

TRAM-LINE FILTERING FOR RETINAL VESSEL SEGMENTATION

Andrew Hunter*, James Lowell* and David Steel**

* University of Lincoln
** Sunderland Eye Infirmary

ahunter@lincoln.ac.uk

Abstract: The segmentation of the vascular network from retinal fundal images is a fundamental step in the analysis of the retina, and may be used for a number of purposes, including diagnosis of diabetic retinopathy. However, due to the variability of retinal images segmentation is difficult, particularly with images of diseased retina which include significant distractors. This paper introduces a non-linear filter for vascular segmentation, which is particularly robust against such distractors. We demonstrate results on the publicly-available STARE dataset, superior to Stare's performance, with 57.2% of the vascular network (by length) successfully located, with 97.2% positive predictive value measured by vessel length, compared with 57% and 92.2% for Stare. The filter is also simple and computationally efficient.

1 Introduction

This paper presents algorithm for the automatic segmentation of the vascular network from retinal images. No user intervention is required. The algorithm is designed to be particularly robust against the type of distractors common in images of diseased retina. We have tested the algorithm against the publicly-available STARE database [12], and against our own diabetic retinopathy database. Results indicate that the algorithm is exceptionally robust, with a very low false positive rate.

1.1 Motivation

The algorithm presented in this paper forms a key component of our system for automated detection and grading of diabetic retinopathy. This is a common complication of diabetes mellitus, with a prevalence of approximately 40% in the diabetic population [18], and is the leading cause of blindness in the working population of Western countries [15]. If appropriate medical and ophthalmic treatment is given in the early stages blindness can be prevented in at least 60% of cases. There is a clear case for screening; automated screening is desirable as manual screening is time-consuming, requires a high degree of training, skill and experience, and is susceptible to subjective observer variation and error.

Retinopathy is primarily a disease of the blood vessels, and the first damage occurs there, although it may not be externally visible at first [13] [17] [16]. As

retinopathy progresses, the rarefaction of small arterioles and capillaries leads to dilatation of larger arterioles. Elevated blood flow rates occur in pre-retinopathy and early retinopathy ([10]), whereas lower rates (in conjunction with increased diameter of large arterioles) have been reported in advanced retinopathy ([6]). As diabetic retinopathy advances to the dangerous "pre-proliferative" stage a number of indicators occur: macular oedema, cotton-wool spots, and venous beading. Venous beading has been described as "the most powerful predictor of conversion to proliferative retinopathy [9]." Another important complication of retinopathy is neovascularization; the new vessels have a characteristically tortuous shape, and several authors have designed methods to detect this tortuosity [14] [3] [11]. The vascular network is also useful in image registration, due to its wide distribution over the retina and relative stability over time [7].

Given the significance of the retinal network, it is unsurprising that a number of authors have suggested methods to segment it. The most popular approach is to use a specially-designed linear filter. This is typically an elongated filter with a gaussian or similar profile, which is rotated through a number of angles and convolved with the image [12] [4] [19] [5]. At any given pixel the highest response filter is selected, and defines the angle of the vessel. These maximum filter strengths may be thresholded to provide an image.

The basic convolution approach suffers a number of drawbacks. First, the retinal background is quite variable, and the contrast with the darker vessels quite low – the vessels are typically about 10 intensity shades deeper than the local surroundings in our images, whereas the background varies by up to 100 intensity levels across the image. Consequently, one needs to compensate, typically by filtering the image with a mean or median filter and subtracting (unsharp masking). The rapid rate of background intensity change means that the unsharp mask cannot be too big (e.g. a radius 50 disk on a 605 × 700 image), which makes it prone to produce artifacts, especially around light objects such as exudates, reflection artifacts and the optic nerve head. Second, the algorithm is very prone to false response around light objects even without unsharp masking, as the contrast between light objects and the retinal background is much greater (perhaps 50 intensity levels) than that between vessels and background. Consequently, and even though the rotating filter may respond only on one half to a light distractor edge, the response is stronger than that at a genuine

vessel. Third, the image is quite noisy, with fine vessels ranging down to sub-pixel level and therefore difficult to detect even by eye. Segmenting these very fine vessels is almost impossible. Some early papers on retinal segmentation do not make the difficulty of the task at all clear, as the authors use images of healthy retinæ [1], or do not use a representative range of pathologies [4].

A number of approaches have been suggested to improve the performance of segmentation. One possibility is to exploit the global connectivity of the vascular network, by searching from some start point. The key problem with this is that large sections of network might be missed if there is some apparent gap, possibly due to poor imaging. Hoover *et al.*, address this issue by constructing a filtered response image, thresholding at a high level, and then progressively lowering the threshold, acquiring segments that are attached to the high threshold segments, and using measures of the additional regions to decide whether or not to accept them. The approach to combining filtering and exploitation of connectivity constraints appears to be very effective.

2 The “Tramline” Filter

In this section we present an algorithm to detect and segment retinal vessels with high reliability and low false-positive rates. The algorithm works primarily by image filtering, with secondary exploitation of global connectivity.

A successful vascular segmentation algorithm must exploit the structure of the vessels; in particular, the characteristic cross-vessel profile (light to either side and darker in the middle, but with a possible light specular highlight in the very center) together with the consistency in the orthogonal direction. As described in the previous section, a linear convolution kernel which matches this profile shape is not adequate, as it responds more strongly to bright distractors than to actual vessels – this occurs as the filter does not impose the constraint of symmetry across the profile line.

We introduce a simple and effective non-linear filter. We reason as follows: if we sample along a line with the same orientation as the vessel, and within the vessel, then the line should be uniformly dark. If we sample along parallel lines with sufficient displacement to either side, then those lines should be lighter. We therefore consider the contrast between a central line and satellite “tram-lines.” If we wish to locate only vessels above a certain width, we may use tram-lines in the internal region too, and these may even span a central light reflex.

What statistic should we use to measure the contrast? The mean contrast on each tramline is not reliable – outliers injected by distractors can too easily distort results, leading to false responses. We require a non-linear filter. An idealized approach is to use a minimax algorithm: calculate the contrast as the *minimum* intensity level sampled on the (supposedly light) outer tram-lines minus the *maximum* intensity on the (supposedly dark) inner tram-

line(s). If this difference has high magnitude, then the entire interior tramline is darker than the entire outer tramline. This is equivalent to the difference of the grey-scale morphological erosion and dilation using structuring elements that match the outer and inner tram-lines respectively, and is thus closely related to the top-hat transform.

In reality, this approach is too sensitive to noise – a single outlying intensity value has a disproportionate effect on the filter. In seeking a robust equivalent, we use the somewhat neglected order statistic filtering [8]: we sort intensity values within the inner and outer filters, choosing the third darkest and third lightest values respectively, and take the difference.

The filter is reasonably robust against noise, and particularly insensitive to distractors. The orientation-sensitivity ensures that it seldom responds strongly off blood-vessels, and the order-statistic on the pair of outer tram-lines makes it non-responsive to the one-sided contrast boundaries found next to white lesions, reflection artifacts and the optic nerve head. A potential disadvantage of the filter is that it fails to respond when the outer tramlines intersect a side branch; consequently junctions are not detected. However, this may turn out to be a positive advantage when attempting to recover “clean” blood vessel segments between junctions, which is required for some applications such as venous beading detection.

The tramline filter is applied at a number of orientations, and the strongest response selected at each pixel; the resulting strength map is thresholded. The reliance on differencing means that, in contrast to most alternative filtering algorithms, no prior contrast equalization is necessary. As with all filters, the resulting image is noisy, and needs to be cleaned up. In particular, many very small isolated false positive segments occur. We also wish to identify, roughly, the center line of each vessel, to initialize our vascular measurement algorithms. Our clean-up procedure therefore consists of: a) speckle removal, to increase the efficiency of later stages; b) morphological thinning, to reduce vascular segments to single pixel-wide segments; c) removal of small segments; d) clean-up of remaining segments, including hole-filling and pruning of small side-spurs.

We tested the tram-line filter against the STARE algorithm [12]. STARE is a particularly sophisticated variant of the linear filtering approach, which includes contrast equalization, a flood-fill like algorithm to exploit structural connectivity of the vascular network, and a liberal growth / conservative prune approach to exploration. The STARE project provide a publicly-available data set of 20 images, including ground-truth labelling. The data set is designed to demonstrate a range of pathologies and other defects, including artifacts such as severe camera blur; consequently, it is not representative of image distribution in a typical medical application. It does, however, present a stern test for vascular segmentation algorithms.

Hoover *et al.* report performance figures of 90% sensitivity and 96% specificity on this dataset, on a per-pixel basis; i.e. the figures are derived by comparing the al-




Figure 1: Vessel centerlines on part of Stare image 002. Top: Stare algorithm. Bottom: Tramline algorithm.

gorithm's final map with the ground-truth on a pixel by pixel basis. There are two problems with this evaluation.

First, it is actually trying to assess two different issues: whether the vessels have been located, and whether their widths have been correctly assessed. There is no correction for the varying width of vessels, and so a wide vessel (which is relatively easy to locate) contributes disproportionately to the sensitivity measure. Measuring the width of a vessel is more usefully treated as a separate stage, and so we prefer a statistic based only on length.

Second, the specificity measure is extremely insensitive. As the majority of the image is non-vessel (approximately 90%) it follows that it is easy to obtain an apparently high specificity measure even when there are many errors (e.g. if the algorithm over-estimated the vascular area by 100%, the specificity would still be nearly 90%). A more useful statistic is the *positive predictive value* (PPV), which gives the proportion of identified vessel length which is true vessel length. The extent to which this falls below 100% indicates the *excess length* – amount of extra vessel found. This corresponds more closely to a subjective evaluation of the amount of error.

To determine the actual vascular length, we thinned the STARE ground truths, to obtain a pixel map that approximates the center line of the vessels. We approximate the total vascular length by counting the number of pixels in this image (with no adjustment for vessel angle). To assess the performance of an algorithm, we thin its output map, dilate the non-thinned ground truths with a 3×3 structuring element, and count coincident pixels - this tells us if the algorithm's foreground pixels lie within or right next to the ground truth vessel pixels. Based on this definition, we calculate the STARE algorithms sensitivity at 57%, and its positive predictive value at 92.2%.

We followed the same procedure for the tram-line algorithm, and adjusted its control parameters to achieve approximately the same sensitivity rate as Stare. With these settings, the algorithm successfully locates 57.2% of the vascular network, with positive predictive value 97.2% – a substantially lower error rate.

3 Vascular Segmentation

The tram-line algorithm extracts a pixel map consisting of sparse lines roughly along vessel centers. This map is further processed to produce a vascular segment map, each section identifying a vessel segment's precise centerline and diameter along the length of the vessel. This stage requires accurate calculation of the diameter of the vessel, resampling of a spline to define the centerline, and the repair of problems such as the presence of short spurs and branches, bridging missed pixels, and misalignment of the centerline due to strong specular highlights.

We first morphologically thin the tramline map, then identify segments. On the thinned image vessel pixels have one neighbor at end-points, two neighbors along vessels, and three or more at junctions and spurs. We remove the junction pixels, and process the remaining well-formed vessel-segments by pixel-to-pixel following, starting at end-points.

Bridging across junctions and small breaks is achieved by calculating the angle of the vessel segment at each clean segment end-point. We construct a pairwise bridging cost matrix, as follows. Consider two end points, \mathbf{e}_i and \mathbf{e}_j . Define the angle of the vessels at the end-points, θ_i and θ_j , as the angle of the lines traced to the end-point from the pixel two places earlier on that segment. Let $\phi_{i,j}$ be the angle of the line joining the two end-points, and $\|\mathbf{e}_i - \mathbf{e}_j\|$ be the distance between the end-points. Then the cost measure, $C(\mathbf{e}_i, \mathbf{e}_j) = \|\mathbf{e}_i - \mathbf{e}_j\| \cdot \sin(\theta_1 - \phi_{i,j}) \cdot \sin(\phi_{i,j} - \theta_2)$. Only end-points within the threshold angle, $\lambda = \pi/4$ of the segment direction vector, and maximum distance, $d = 5$, are considered for bridging. We seek a minimum cost assignment of bridges using a simple greedy algorithm, which iterates through the set of end-points. Match costs are initialized to infinity. Each end-point, \mathbf{e}_i is compared to each of its feasible matches, \mathbf{e}_j . If the cost is lower than the current match costs for both \mathbf{e}_i and \mathbf{e}_j , they are matched and any previous match is broken. The algo-

rithm is guaranteed to converge rapidly, and is effective at resolving situations such as two crossing vessels (producing four proximate end points) and branching vessels (three proximate end points).

The resulting segments are suitable for initializing a measurement algorithm such as that described in [2].

4 Conclusion

We have described a new algorithm for the segmentation of blood vessels from retinal images. The algorithm uses a difference of non-linear order statistic filters across two regions of interest: a line oriented along the vessel, and a “tram-line” with rails to either side of this line. The filter is rotated to various angles, and probes the image; the response is thresholded and morphologically cleaned up. Experiments on the publicly-available Stare image set indicate that performance is extremely good – the method is particularly robust against exudate edge distractors. After thinning, the vessel map is converted into segments using line following and bridging algorithms; the resulting segments are suitable for submission to algorithms for vascular measurement, and ultimately for the detection of vascular diseases.

Acknowledgments

This project was sponsored by Diabetes UK, under grant no. BDA:RD00/0002033.

References

- [1] K. Akita and H. Kuga. A computer method of understanding ocular fundus images. *Pattern Recognition.*, 15:431–443, 1982.
- [2] B. Al-Diri and A. Hunter. A ribbon of twins for extracting vessel boundaries. In *Proc. EMBEC 2005 Conference, Special Session SS10 on Retinal Imaging*, 2005.
- [3] J.J. Capowski, J.A. Kylstra, and S.F. Feedman. A numeric index based on spatial frequency for the tortuosity of retinal vessels and its application to plus disease. *Retinopathy of Prematurity*, 15:490–500, 1995.
- [4] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Trans. Medical Imaging*, 8:263–369, 1989.
- [5] R. Collorec and J.L. Coatrieux. Vectorial tracking and directed contour finder for vascular network in digital subtraction angiography. *Pat. Rec. Let.*, 8:353–358, 1988.
- [6] M.H. Cuypers, J.S. Kasanardjo, and B.C. Polak. Retinal blood flow changes in diabetic retinopathy measured with the heidelberg scanning laser doppler flowmeter. *Graefes Arch Exp Ophthalmol*, 238:935–41, 2000.
- [7] Michael H. Goldbaum, Valentina Kouznetsova, Brad L. Cote, William E. Hart, and Mark Nelson. Automated registration of digital ocular fundus images for comparison of lesions. *SPIE Vol. 1877 Ophthalmic Technologies III*, 1877:94–99, 1993.
- [8] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing (Second Edition)*. Prentice Hall, Upper Saddle River, NJ, 2001.
- [9] P.H. Gregson, Z. Shen, R.C. Scott, and V. Kouzousek. Automated grading of venous beading. *Computers and Biomedical Research*, 28:291–304, 1995.
- [10] J.E. Grunwald, J. DuPont, and C.E. Riva. Retinal haemodynamics in patients with early diabetes mellitus. *Br. J. Ophthalmol*, 80:327–31, 1996.
- [11] W.E. Hart, M. Golbaum, B. Cote, P. Kube, and M.R. Nelson. Measurement and classification of retinal tortuosity. *Int J Med Inf*, 53:239–52, 1999.
- [12] A. Hoover, V. Kouznetsova, and M. Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Transactions on Medical Imaging*, 19:203–210, 2000.
- [13] T. Kawagishi, Y. Nishizawa, M. Emoto, T. Konishi, K. Maekawa, S. Hagiwara, Y. Okuno, H. Inada, G. G. Isshiki, and H. Morii. Impaired retinal artery blood flow in iddm patients before clinical manifestations of diabetic retinopathy. *Diabetes Care*, 18:1544–9, 1995.
- [14] W. Lotmar, A. Freiburghaus, and D. Bracher. Measurement of vessel tortuosity on fundus photographs. *Graefes Arch. Clin. Exp. Ophthalmol.*, 211:49–57, 1979.
- [15] A. Manivannan, P.F. Sharp, R.P. Phillips, and J.V. Forrester. Digital fundus imaging using a scanning laser ophthalmoscope. *Physiol Meas*, 14:43–56, 1993.
- [16] M. Quigley and S. Cohen. A new pressure attenuation index to evaluate retinal circulation. a link to protective factors in diabetic retinopathy. *Arch Ophthalmol*, 117:84–9, 1999.
- [17] S.M. Rassam, V. Patel, and E.M. Kohner. The effect of experimental hypertension on retinal vascular autoregulation in humans: a mechanism for the progression of diabetic retinopathy. *Exp Physiol*, 80:53–68, 1995.
- [18] D.E. Singer, D.M. Nathan, H.A. Fogel, and A.P. Schachat. Screening for diabetic retinopathy. *Ann Intern Med*, 116:660–71, 1992.
- [19] B.D. Thackray and A.C. Nelson. Semi-automatic segmentation of vascular network images using a rotating structuring element (rose) with mathematical morphology and dual feature thresholding. *Pattern Recognition.*, 3:385–392, 1993.