Fuzzy Convergence

Adam Hoover
Electrical & Computer Engineering Department
University of California, San Diego
La Jolla, CA 92093-0407
hoover@vision.ucsd.edu

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

Abstract

This work considers the problem of discovering areas of convergence of line-like shapes in an image. The motivating application is to use the convergence of the blood vessel network to automatically locate the optic nerve in an ocular fundus image. A fuzzy segment model is proposed, based on a conjecture that line-like shapes only contribute to a perception of convergence in their near neighborhood. Using this model, a voting-type method is described to compute a convergence image*, which can be searched for one absolute, or one or more relative, strongest points of convergence. Results are presented for twenty ocular fundus images, with a 65% success rate for finding the optic nerve.*

Key words: low-level processing, fuzzy geometry, medical image processing

1 Introduction

This work considers the problem of discovering areas of convergence of line-like shapes in an image. For this work, the term convergence is defined as a common near neighborhood of intersection of line-like shapes, that may be modeled by a point. A line-like shape is defined as any elongated area, with an appropriate minimum of global curvature, that may be modeled by a line. The motivating application is to locate the optic nerve in an image of the ocular fundus (interior of the eye). The optic nerve is an important feature for visual health diagnosis, screening, and image registration. The green plane of an example image is shown in Figure 1. The optic nerve appears in the right side of this image as a circular area, roughly one-seventh the width of the image in diameter, brighter than the surrounding area, as the convergent area of the blood vessel network. Although each of these properties (shape, color, size, convergence) contributes to one's perception of the identity and location of the nerve, this work pursues the use of convergence, for the following reasons.

Figure 1: An ocular fundus image.

The appearance of the nerve can vary in location, depending on the eye orientation relative to the camera. The nerve also varies substantially in color, and somewhat in size, depending on the subject and the subject's health. If hemorrhaging is present, the nerve can be entirely obscured, as in Figure 2. In cases like this (which may be categorized amongst the most important to detect), vessel convergence seems to be a necessary property for locating the optic nerve.

In other work [5], a process was developed to automatically segment blood vessels in fundus imagery. The result of this process on the image in Figure 1 is presented in Figure 3.

This binary image presents the appearance of blood vessels only. It does not contain any shape, size, or color cues for the optic nerve. Twenty similar images were shown to five people who had no familiarity with fundus imagery or this work. The subjects were asked

Figure 2: The nerve is obscured.

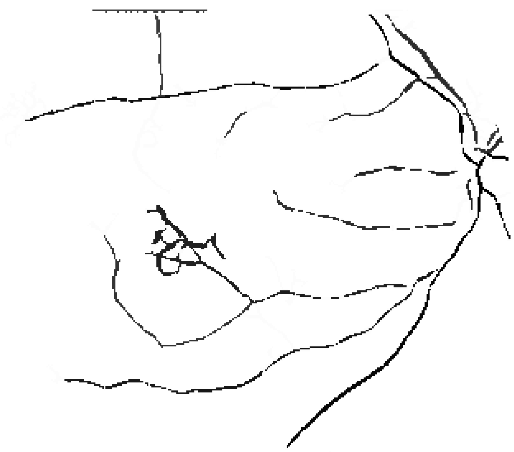

Figure 3: Segmentation of blood vessels, for Fig. 1.

to indicate the strongest visible points of convergence. The term convergence was explained by analogy to the joining of spokes in a bicycle wheel. The five subjects selected a point that was within the optic nerve 9/20, 11/20, 14/20, 16/20 and 19/20 times.[1] It is important to reiterate that this feat was accomplished by people who had never seen the original images, and who were not told what the line-like shapes represented. This suggests that vessel network convergence can be a sufficient property for locating the optic nerve.

This work presents a method to automatically find points of convergence in images like that shown in Figure 3. The solution presented is a voting-type method. The input is a binary image marking pixels that belong to line-like shapes. The output is an image of the same spatial resolution, where each pixel contains a value proportionate to the strength of convergence of the shapes at that pixel. This image may then be thresholded for one absolute, or one or more relative locations of strongest convergence. Although this method was developed for a specific task, the extension to further applications seems promising. Besides other tasks which require intersection solutions, a fuzzy convergence may prove useful (for instance) as an image database feature.

2 Related work

Least squares solutions (see [8]) 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 [7]). 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 [6]) 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, although methods to generalize or fuzzify the Hough space (see [1, 4]) have improved their robustness.

The solution presented in this paper is a voting-type 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

[1] This experiment is further discussed in Section 6. Here it is described only for motivation.

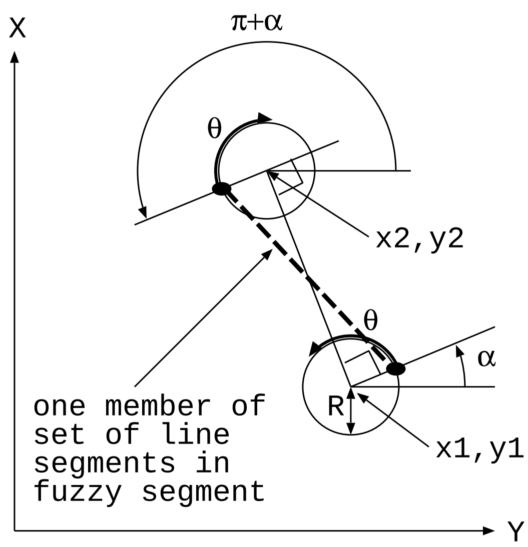

Figure 4: Fuzzy segment model.

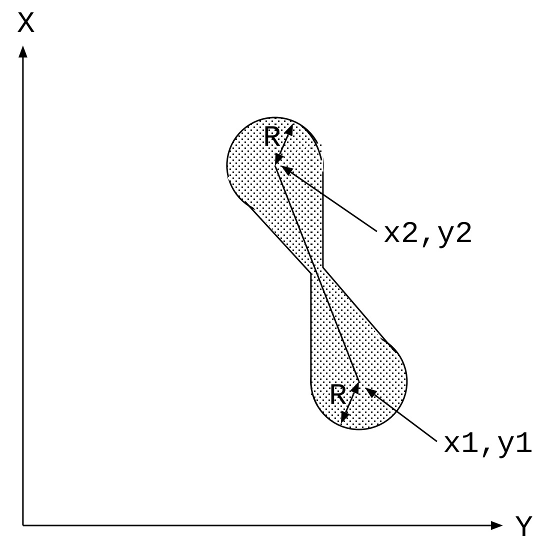

Figure 5: Area covered by fuzzy segment.

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-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.

3 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 $\mathcal{F}$, is defined by a set of parametric line segments:

$$\begin{aligned} x(t) &= x_1 + r\cos(\alpha+\theta) + \\ &\quad (x_2 - x_1 - 2r\cos\theta\cos\alpha)t \\ y(t) &= y_1 + r\sin(\alpha+\theta) + \\ &\quad (y_2 - y_1 - 2r\cos\theta\sin\alpha)t \\ &\quad 0 \le t \le 1, \\ &\quad 0 \le \theta \le 2\pi, \\ &\quad 0 \le r \le R. \end{aligned} \tag{1}$$

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

$$\alpha = \frac{\pi}{2} - \arctan\frac{y_2 - y_1}{x_1 - x_1}. \tag{2}$$

The fuzzy segment $\mathcal{F}$ 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 4 illustrates the enumeration for the subset of $\mathcal{F}$ where $r = R$. It may be visualized as a moving line segment, whose initial endpoints are marked in Figure 4 at $\theta = 0$. As theta moves from zero to 2π, the endpoints trace the boundaries of circles, one in the clockwise direction, the other counter-clockwise. By starting $\theta = 0$ at α, the shape of the fuzzy segment, pictured in Figure 5, remains invariant to orientation.

Relevant previous work in fuzzy geometry (see [2, 3, 9]) deals mainly with fuzzy line models. The fuzzy segment $\mathcal{F}$ 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

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

the width at its endpoints is R. The model also proposes that a line-like shape generally only contributes to a convergence in its 'near' neighborhood. This is further discussed in Section 6. The model allows for an interpretation of what is near via the parameter R.

4 Finding convergences

Given a binary input image, like the one depicted in Figure 3, a process for finding convergences works

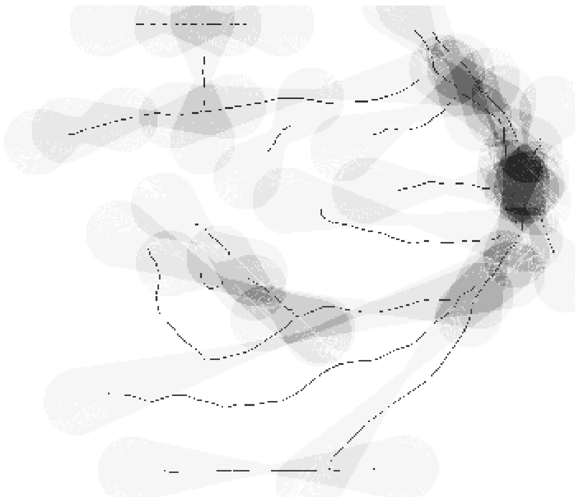

Figure 6: Convergence image.

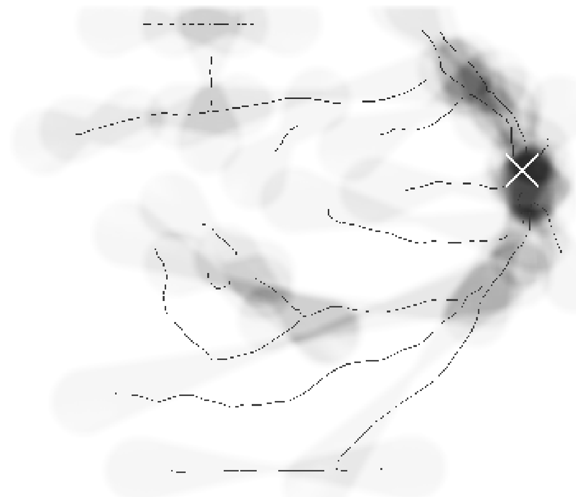

Figure 7: Blurred convergence image, with 'X' over strongest pixel.

as follows:

1. Thin the image (for instance, using the algorithm given in [6], 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. Erase (relabel as background) any segments whose size (in connected pixels) is below an absolute threshold S.

4. Extend each segment a distance of R pixels in both directions. The extension is done along the vector made by the segment's endpoints.

5. Model each segment, via its two extended endpoints, with a fuzzy segment $\mathcal{F}$. The image area covered by $\mathcal{F}$ may be found by enumerating $\mathcal{F}$ (see eq. 2) at $r = R$ with suitable discretizations of θ and t. For the experiments reported in Section 5, 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 $\mathcal{F}$. This image is cleared after each fuzzy segment enumeration and voting is completed.

6. For each pixel enumerated in $\mathcal{F}$, a vote is cast. The image used to tally these votes may be 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 $\mathcal{F}$.

7. Figure 6 shows the convergence image for the input shown in Figure 3, superimposed on the thinned segments. The measures of convergence have been normalized so that the highest vote yields the darkest pixel value. Figure 7 shows the image from Figure 6 after 11×11 mean filter blurring, and renormalization, along with an 'X' marking the point of highest convergence.

5 Experimental results

The algorithm detailed in Section 4 was run on twenty images like the one presented in Figure 3. The input images were 605×700 pixels in size. The algorithm's parameters were set at $R = 40, S = 30$. Figure 8 presents the input binary vessels segmentation, and Figure 9 presents the resulting blurred convergence image, for the fundus image shown in Figure 2.

Out of twenty images, the strongest point of convergence was located inside the optic nerve thirteen times. In two cases, the second (locally) strongest point of convergence was located inside the optic nerve. In five cases, the convergence image did not provide a strong indicator inside the optic nerve. This success rate (13/20 = 65%) compares favorably with the av-

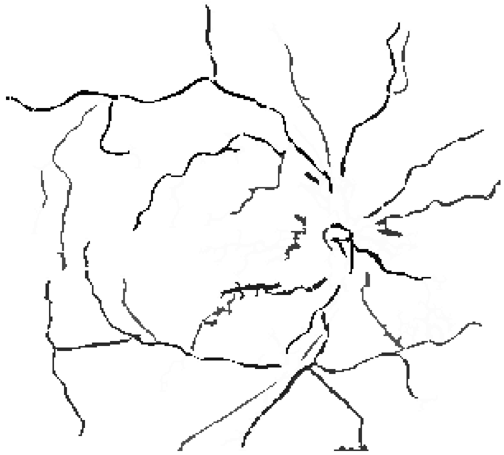

Figure 8: Segmentation of blood vessels, for Fig. 2.

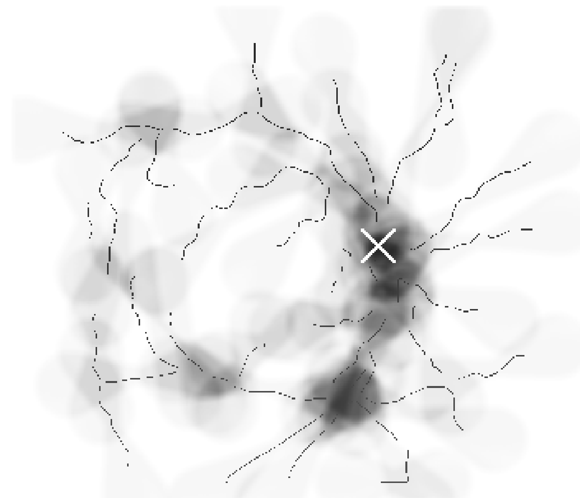

Figure 9: Blurred convergence image, with 'X' over strongest pixel.

erage success rate of the people in the experiment described in Section 1 (69/100 = 69%). There was also a noticeable correlation in performance on individual images. For the thirteen cases in which fuzzy convergence was successful, the five people had an 82% success rate (average). For the seven cases in which fuzzy convergence was not successful, the five people had only a 46% success rate (average).

It was noted in Section 4 that changes in the voting function had little effect on average performance. Changes in values for the parameters R (in the range 20-60) and S (in the range 20-50) were also found to have little effect, on average. Finally, a fuzzy segment model was also developed and tested for a parametric quadratic curve. For some individual segments, the curved fuzzy segment modeled the shape noticeably better than the straight fuzzy segment, but for this particular application, there was again (surprisingly) no improvement on average performance. Further applications and test cases need to be explored.

6 Discussion

The experiment discussed in Section 1 could obviously be improved in control and formalness. However, for motivational purposes, it is illuminating. The vessel maps shown to the subjects contained breaks, missing pieces, spurious pieces, and other imperfections, because they were produced via automatic segmentation. The subjects were not familiar with this work, and were not told what the shapes in the images represented. Nevertheless, the subjects located the optic nerves (as points of convergence) with relative ease. It should be noted that in the majority of cases where points were selected that were not within the optic nerve, an interpretation of the word 'convergence' was often the cause. In many cases, a subject selected a point within a group of smaller, but more numerous and better connected, group of segments. The cluster towards the middle-left of the image in Figure 3 is an example.

Based upon this experiment, a group discussion led to the following conjecture:

> *A line-like shape generally only contributes to the perception of a convergence in its near neighborhood.*

Figure 10(a) shows a jumble of lines, which is duplicated in Figure 10(b) with numeric labels, for reference. When casually examining Figure 10(a), it is evident that lines 15, 16, 17, 19, and possibly 21 contribute to a sense of convergence. Interestingly, lines 3, 4, 6, 9, 17, and 24, although they all extend to a common point of intersection, do not yield nearly as strong a sense of convergence.[2]

One implication of this conjecture is that solutions for convergence which model line-like shapes with lines (which includes the majority of known methods) may often be inappropriate. The visual field, after projection and field-of-view clipping, can not in general record a line (meaning its infinite length). Yet people can easily find convergences. Again, the implication is that there likely exists a method(s) for solution that does not use lines.

[2]Of course, this is only one simple example, subject to interpretation; proper perceptual experimentation is necessary.

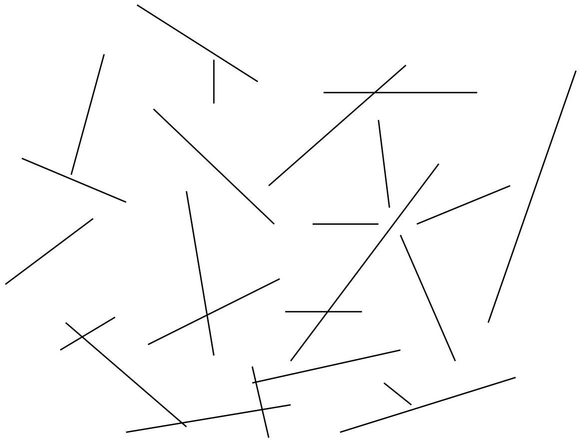

(a) jumble of lines

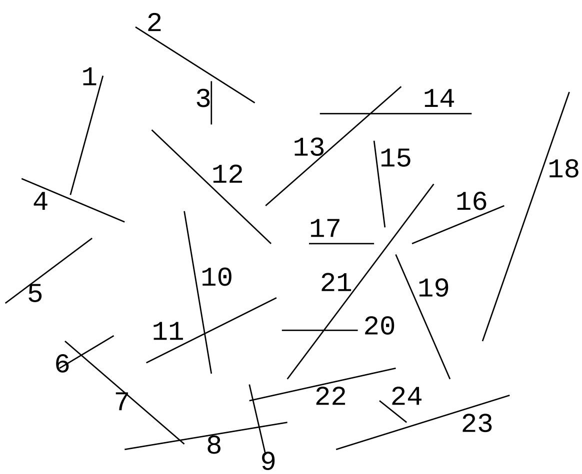

(b) labeled, for reference in text

Figure 10: An informal experiment in convergence perception.

This work proposes a method for finding convergences of line-like shapes in an image. The method models each shape with a fuzzy segment, which is defined as a set of line segments. Each fuzzy segment contributes votes for convergence in an area of the image space. The summation of votes may be visualized as a *convergence image*. This image may then be blurred and thresholded for a single strongest point of convergence, or thresholded relative to some absolute value to locate multiple (or zero) areas of convergence.

Besides further perceptual experimentation, further applications for this technique should be explored. Fuzzy convergences may prove useful as image features, for instance for image matching or registration. To improve optic nerve detection, additional perceptual cues (size, color, shape) should be incorporated. In some cases, a curved fuzzy segment model may prove appropriate. Insights might also be gained by performing application-based comparisons between fuzzy convergence, least-squares, and Hough space approaches.

7 Acknowledgments

The authors thank Simone Santini for useful discussions during the development of this work. This work was supported by the NIH Library of Medicine, grant #LM 05759-09.

References

[1] S. M. Bhandarkar, "A Fuzzy Probabilistic Model for the Generalized Hough Transform", in *IEEE Trans. on Systems, Man & Cybernetics*, **24:5**, May '94, pp. 745-759.

[2] B. B. Chaudhuri, "Some shape definitions in fuzzy geometry of space", in *Pattern Recognition Letters*, **12**, Sep '91, pp. 531-535.

[3] K. C. Gupta and S. Ray, "Fuzzy plane projective geometry", in *Fuzzy Sets & Systems*, **54**, 1993, pp. 191-206.

[4] J. H. Han, L. T. Koczy and T. Poston, "Fuzzy Hough Transform", in *Pattern Recognition Letters*, **15**, July '94, pp. 649-658.

[5] A. Hoover, V. Kouznetsova and M. Goldbaum, "Locating Blood Vessels in Retinal Images by Piece-wise Threshold Probing of a Matched Filter Response", submitted to *American Medical Informatics Association Annual Symposium*, 1998.

[6] R. Jain, R. Kasturi and B. G. Schunck, Machine Vision, McGraw-Hill, 1995.

[7] P. Meer, D. Mintz and A. Rosenfeld, "Robust Regression Methods for Comptuer Vision: A Review", in *Int'l J. of Computer Vision*, **6:1**, 1991, pp. 59-70.

[8] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, Numerical Recipes in C: The Art of Scientific Computing, 2nd edition, Cambridge Press, 1992.

[9] A. Rosenfeld, ""Geometric properties" of sets of lines", in *Pattern Recognition Letters*, **16**, May '95, pp. 549-556.