prematurity employing the RetCam 120 – sensitivity and specificity. *Arch. Ophthalmol.* 119, 268–272.
- Saunders, R.A., Donahue, M.L., Berland, J.E., Roberts, E.L., Von Powers, B., Rust, P.F., 2000. Non-ophthalmologist screening for retinopathy of prematurity. *Brit. J. Ophthalmol.* 84, 130–134.
- Serra, J., 1982. *Image Analysis and Mathematical Morphology*. Academic Press, London.
- Soille, P., 1999. *Morphological Image Analysis: Principles and Applications*. Springer-Verlag, Berlin.
- Soille, P., Breen, E., Jones, R., 1996. Recursive implementation of erosions and dilations along discrete lines at arbitrary angles. *IEEE Trans. Patt. Anal. Mach. Intell.* 18, 562–567.
- Sonka, M., Hlavac, V., Boyle, R., 1998. *Image Processing, Analysis, and Machine Vision*, 2nd Edition. Brooks/Cole, Pacific Grove, CA.
- Sun, Y., 1989. Automated identification of vessel contours in coronary angiograms by an adaptive tracking algorithm. *IEEE Trans. Med. Imag.* 8, 78–88.
- Van Herk, M., 1993. A fast algorithm for local minimum and maximum filters on rectangular and octagonal kernels. *Patt. Recogn. Lett.* 13, 517–521.
- Vincent, L., 1993. Morphological greyscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Trans. Image Proc.* 2, 176–201.
- Vogele, C., Seiberth, V., Jendritza, W., Akkoyun-Vardarli, I., Liesenhoff, H., 1999. Software for retinopathy of prematurity screening. *Ophthalmologie* 96, 82–86.
- Wallace, D.K., Kylstra, J.A., Chesnutt, D.A., 2000. Prognostic significance of vascular dilation and tortuosity insufficient for plus disease in retinopathy of prematurity. *Journal of AAPOS* 4, 224–229.
- Zana, F., Klein, J.C., 1997. Robust segmentation of vessels from retinal angiography. In: *Proc. Int. Conf. Digital Signal Processing*, Santorini, Greece, pp. 1087–1091.
- Zana, F., Klein, J.C., 1999. A multimodal registration algorithm of eye fundus images using vessels detection and Hough transform. *IEEE Trans. Med. Imaging* 18, 419–427.
- Zhou, L., Rzeszutarski, M., Singerman, L., Cokreff, J., 1994. The detection and quantification of retinopathy using digital angiograms. *IEEE Trans. Med. Imaging* 13, 619–626.