

Locating the Optic Nerve in a Retinal Image Using the Fuzzy Convergence of the Blood Vessels

Adam Hoover
Electrical and Computer Engineering Department
Clemson University
Clemson, SC 29634-0915
ahoover@clemson.edu

Michael Goldbaum
Department of Ophthalmology
University of California, San Diego
La Jolla, CA 92093
mgoldbaum@ucsd.edu

This work was supported by the NIH NEI grant #EY13235-11.

Submitted September 2001, revised June 2002, to the IEEE Trans. on Medical Imaging.

Abstract

We describe an automated method to locate the optic nerve in images of the ocular fundus. Our method uses a novel algorithm we call fuzzy convergence to determine the origination of the blood vessel network. We evaluate our method using 31 images of healthy retinas and 50 images of diseased retinas, containing such diverse symptoms as tortuous vessels, choroidal neovascularization, and hemorrhages that completely obscure the actual nerve. On this difficult data set, our method achieved 89% correct detection. We also compare our method against three simpler methods, demonstrating the performance improvement. All our images and data are freely available for other researchers to use in evaluating related methods.

Key words: retinal imaging, optical image processing, blood vessels, segmentation



Figure 1: In a healthy retina, the optic nerve has a readily identifiable size, shape, color, and location relative to the blood vessels.

1 Introduction

The optic nerve is one of the most important organs in the human retina. The central retinal artery and central retinal vein emanate through the optic nerve, supplying the upper layers of the retina with blood. The optic nerve also serves as the conduit for the flow of information from the eye to the brain. Most retinal pathology is local in its early stages, not affecting the entire retina, so that vision impairment is more gradual. In contrast, pathology on or near the nerve can have a more severe effect in early stages, due to the necessity of the nerve for vision [16] (pg. 719).

The optic nerve can be examined in a viewing of the retinal fundus. The portion of the nerve that is visible in such a view is called the optic disc, referring to the two-dimensional appearance of the part of the nerve that is visible. Fundus imaging is a common clinical procedure used to record a viewing of the retina. This image may be used for diagnosis, treatment evaluation, and the keeping of patient history. In this work we describe a process to automatically locate the optic nerve in a retinal image. Such a tool could be used for automated patient screening, eye orientation tracking, image sequence registration, and automated measurements for treatment evaluation or diagnosis.




Figure 2: A retina containing lesions of the same brightness as the nerve.

Figure 1 shows the green plane of an example retinal image. The optic nerve appears towards the left side of this image as a circular area, roughly one-sixth the width of the image in diameter, brighter than the surrounding area, as the convergent area of the blood vessel network. In an image of a healthy retina, all these properties (shape, color, size, convergence) help contribute to the identification of the nerve.

However, these features show a large variance that makes simple detection methods brittle, particularly in the presence of retinal disease. Figure 2 shows an image of a retina containing drusen. The brightness of these lesions overlaps the brightness of the nerve, so that using brightness as a lone feature for detection is difficult. Figure 3 shows an example of a swollen optic nerve, where the circular shape and size are distorted. Figure 4 shows a retina exhibiting choroidal neovascularization and subretinal fibrosis, with a bright circular lesion that appears similar to a nerve. Figure 5 shows a retina exhibiting central artery and vein occlusion, where the nerve is completely obscured by hemorrhaging. These cases demonstrate the difficulty in optic nerve detection, and support the need for a robust method that can work in the presence of a variety of retinal diseases.

In Figures 1-5, the only consistently visible property of the optic nerve is that it is the convergent point of the blood vessel network. Therefore we base our method of optic nerve




Figure 3: A swollen nerve, showing a distorted size and shape.




Figure 4: A bright circular lesion that looks similar to an optic nerve.

detection upon finding the convergence of the blood vessels. In the absence of a unique and strongly identifiable convergence, our method uses brightness as a secondary feature for optic nerve detection. We also compare the success of each of these features used alone.

In previous work we described a method to compute a binary segmentation of blood vessels [10]. The parameters of this segmenter can be adjusted to produce vessel segmentations at different scales. Figure 6 shows an example, where the segmentation of the vessels was produced using two different parameter sets. In this work we use multiple vessel segmentations of the same image in order to reinforce the detection of convergent points. The idea is



Figure 5: A nerve that is completely obscured by hemorrhaging.

that the convergence should be detectable using vessel segmentations at different scales.

To find the vessel network convergence, we describe a novel algorithm we call fuzzy convergence. The algorithm is a voting type method that works in the spatial domain of the image. The input to the algorithm is a binary segmentation of the blood vessels. Each vessel is modeled by a fuzzy segment, which contributes to a cumulative voting image. The output from the algorithm is a convergence image, which is thresholded to identify the strongest point(s) of convergence.

We test our method on 31 images of healthy retinas and 50 images of diseased retinas. We report the success of our method to detect the optic nerve using fuzzy convergence alone, and in conjunction with using brightness as a salient feature.

1.1 Related work

The problem of optic nerve detection has rarely received unique attention. It has been investigated as a precursor for other issues, for example as identifying a starting point for blood vessel segmentation [23, 24]. It has also been investigated as a byproduct of general retinal image segmentation, for instance into separate identifications of arteries, veins, the nerve, the fovea, and lesions [1, 7, 13, 17]. Here we review these related works.




Figure 6: Binary segmentation of blood vessels, of image in Figure 1.

In [13] a method is presented to segment a retinal image into arteries, veins, the optic disk, the macula, and background. The method is based upon split-and-merge segmentation, followed by feature based classification. The features used for classification include region intensity and shape. The primary goal of the paper was vessel measurement; the nerve was identified only to prevent its inclusion in the measurement of vessels. Ten healthy retinas and ten retinas with arterial hypertension were used for experiments. Quantitative results for nerve detection were not provided. A similar approach was taken in [7], in which the segmentation was accomplished using matched spatial filters of bright and dark blobs. Quantitative results for nerve detection were not provided.

In [17] a method is presented to segment a retinal image into vessels, the nerve, the fovea, scotomas, and subretinal leakages. Nerve detection is based upon the transform of gradient edges into a Hough space describing circles. The search is restricted to one-third of the image based upon apriori knowledge of the expected general location of the nerve. Eleven retinas with age-related macular degeneration (ARMD) were used for experiments. In 10 out of 11 cases, the nerve was successfully detected. In [1] a method is presented to segment a retinal image into arteries, veins, the optic disk, and lesions. Nerve detection is

based upon tracking the vessel network to a common starting point. The tracking process uses the angles between vessels at branch points to identify the trunk. A result is shown for two images; quantitative results were not provided.

In [24] a method is described to detect the optic nerve as a starting point for tracking-based segmentation of blood vessels. Nerve detection is based upon locating the brightest region in a restricted third of the image. Three retinal images were used for experiments. Although the health of these retinas is not reported, the nerves show readily identifiable properties similar to the nerve shown in Figure 1 of this paper. In all 3 cases, the nerve was successfully detected. A similar method is described in [20].

In [23], the nerve was detected using the transform of gradient edges into a Hough space describing circles. Quantitative results were not provided. A similar method is described in [2].

In [14] a method is described to track eye movements based upon retinal features, including the nerve. Nerve detection is accomplished using basic thresholding for bright areas. Quantitative results were not provided.

In [5] a method is described to detect the optic nerve in a retinal image. Nerve detection is based upon three features: bright circularity, vertical vessel detection, and discrete vessel convergence. Bright circle detection is accomplished using least-median-of-squares fitting of gradient points. Vertical vessels contribute to bins distributed horizontally across the image; the idea is that the primary four vessels emanate near-vertically from the nerve in a standard retinal image. Discrete vessel convergence is accomplished using a method similar to the fuzzy convergence reported herein, except that each vessel is modeled by a single line in a coarse voting space. In later testing [3], 133 retinas containing various health disorders were used for experiments. In 106 cases (80%) the nerve was successfully detected.

Our work differs from previous methods in that we use blood vessel convergence as the primary feature for detection. We test our method on 81 images showing a variety of retinal

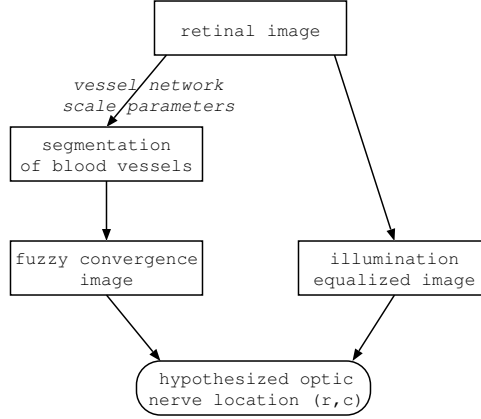


Figure 7: Outline of our method for optic nerve detection.

diseases, confusing lesions and manifestations. On this difficult data set, our method achieves an 89% correct detection rate.

2 Methods

Our method for optic nerve detection is outlined in Figure 7. It is based primarily upon a novel algorithm we call fuzzy convergence. This algorithm identifies the optic nerve as the focal point of the blood vessel network. In the absence of a strong convergence, our method identifies the optic nerve as the brightest region in the image after illumination equalization. Section 2.1 describes fuzzy convergence, Section 2.2 describes illumination equalization, and Section 2.3 describes our hypothesis generation and consensus.

2.1 Fuzzy convergence

The blood vessel segments may be modeled by line segments. The problem of finding the convergence of the vessel network may then be modeled as a line intersection problem. Our approach differs from traditional methods in that we model each vessel with a line segment (of finite length), as opposed to a line (of infinite length). Before describing our approach we review the least squares and Hough space families of solutions, which are the two most traditional approaches to line intersection.

Least squares solutions (see [18]) and related approaches work by minimizing the mean square (or some other function) of errors between a set of equations and a single solution. In physical terms, a least squares solution to the intersection of a set of lines finds the point simultaneously minimally distant from all the lines. An assumption common to these methods is that the data set being fitted is uniformly distributed around the optimal single solution. Subsets of data that do not meet this criterion are termed outliers, and cause wrong solutions. Several minimization-based solutions that overcome various types and amounts of outlier data have been proposed (see [15]). Generally speaking, these methods work by partitioning the data set into inliers and outliers, so they are computationally expensive, and require $> 50\%$ inliers.

Hough space methods (see [12]) and related approaches work by transforming data points from an image space to a quantized parameter space. Each data point in effect votes for a finite number of parameter sets. The parameter set with the highest total vote is taken for the solution. In physical terms, a Hough transform solution to the intersection of a set of lines finds the bin (point at the resolution of the Hough space) through which the largest number of lines passes. Generally speaking, these methods model line-like shapes with lines (i.e., of infinite length), so any sense of convergence contributed by the endpoints of the line-like shapes is lost. Simple Hough spaces also tend to be very sensitive to the chosen resolution.

Fuzzy convergence [9] is a voting-based method. The voting takes place on the integer grid of the original image. Each line-like shape is modeled by a fuzzy segment, whose area contributes votes to its constituent pixels. The summation of votes at each pixel produces an image map where each pixel contains a value proportionate to its strength of convergence. The map is then blurred (to reduce the effects of quantization) and thresholded to produce one (or more, if desired) points of strongest convergence. The proposed method runs in $O(n)$ time, where n is the number of line-like shapes. It does not require any amount of inliers;

instead, an absolute threshold for strength may be applied to determine if any area should be deemed convergent.

2.1.1 Fuzzy segment model

A line segment is defined by its two endpoints (x_1, y_1) and (x_2, y_2) . In this section a fuzzy segment model is proposed. The fuzzy segment, henceforward denoted as \mathbf{F} , is defined by a set of parametric line segments:

$$\begin{aligned} x(t) &= x_1 + r \cos(\alpha + \theta) + (x_2 - x_1 - 2r \cos \theta \cos \alpha)t \\ y(t) &= y_1 + r \sin(\alpha + \theta) + (y_2 - y_1 - 2r \cos \theta \sin \alpha)t \end{aligned} \quad (1)$$

where

$$\begin{aligned} 0 &\leq t \leq 1 \\ 0 &\leq \theta \leq 2\pi \\ 0 &\leq r \leq R \end{aligned} \quad (2)$$

The amount of “fuzziness” is controlled by the parameter R , which, at zero, reduces the fuzzy segment to the single line segment from (x_1, y_1) to (x_2, y_2) . The parameter α corresponds to the orientation of that single segment, and is computed as

$$\alpha = \frac{\pi}{2} - \arctan \frac{y_2 - y_1}{x_2 - x_1} \quad (3)$$

The fuzzy segment \mathbf{F} defines a set of segments of orientations and lengths ranging about a line segment. The motivation for the fuzzy segment is best demonstrated through its visualization. Figure 8 illustrates the enumeration for the subset of \mathbf{F} where $r = R$. It may be visualized as a moving line segment, whose initial endpoints are marked in Figure 8 at $\theta = 0$. As theta moves from zero to 2π , the endpoints trace the boundaries of circles, one




Figure 8: Fuzzy segment model.

in the clockwise direction, the other counter-clockwise. By starting $\theta = 0$ at α , the shape of the fuzzy segment remains invariant to orientation.

Relevant previous work in fuzzy geometry (see [4, 8, 19]) deals mainly with fuzzy line models. The fuzzy segment \mathbf{F} is proposed as a model for the area in which an observed line-like shape contributes to a sense of convergence. The model proposes that the contribution of the midpoint of the line-like shape is more compact than the endpoints. The width of the fuzzy segment at its midpoint is

$$\frac{R^2}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \quad (4)$$

while the width at its endpoints is $2R$. The model also proposes that a line-like shape generally only contributes to a convergence in its "near" neighborhood. The term "near" is subjectively applicable and implies the area within some perceptually relevant distance surrounding the line-like shape. The model allows for an interpretation of what is near via the parameter R .

2.1.2 Convergence image

Given a binary input image, like the one depicted in Figure 6, a process for finding convergences works as follows:

1. Thin the image (for instance, using the algorithm given in [12], pg. 59).
2. Erase (relabel as background) all branchpoints, breaking up the entire foreground into segments that contain two endpoints each. In a thinned image, endpoints may be discovered as any pixel for which a traverse of the eight bordering pixels in clockwise order yields only one foreground-to-background transition. Similarly, branchpoints may be discovered as any pixel for which the same traverse yields more than two transitions.
3. Extend each segment a distance of R pixels in both directions. The extension is done along the vector made by the segment's endpoints. For our experiments, we used $R = 15$ (the average distance between vessels in the nerve in our data set).
4. Model each segment, via its two extended endpoints, with a fuzzy segment \mathbf{F} . The image area covered by \mathbf{F} may be found by enumerating \mathbf{F} at $r = R$ with suitable discretizations of θ and t . For the experiments reported herein, θ was enumerated to produce unique pixel coordinates for endpoints. For each θ , t was also enumerated to produce unique pixel coordinates. Since the line segments at different θ 's overlap, a second binary image is used to keep track of which pixels are found to lie in \mathbf{F} . This image is cleared after each fuzzy segment enumeration and voting is completed.
5. For each pixel enumerated in \mathbf{F} , a vote is cast. The image used to tally these votes is termed the convergence image. Several voting functions were explored, including weighting by segment size, and weighting by distance from midpoint. Interestingly, no function seemed to work better, overall, than the simplest: equal voting (+1) for all pixels in each \mathbf{F} .




Figure 9: Fuzzy convergence of blood vessels shown in Figure 6.

6. Smooth the convergence image to identify the center of the peak of convergences.

Figure 9 shows the smoothed convergence images for the vessel segmentations shown in Figure 6 superimposed on the thinned segments. The measures of convergence have been normalized so that the highest vote appears darkest. For our experiments, we used an 11×11 mean filter for smoothing.

2.2 Illumination equalization

The illumination in a retinal image is uneven [11]. The imbalance is primarily due to an optical aberration called vignetting (see for example [21], pg. 490). Vignetting is the result of an improper focusing of light through an optical system. The result is that the brightness of the image generally decreases radially outward from near the center of the image. Figure 10(a) shows a retinal image demonstrating uneven illumination. Figure 10(b) shows the local average intensity, displayed using random shades of grey to show iso-intensity contours.

In our case, the imbalance is complicated by the optics of the system used to capture the image. A retinal image is captured by viewing the inner rear surface of the eyeball through the pupil. The lens of the camera works in conjunction with the lens of the eyeball to form

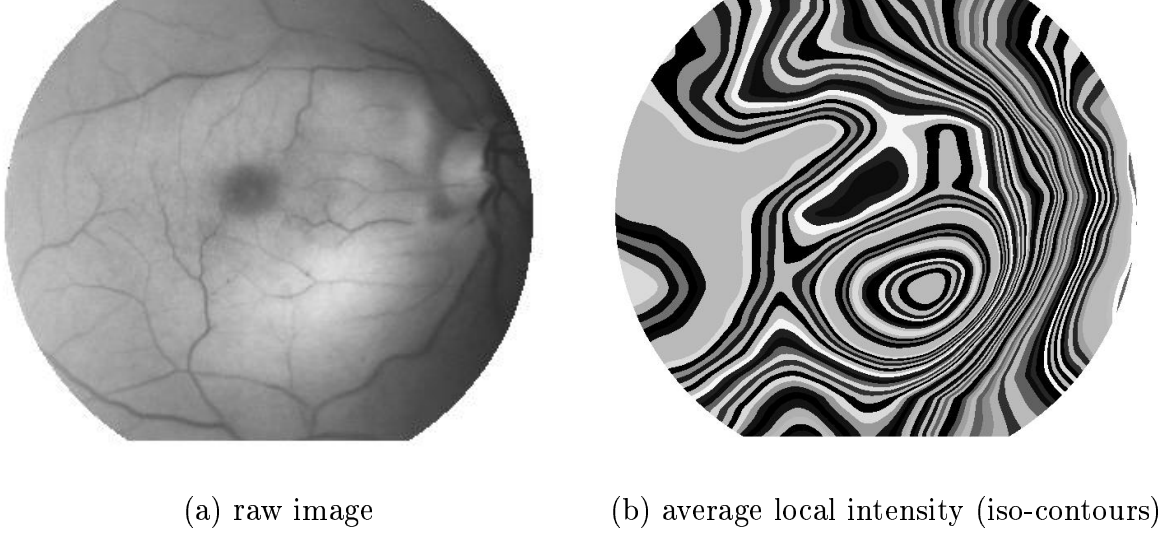


Figure 10: The illumination across the image is uneven, due to the imaging process.

the image. Since the position of the eye relative to the camera varies from image to image, the exact properties of the vignetting also vary from image to image.

The uneven illumination hinders absolute interpretation of the intensities in the image. In a healthy retina, the optic nerve is usually the brightest feature. Applying a simple high threshold to an image of a healthy retina should yield pixels inside the optic nerve. However, a retinal image is often captured so that the fovea appears mid-image, with the nerve to one side. Because of the vignetting, the nerve may appear darker than areas central to the image (as in Figure 10).

In order to undo the vignetting, we apply an illumination equalization to the image. Each pixel $I(r, c)$ in the image is adjusted as follows:

$$I_{eq}(r, c) = I(r, c) + m - A(r, c) \quad (5)$$

where m is the desired average intensity (128 in an 8-bit greyscale image) and $A(r, c)$ is the local average intensity. The local average intensity is computed independently for each pixel as the average intensity of the pixels within an $N \times N$ window. The window size N is variable, so that averages near the border of the image use the same number of pixels (between 30 and 50 in our experiments) as averages in the center of the image. The local average intensities




Figure 11: The result of illumination equalization on the image in Figure 10.

are also smoothed using the same windowing, to prevent blurring the image features. This process can be implemented using a sliding window algorithm (see for instance [21], pg. 75), so that the computations are fast. Figure 11 shows the result of illumination equalization upon the image in Figure 10.

2.3 Hypothesis generation

Both the fuzzy convergence and illumination equalization algorithms produce images in which the brightness of a pixel is indicative of the presence of the optic nerve. In order to generate a hypothesis of nerve location, either image is thresholded to identify the brightest 1,500 pixels (or darkest 1,500 pixels as depicted in Figure 9). The quantity 1,500 was chosen based upon the observation that this is the average number of pixels that “stand out” in an optic nerve in our images (our imaging resolution is 605×700 pixels on a 35° field-of-view). These pixels are grouped into regions using standard 8-connected component analysis. Any regions within 5 pixels of each other are grouped using standard morphological operators. This last step groups areas that may be separated by blood vessels, as commonly occurs in the nerve.

The regions are sorted and separated into two classes, according to region size, using Fisher’s linear discriminant (see for instance [6], pg. 114). This statistical test finds the

best separation of the regions into those that are “large” and those that are “small” or insignificant. It works as follows. The regions are sorted by size, and repeatedly partitioned into two sets, A and B . The initial partition has only the largest element (region size) in set A , and all other elements in set B . Each partitioning moves the largest remaining element from set B into set A , until there is only one element left in set B . Thus, for n total elements (regions) there will be $n - 1$ total partitions.

For each partition p , the discriminant statistic F_p is computed as

$$F_p = \frac{(\mu_A - \mu_B)^2}{\sigma_A^2 + \sigma_B^2} \quad (6)$$

where μ and σ are the mean and standard deviation of each set. The largest value of F_p indicates the best partition. If only one region passes this test (resides in set A for the largest value of F_p), its centroid is deemed a hypothesis of the optic nerve location. If more than one region passes this test, then the result is deemed inconclusive, and no hypothesis is generated.

3 Results

We test our methods on 31 images of healthy retinas and 50 images of retinas with disease (all the images shown in this paper are included in our test set). All these images were acquired using a TopCon TRV-50 fundus camera at 35° field-of-view, and subsequently digitized at 605×700 pixels in resolution, 24-bits per pixel (standard RGB). In this work, only the green band of the image was used (the red band of a retinal image tends to be saturated, and the blue band tends to be empty). The nerve is visible in all 81 images, although in 14 images it is on the border of the image so that only a fraction of the nerve is visible, and in 5 of the images the nerve is completely obscured by hemorrhaging. All our images and results may be viewed at [22].

We tested our methods in four variations:

1. Equalized brightness. This variation uses illumination equalization followed by hypothesis generation.
2. Fuzzy convergence, single scale. This variation applies fuzzy convergence to a single blood vessel segmentation (the same scale is used for all images), followed by hypothesis generation.
3. Fuzzy convergence, multiple scales. This variation applies fuzzy convergence to six segmentations of the blood vessels, each at a different scale (see Figure 6). The scales correspond to the first six of the ten parameter values given in [10]. The distances between the hypotheses generated for each scale are measured relative to the radius of the average optic nerve (60 pixels in our images). If more than half of the hypotheses (3 or more out of 6) are all within this distance, the centroid of these locations is hypothesized as the optic nerve location. If there is no such consensus, then the result for the image is deemed inconclusive, and no hypothesis is generated.
4. Fuzzy convergence and equalized brightness. Variations numbers 1 and 3 are combined. If the result from fuzzy convergence at multiple scales is deemed inconclusive, then the result from equalized brightness is used.

The center point of the nerve in each image was recorded manually, for ground truth. A nerve is considered successfully detected if the hypothesis generated by the automated method is within the optic nerve, measured as within 60 pixels of the ground truth location. The nerve detection is considered unsuccessful if either the hypothesized location is wrong, or if our method did not produce a hypothesis (see “inconclusive”, above). Figure 12 shows the results for all the images shown in this paper, all of which are considered successes. Table 1 tabulates the success rates of the methods on all 81 images.

The results for the equalized brightness method demonstrate the limit on the usefulness of brightness as a lone feature for nerve detection (62%). They also demonstrate that the




Figure 12: Result on each image shown in this paper.

brightness of the nerve is more easily confused in diseased retinas than in healthy retinas (52% vs. 77%).

The results for fuzzy convergence at a single scale show that the convergence of the vessel network is a more stable feature of the nerve than the brightness (74% vs. 62%). They also demonstrate that the visibility of the convergence is apparently independent of the health of the retina (74% vs. 74%).

The results for fuzzy convergence at multiple scales shows that the persistence of this feature at multiple scales of vessels improves the detection of the nerve (79% vs. 74%).

Finally, the results for fuzzy convergence at multiple scales in combination with equalized brightness shows the highest performance overall (89%), and complete success on all our healthy retina test cases (100%).

In the cases where our method failed, the cause was usually a convergence of blood vessels around a bright lesion. In these cases, an additional operator that searched for circularity and appropriate size of the bright region (hypothesized nerve location) may improve performance.

	equalized brightness	fuzzy convergence, single scale	fuzzy convergence, multiple scales	fuzzy convergence and equalized brightness
healthy retinas	77%	74%	87%	100%
diseased retinas	52%	74%	74%	82%
all retinas	62%	74%	79%	89%

Table 1: Results of our methods on 31 healthy retinas and 50 diseased retinas.

4 Conclusions

We have presented novel methods to automatically locate the optic nerve in a retinal image. Our methods use the convergence of the blood vessel network as the primary feature for detection, in conjunction with the brightness of the nerve as a secondary feature. We have tested our methods on 31 images of healthy retinas and 50 images of diseased retinas, exhibiting a wide variety of lesions and confusing manifestations. On this difficult data set, our methods successfully detect the nerve in 89% of the cases, and in 100% of the healthy cases. We believe this improves upon all previously published methods, although few of these previous works reports quantitative results against which we can compare. All of our images and results are freely available [22] for other researchers interested in related work.

Our method for finding the convergence of the blood vessel network is based upon a novel algorithm we call fuzzy convergence. Unlike least-squares and Hough space based solutions, fuzzy convergence uses the endpoints of the linear shapes (in our case blood vessel segments) to help find the solution. Identifying the intersection of a number of lines (a convergence) is a fundamental geometric problem, with applications ranging from astronomy to engineering, such as model fitting and prediction. We also believe that fuzzy convergence may find usefulness in other medical applications besides optic nerve detection. We intend to explore these issues in future work.

References

- [1] K. Akita and H. Kuga, “A Computer Method of Understanding Ocular Fundus Images”, in *Pattern Recognition*, vol. 15 no. 6, 1982, pp. 431-443.
- [2] S. Barrett, E. Naess and T. Molvik, “Employing the Hough Transform to locate the optic disk”, in *Biomedical Sciences Instrumentation*, vol. 37, 2001, pp. 81-86.
- [3] J. Boyd, “STARE Software Documentation: disk - Optic Disk Locator”, technical report, Visual Computing Laboratory, Department of Electrical and Computer Engineering, University of California, San Diego, September 1996.
- [4] B. B. Chaudhuri, “Some shape definitions in fuzzy geometry of space”, in *Pattern Recognition Letters*, vol. 12, September 1991, pp. 531-535.
- [5] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson and M. Goldbaum, “Automatic detection of the optic nerve in retinal images”, in the proceedings of the *IEEE International Conference on Image Processing*, 1989, pp. 1-5.
- [6] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*, Wiley & Sons, 1973.
- [7] M. Goldbaum, S. Moezzi, A. Taylor, S. Chatterjee, J. Boyd, E. Hunter and R. Jain, “Automated Diagnosis and Image Understanding with Object Extraction, Object Classification, and Inferencing in Retinal Images”, in the proceedings of the *IEEE International Conference on Image Processing*, 1996, pp. 695-698.
- [8] K. Gupta and S. Ray, “Fuzzy plane projective geometry”, in *Fuzzy Sets & Systems*, vol. 54, 1993, pp. 191-206.
- [9] A. Hoover and M. Goldbaum, “Fuzzy Convergence”, in the proceedings of *IEEE Conference on Computer Vision and Pattern Recognition*, 1998, pp. 716-721.

- [10] A. Hoover, V. Kouznetsova and M. Goldbaum, “Locating Blood Vessels in Retinal Images by Piecewise Threshold Probing of a Matched Filter Response”, in *IEEE Trans. on Medical Imaging*, vol. 19 no. 3, March 2000, pp. 203-210.
- [11] A. Hoover and M. Goldbaum, “Illumination Equalization of a Retinal Image Using the Blood Vessels as a Reference”, abstract in the annual meeting of the *Association for Research in Vision and Ophthalmology (ARVO)*, 2001.
- [12] R. Jain, R. Kasturi and B. Schunck, *Machine Vision*, McGraw-Hill, 1995.
- [13] A. Kaupp, A. Dolemeyer, R. Wilzeck, R. Schlosser, S. Wolf and D. Meyer-Ebrecht, “Measuring Morphological Properties of the Human Retinal Vessel System Using a Two-Stage Image Processing Approach”, in the proceedings of the *IEEE International Conference on Image Processing*, 1994, pp. 431-435.
- [14] H. Kawai and S. Tamura, “Eye Movement Analysis System Using Fundus Images”, in *Pattern Recognition*, vol. 19 no. 1, 1986, pp. 77-84.
- [15] P. Meer, D. Mintz and A. Rosenfeld, “Robust Regression Methods for Computer Vision: A Review”, in *Int’l Journal of Computer Vision*, vol. 6 no. 1, 1991, pp. 59-70.
- [16] C. Oyster, *The Human Eye: Structure and Function*, Sinauer Associates Publishing, 1999.
- [17] A. Pinz, S. Bernogger, P. Datlinger and A. Kruger, “Mapping the Human Retina”, in *IEEE Trans. on Medical Imaging*, vol. 17 no. 4, August 1998, pp. 606-619.
- [18] W. Press, S. Teukolsky, W. Vetterling and B. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edition, Cambridge Press, 1992.
- [19] A. Rosenfeld, “‘Geometric properties’ of sets of lines”, in *Pattern Recognition Letters*, vol. 16, May 1995, pp. 549-556.

- [20] C. Sinthanayothin, J. Boyce, H. Cook and T. Williamson, “Automated localisation of the optic disc, fovea, and retinal blood vessels from digital colour fundus images”, in *British J. of Ophthalmology*, vol. 83, 1999, pp. 902-910.
- [21] M. Sonka, V. Hlavac and R. Boyle, *Image Processing, Analysis and Machine Vision*, 2nd edition, Chapman & Hall, 1999.
- [22] STARE project website, <http://www.ces.clemson.edu/~ahoover/stare>.
- [23] S. Tamura, Y. Okamoto and K. Yanashima, “Zero-Crossing Interval Correction in Tracing Eye-Fundus Blood Vessels”, in *Pattern Recognition*, vol. 21 no. 3, 1988, pp. 227-233.
- [24] Y. Tolias and S. Panas, “A Fuzzy Vessel Tracking Algorithm for Retinal Images Based on Fuzzy Clustering”, in *IEEE Trans. on Medical Imaging*, vol. 17 no. 2, April 1998, pp. 263-273.