Locating Blood Vessels in Retinal Images by Piecewise Threshold Probing of a Matched Filter Response

Adam Hoover*, Valentina Kouznetsova, and Michael Goldbaum

Abstract—We describe an automated method to locate and outline blood vessels in images of the ocular fundus. Such a tool should prove useful to eye care specialists for purposes of patient screening, treatment evaluation, and clinical study. Our method differs from previously known methods in that it uses local and global vessel features cooperatively to segment the vessel network. We evaluate our method using hand-labeled ground truth segmentations of 20 images. A plot of the operating characteristic shows that our method reduces false positives by as much as 15 times over basic thresholding of a matched filter response (MFR), at up to a 75% true positive rate. For a baseline, we also compared the ground truth against a second hand-labeling, yielding a 90% true positive and a 4% false positive detection rate, on average. These numbers suggest there is still room for a 15% true positive rate improvement, with the same false positive rate, over our method. We are making all our images and hand labelings publicly available for interested researchers to use in evaluating related methods.

Index Terms—Adaptive thresholding, blood vessel segmentation, matched filter, retinal imaging.

I. INTRODUCTION

BLOOD vessel appearance is an important indicator for many diagnoses, including diabetes, hypertension, and arteriosclerosis. Veins and arteries have many observable features, including diameter, color, tortuosity (relative curvature), and opacity (reflectivity). Artery–vein crossings and patterns of small vessels can also serve as diagnostic indicators. An accurate delineation of the boundaries of blood vessels makes precise measurements of these features possible. These measurements may then be applied to a variety of tasks, including diagnosis, treatment evaluation, and clinical study. We describe an automated method to locate and outline blood vessels in images of the ocular fundus. With this tool, eye care specialists can potentially screen larger populations for vessel abnormalities. Precise measurements may be more easily recorded, for instance, for evaluation of treatment or for clinical study (such as reported in [2]). Observations based upon such a tool would also be more systematically reproducible.

Previous methods to segment blood vessels automatically have concentrated primarily on their local attributes. Vessels may be characterized by the expected color (reddish), shape (curvilinear), gradient (strength of boundary), and contrast (with background). Unfortunately, this description is not exclusive. For suitable ranges of these attributes, other image manifestations, such as the boundaries of the optic nerve and some hemorrhages and lesions, can exhibit the same local attributes as vessels.

Fig. 1 shows an example retinal image, along with an image showing the result of the matched filter convolution described in [3]. The strength of the matched filter response (MFR) is coded in greyscale: the darker a pixel, the stronger the response. Notice that the strong responses in the center of the MFR image, which are obviously not vessel, are unfortunately much stronger than the responses on the left side of the MFR image, which are vessel. Therefore, applying a single global threshold does not provide adequate classification, as shown in Fig. 2. A bilevel threshold (such as hysteresis) is also inadequate, because the vessel and nonvessel pixels with strong MFR's are usually spatially connected, as in Fig. 1.

We propose a novel method to segment blood vessels that compliments local vessel attributes with region-based attributes of the network structure. A piece of the blood vessel network is hypothesized by probing an area of the MFR image, iteratively decreasing the threshold. At each iteration, region-based attributes of the piece are tested to consider probe continuation, and ultimately to decide if the piece is vessel. Pixels from probes that are not classified as vessel are recycled for further probing. The strength of this approach is that individual pixel labels are decided using local and region-based properties.

*A. Hoover is with the Electrical and Computer Engineering Department, Clemson University, Clemson, SC 29634-0915 USA (e-mail: ahoover@clemson.edu).

V. Kouznetsova is with the Visual Computing Lab., Electrical and Computer Engineering Department, University of California, San Diego, La Jolla, CA 92093-0407 USA (e-mail: vkouznet@vision.ucsd.edu).

M. Goldbaum is with the Department of Ophthalmology, University of California, San Diego, La Jolla, CA 92093 USA (e-mail: mgoldbaum@ucsd.edu).

II. RELATED WORK

Previous methods to segment blood vessels generally fall into three categories: window-based [3], [16], [17], [19], classifier-based [5], [21], and tracking-based [20]–[23]. Window-based methods, such as edge detection, estimate a match at each pixel for a given model against the pixel's surrounding window. In [3], the cross section of a vessel in a retinal image was modeled by a Gaussian shaped curve and then detected using rotated matched filters. In [17], a similar method was used for artery detection in angiograms. In [19], a standard gradient filter was

(a)   (b)

Fig. 1.   (a) An example retinal image with obscured vessels. (b) MFR. The response is coded such that a darker value represents a stronger response.

(a)   (b)
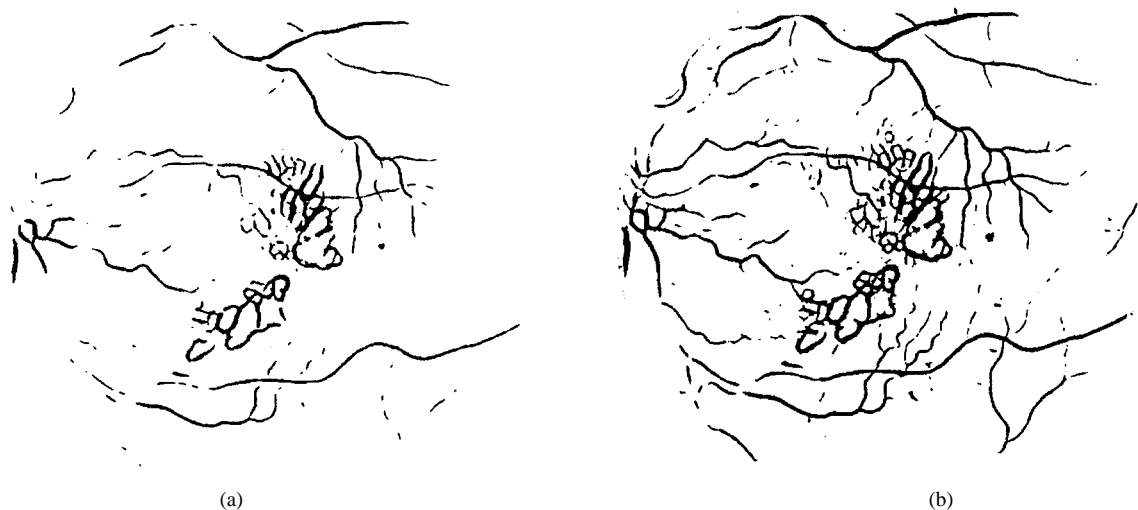
Fig. 2.   (a) and (b) MFR thresholded at two different values. There is a strong overlap between true positive and false positive responses.

used to detect pixels on the boundary of retinal vessels for subsequent grouping. In [16], a window surrounding a vessel pixel was modeled by a neural network trained on user selected examples. The drawback of these methods is that the large-scale properties of vessels (i.e., their network structure) must be ignored to insure computational feasibility.

Classifier-based methods proceed in two steps. First, a low-level algorithm produces a segmentation of spatially connected regions. These candidate regions are then classified as being vessel or not vessel. In [21], regions segmented by user-assisted thresholding were classified as blood vessel or leakage according to their length to width ratio. In [5], regions segmented by the method in [3] were classified as vessel or not vessel according to many properties, including their response to a classic operator designed to detect roads in aerial imagery [8]. The drawback of these methods is that the large-scale properties of vessels cannot be applied to the problem until after the low-level segmentation has already finished. Therefore, these properties cannot be used to drive the segmentation, merely to evaluate it.

Tracking-based methods utilize a profile model to incrementally step along and segment a vessel. In [22], a Hough transform is used to locate the papilla in a retinal image. Vessel tracing proceeds iteratively from the papilla, halting when the response to a one-dimensional (1-D) (cross-section) matched filter falls below a given threshold. In [20], a similar method was employed to detect vessels in coronary arteriograms, from user-given starting points. In [23], the tracking method was driven by a fuzzy model of a 1-D vessel profile. One drawback to these approaches is their proclivity for termination at branch points (whether real or caused by pathology), which are not detected well by 1-D filters. Another drawback is their reliance upon unsophisticated methods for locating starting points, which must always be either at the optic nerve or at subsequently detected branch points.

In [6], a method for tracking edge paths is used to segment arteries in cineangiograms. Edge paths are modeled as Markov chains. A sequential edge linking (SEL) algorithm is introduced to search the possible set of paths for the best fit to the Markov model. The probabilities of the model are adjusted to reflect the properties of the desired path, such as the tolerance to local curvature. A strength of this approach is that the grouping operation works upon actual gradient values, as opposed to a

thresholded response. Therefore, a segmentation decision is not reached until an arbitrary number of pixels is available for classification. A drawback to the approach is that branches are not modeled, so that each branch must be traced and classified independently.

In this work, we propose a new method for segmenting blood vessels in a retinal image. The MFR image, computed as described in [3], is thresholded using a novel probing technique. The probe examines the image in pieces, testing a number of region-based properties. If the probe decides a piece is vessel, then the constituent pixels are simultaneously segmented and classified. Contrasted against classifier-based methods, our probing method allows a pixel to be tested in multiple region configurations before final classification. Contrasted against tracking-based methods, our probing method is driven by a two-dimensional (2-D) MFR. Contrasted against [6], our probing method is region based and so naturally allows for multiple branches.

This paper expands upon a preliminary report given at the 1998 American Medical Informatics Association Annual Symposium [14].

III. Algorithm

We first review the matched filter construction and convolution, described in [3], upon which our algorithm builds. We then present threshold probing and its application to blood vessel segmentation in a retinal image.

A. Matched Filter for Blood Vessels

A matched filter describes the expected appearance of a desired signal, for purposes of comparative matching. In [3] a Gaussian function is proposed as a model for a blood vessel profile. The model is extended to two dimensions by assuming a vessel has a fixed width and direction for a short length. Since vessels may appear in any orientation, a set of 2-D segment profiles in equiangular rotations is used as a filter bank. The filters are implemented using twelve 16×16 pixel kernels. The details for computing the actual values in the kernels may be found in [3].

The matched filter is applied by convolving a retinal image with all twelve kernels.[1] The MFR is taken as the value for the highest scoring kernel at each pixel. On a Sun SPARCstation 20, the computation of the MFR image for a 700×605-pixel retinal image takes approximately 5 min. For purposes of threshold probing, the MFR image is normalized and quantized to eight bits per pixel.

B. Threshold Probing

The basic operation of our algorithm is to probe regions in an MFR image. During each probe, a set of criteria is tested to determine the threshold of the probe, and ultimately to decide if the area being probed (termed a piece) is blood vessel. A flowchart for the algorithm is shown in Fig. 3. A queue of points is initialized, each of which will be used for a probe. Upon a probe's completion, if the piece is determined to be vessel, then

---

[1]A retinal image is normally captured in full color. However, using an RGB color model, the blue band in a retinal image is often empty and the red band is often saturated. The matched filter convolution therefore uses only the green band.
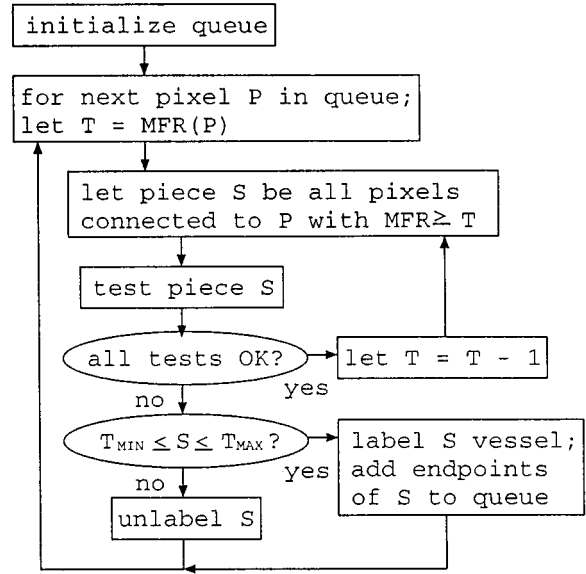

Fig. 3. Flowchart of algorithm.

the endpoints of the piece are added to the queue. In this way, different probes (and thus different thresholds) can be applied throughout the image.

The following steps initialize a queue of pixels which are to be used as starting points for probing.

- Convolve the matched filter described in [3] with the image, producing an MFR image.
- Using a histogram of the MFR image, threshold the image such that $>T_{THRESH}$ pixels are above the threshold.
- Thin the thresholded image (for instance, using the algorithm given in [15, p. 59]).
- In the thinned image, erase (relabel as background) all branchpoints, breaking up the entire foreground into segments that contain two endpoints each. 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.
- Discard segments with less than ten pixels.
- All remaining endpoints are placed in the probe queue.

The segments created by simple thresholding (above) are used only to locate a set of starting points to initialize the probe queue. The segments themselves will not appear in the final segmentation unless the probing procedure (below) causes their reappearance and classification as vessel. This process of initialization allows the pixels with a strong response to the matched filter to act as candidate vessels, with the design that not all need necessarily become part of the final vessel segmentation. Unlike tracking-based methods [20], [22], [23], these starting points can be anywhere in the vessel network, so that pathology and branches do not cause parts of the network to be missed.

Each pixel in the probe queue is used as a starting point for threshold probing. The probing is iterative. The iterations are used to determine an appropriate threshold for the area being probed. The initial threshold is the MFR image value at the starting pixel. In each iteration, a region is grown from the start pixel, using a
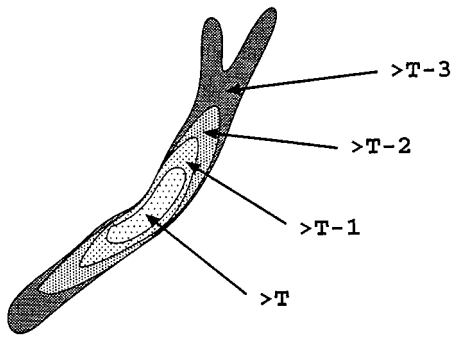
Fig. 4. The basis of threshold probing. At each iteration, a set of region tests is applied to determine if the threshold may be decreased an additional step. Local probing halts when any of the region tests fail.

conditional paint-fill technique. The paint fill spreads across all connecting pixels that are not already labeled and that are above the current threshold. Once the paint-fill is complete, the desired attributes of the grown region are tested. If the region passes the tests, then the threshold is decreased by one and a new iteration begins. The technique is illustrated in Fig. 4.

Each probe iteration conducts the following tests.

- If the piece size (in pixels) exceeds T_{MAX}, then the probe halts. This requires multiple pieces (and thus potentially multiple thresholds) to segment the entire image. The effect is that the probe adapts to the local strength of the MFR image.
- If the threshold reaches zero, then the probe halts. This happens when probing a small area (even one pixel) interior to an area already classified as vessel.
- If the piece touches (on its border) more than one previously vessel-classified pieces, then the probe halts. This is particularly useful for bridging gaps along vessels exhibiting weak MFR values.
- If the ratio $(border\text{-}pixels\text{-}touching\text{-}another\text{-}piece/total\text{-}pixels\text{-}in\text{-}piece) > T_{FRINGE}$ then the piece is fringing, and the probe halts. This prevents a probe from searching along the borders of vessel pieces already segmented.
- If the ratio $(total\text{-}pixels\text{-}in\text{-}piece/branches\text{-}in\text{-}piece) < T_{TREE}$, then the probe halts. This requires a piece to have a minimum span of vessel(s) per branch, and thus prevents over-branching down false paths. The count of *branches-in-piece* is found by calculating the skeleton of the piece at each iteration. The computational cost of this step is kept low by using indirect image indexing (a list of the image coordinates of the pixels in the piece).

None of the tests relying upon thresholds (T_{MAX}, T_{FRINGE}, T_{TREE}) is performed until the piece reaches at least 30 pixels in size. However, the other tests may cause the probe to halt before this minimum is reached.

Once the probe is complete, if the resulting region has at least T_{MIN} pixels but less than T_{MAX} pixels, or connects two previously probed pieces, then the region is labeled as vessel. The endpoints of the vessel piece are added to the queue. If the region is not determined to be vessel, then its pixels are left unlabeled. In either case, the next point in the queue is selected for probing. When the queue is empty, the algorithm is complete.

Probes that begin at the endpoints of previously grown pieces have one additional constraint. An eight-pixel long artificial boundary is placed perpendicular to the end of the previously grown piece, to prevent the new piece from probing back along the sides of the piece already grown. This forces the growing piece to probe in a new direction. These artificial boundaries are removed at the completion of the algorithm.

Some tests besides those listed above were explored during the development of this algorithm. Of particular note is a test for spatial looping, which is often caused by lesions and hemorrhages. Halting the probe when a loop is detected eliminates many of the false positive responses to these pathologies. However, it also halts probing wherever vessels cross each other at different depths of the retina. Without an explicit depth perception, these crossings appear exactly like loops. In the final analysis we abandoned the loop test because of this problem.

IV. Images

Twenty retinal fundus slides were selected for testing the described method. The slides were captured by a TopCon TRV-50 fundus camera at 35° field of view. Each slide was digitized to produce a 605×700 pixel image, 24 bits per pixel (standard RGB). Ten of the images are of patients with no pathology (normals). Ten of the images contain pathology that obscures or confuses the blood vessel appearance in varying portions of the image (abnormals). This selection was made for three reasons. First, most of the referenced methods have only been demonstrated upon normal vessel appearances, which are easier to discern. Second, some level of success with nonnormal vessel appearances must be established to recommend clinical usage. Third, we desired to evaluate the performance difference (if any) of our algorithm on normal and abnormal cases.

Each of these 20 images was carefully labeled by hand, to produce a ground truth vessels segmentation. An example is shown in Fig. 5. The tool used for hand labeling is adapted from the tool described in [12], which was used to create hand-labeled images for evaluating range image segmentation algorithms [13]. The tool allows the user to magnify the image to a level appropriate for labeling individual pixels, one at a time, as being vessel or not vessel. The tool also allows the user to apply various histogram transformations, to better visualize the original image data. The process of labeling an image takes several hours, depending on the user and image.[2]

Fig. 6 shows the distribution of MFR values for pixels hand labeled as vessel. Fig. 6(a) shows the distribution for the ten normal cases, Fig. 6(b) shows the distribution for the ten abnormal cases. Although there is a better separation between vessel and nonvessel pixels in the normal cases, there is a significant overlap in both normal and abnormal cases. The results from basic thresholding on an abnormal image, presented in Fig. 2, are explained by this overlap.

The classification of a majority of the pixels is often clear to a human observer. However, some of the pixels, such as those on the boundary of a vessel, those for small vessels, and those

---

[2]Lengthy breaks are often required to maintain a reasonable perspective, so the process takes longer than one might imagine.
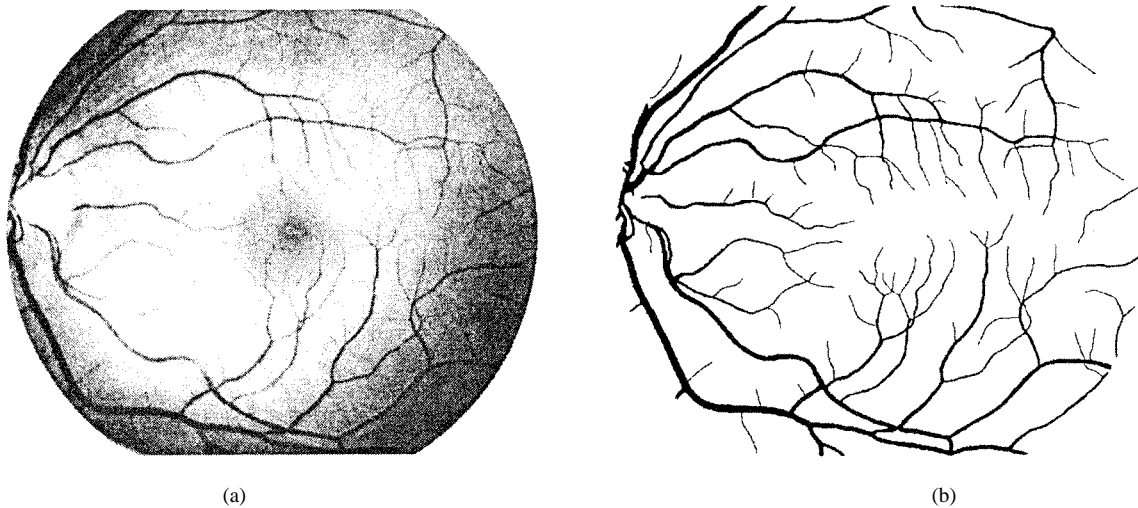
Fig. 5.   (a) An example retinal image showing normal vessels. (b) A hand-labeled ground truth vessels segmentation.

for vessels near pathology, are less easily labeled. To estimate this variance in observation, a second person[3] produced an additional set of hand labelings for the 20 test images. For the results reported in Section V, this second labeling is used to establish a reference for performance comparison.

On average, the first person labeled 32 200 pixels in each image as vessel, while the second person labeled 46 100 pixels in each image as vessel. Subsequent review indicated that the first person took a more conservative view of the boundaries of vessels and in the identification of small vessels than the second person. Both labelings for one of the normal images are shown in Fig. 7. We are making all the original retinal images and hand labelings available to any interested researchers for development and evaluation of related methods.[4]

V. EXPERIMENTS

The MFR images for all 20 images were processed using basic thresholding. Performance was established as follows. Any pixel which was hand labeled as vessel, whose MFR is also above the given threshold was counted as a true positive. Any pixel which was hand labeled as not vessel whose MFR is also above the given threshold was counted as a false positive. The true positive rate is established by dividing the number of true positives by the total count of pixels hand labeled as vessel. The false positive rate is established by dividing the number of false positives by the total count of pixels hand labeled as not vessel. Fig. 8 shows the true positive and false positive detection rates across the range of possible thresholds.

Note that the false positive detection rate is considerably worse for the abnormal cases than for the normal cases. By comparing the second hand-labeled images to the first hand-labeled images (using the same method as outlined for thresholding, above) we can establish a target performance level. This level is indicated by three isolated marks in Fig. 8, showing the second person's performance on the normal, abnormal, and average case. Note that the grouping of these

[3]For background, both of the ground truthers are experts in image processing and had spent at least two years studying retinal imagery before this work.

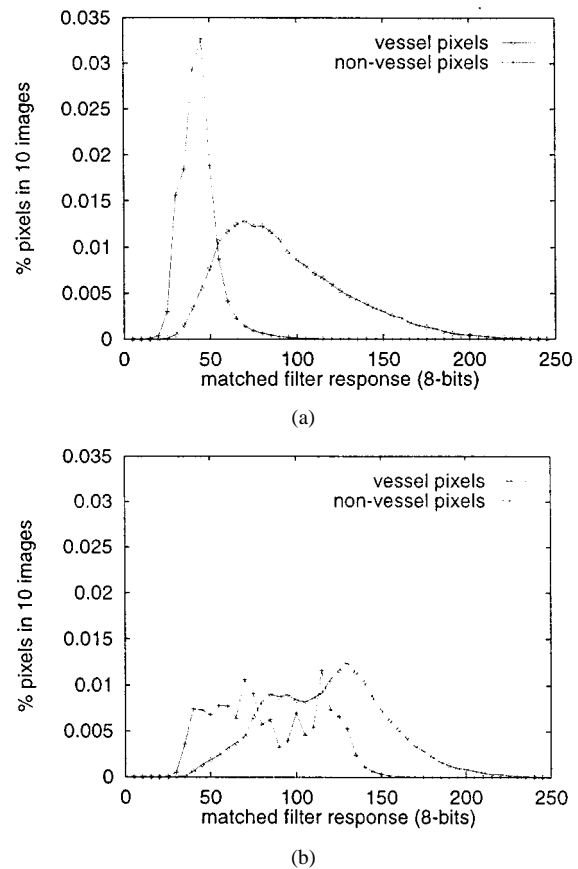[4]Available at www.ces.clemson.edu/~ahoover/stare.

Fig. 6.   Distribution of matched filter response (MFR) for pixels labeled vessel and nonvessel. (a) The distribution for ten normal images. (b) The distribution for ten abnormal images.

three marks shows a small distribution, indicating that people may in fact be somewhat affected (for this task) by the presence of pathology.

The shapes of the curves in Fig. 8 are explained by reexamining the distributions of pixels shown in Fig. 6. The abnormals not only have a greater overlap, but also a bimodal distribution of nonvessel pixels. The bimodal distribution is caused by strong responses of the MFR to the boundaries of lesions, hemorrhages, and other pathology. This causes the dent in the abnor-

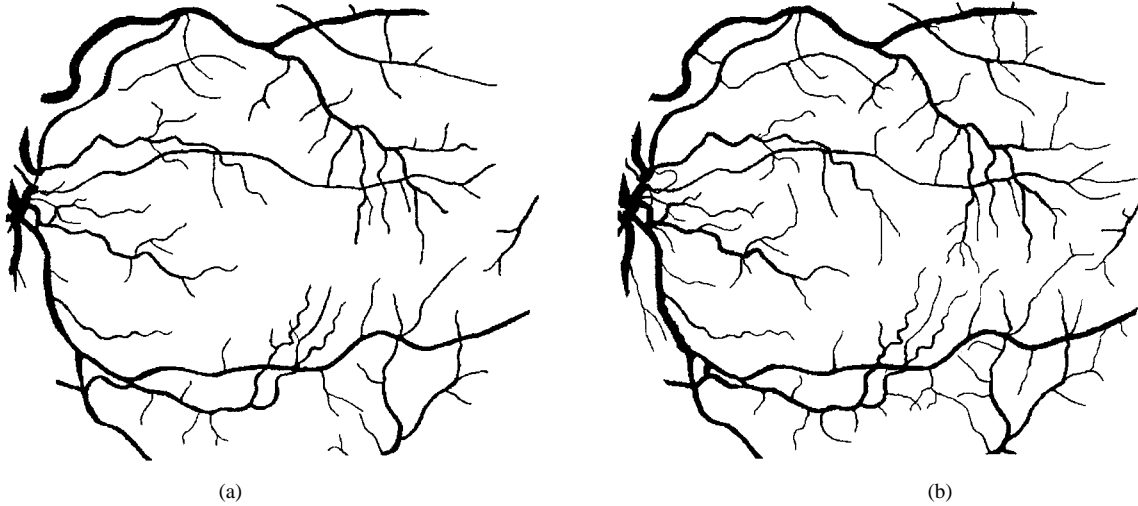(a)   (b)

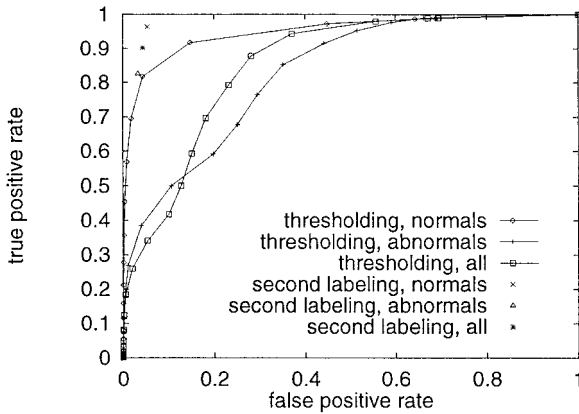Fig. 7.   An example hand-labeled vessels segmentation from each of two persons.

Fig. 8.   Average performance for basic thresholding across the range of possible thresholds. For reference, the comparison of the second hand-labeled images to the first hand-labeled images is also shown.

mals' curve. Note also that the actual number of nonvessel pixels outnumbers the number of vessel pixels by a factor of ten. The appearance of a substantial number of false negatives occurs at a much higher threshold for the abnormals than for the normals (see Fig. 6), so that the average performance curve actually dips below the abnormals performance curve for a short range.

There are five parameters for our algorithm: T_{THRESH}; T_{MIN}; T_{MAX}; T_{FRINGE}; and T_{TREE}. We report results processing all 20 of our images using ten sets of values for these parameters

$$
\left.\begin{array}{r}
T_{THRESH} = 5000 + 5000i \\
T_{MIN} = 180 - 20i \\
T_{MAX} = 1000 + 1000i \\
T_{FRINGE} = 0.1 + 0.04i \\
T_{TREE} = 300 - 25i
\end{array}\right\} \forall i = 0 \cdots 9 \qquad (1)
$$

where each value of i represents one tested set of values. An example result, processed at values in the middle of these sets, is

shown in Fig. 9. Several functions similar to (1) were explored, by varying the initial values and increments. This strategy was taken in lieu of a full five-parameter search for the best performance curve, which is computationally prohibitive. All 20 images were used to select the best parameter curve. However, the additional parameter curves explored produced very similar results. Based on this observation we believe that the overestimation of performance caused by the absence of separate train and test sets is minimal in this case.

The performance curves for our algorithm on the normals, abnormals, and all images are shown in Fig. 10. For reference, the average performance mark for the second set of hand-labeled images is included, as is the average performance curve for basic thresholding. Note that there is virtually no difference in the performance of our algorithm on normals or abnormals. Also note that the performance of our algorithm reduces the number of false positives by as much as 15 times over basic thresholding of an MFR, at up to a 75% true positive rate. For these experiments, our algorithm appears to have a breaking point at an approximately 80% true positive rate. Our algorithm produces the same number of false positives at a 75% true positive rate as the second set of hand-labeled images produces at a 90% true positive rate. This suggests room for an improvement of 15% in the true positive rate over our method.

Fig. 11(a) shows another example result from our method. In Fig. 11(b) the pixels are greycoded to represent the final local threshold used during probing (the darker the pixel, the higher the threshold). Fig. 12 shows the distribution of thresholds relative to the number of pixels labeled as vessel using that threshold. These figures show how our method adapts to the local maxima in the MFR image. In particular, the large vessels close to the optic nerve give a strong response to the matched spatial filter, so that the threshold applied locally is high [shown as darker segments in Fig. 11(b)]. The smaller vessels further down the vessel network give a continually weaker response to the matched spatial filter, so that the threshold applied locally continues to decrease [shown as brighter segments in Fig. 11(b)].
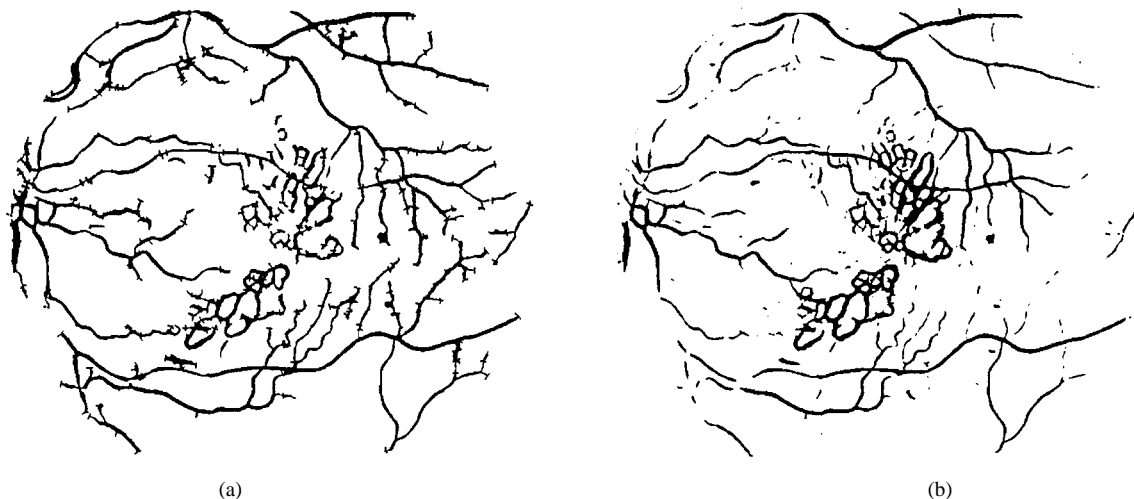
Fig. 9.   (a) An example result (abnormal case) from threshold probing. (b) A result on the same image for basic thresholding is shown for comparison.
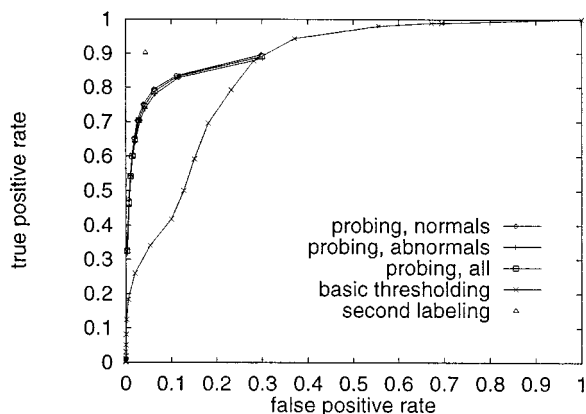

Fig. 10.   Average performance for our method across the given sets of parameter values (see text). For reference, the comparison of the second hand-labeled images to the first hand-labeled images is also shown, as is the average performance curve for basic thresholding.

The most convincing demonstration of the proposed algorithm is obtained by observing its iterative dynamic operation. For this purpose, a Windows O/S executable has been made available[5] that shows the real-time execution of our algorithm on an image. This executable may be observed like a movie clip. Interested readers are also directed to that site for the original data, ground truths, and our results for all 20 images.

VI. Conclusions and Discussion

The described method segments roughly 3/4 of the vessels in a retinal fundus image at a false positive rate comparable to a human observer. Compared to basic thresholding of an MFR, as reported in [3], we have shown that our method reduces the false positive rate by a factor of 15 times. However, empirical evaluation also indicates the potential for a further 15% improvement in the true positive detection rate.

[5]At www.ces.clemson.edu/~ahoover/stare.

Threshold probing is related to the watershed algorithm (see for instance [9]). In fact, threshold probing is a generalization of the watershed algorithm. In the scheme of threshold probing, the watershed algorithm may be characterized by tests which halt the local probe at either: 1) any local maxima or 2) any contact with the boundary of another probe. Threshold probing may be considered a generalization of this approach in which more complex tests, for instance upon the structure of the local probe, are used to determine halting conditions.

The core of our algorithm, threshold probing, is independent of the actual region tests applied. We suppose that this algorithm may be applied to other tasks in which thresholding of a 2-D image is required, for instance thresholding gradients for edge detection. The design of suitable tests for this task is a topic of current interest.

The empirical evaluation of image processing algorithms is currently receiving a wellspring of attention (see, for instance, [1] and [4]), as is its difficulty. One aspect of our approach not captured in our evaluation is the property of connectedness. By design, our approach produces a labeling (right or wrong) of continuous segments. Basic thresholding is likely to produce small groups of isolated pixels, as in Fig. 2. Although such pixels may in fact be correctly labeled, their utility for measurement is probably limited. For future work, we plan on exploring the evaluation of this property.

The application of our method to three-dimensional (3-D) vessel segmentation on data from other imaging modalities (see for instance [7], [10], or [18]) is also a topic of current interest. Finally, we are examining applications of our method for making actual measurements and classifications of retinal blood vessels. Some recent results from measuring tortuosity [11] using automated methods are encouraging.
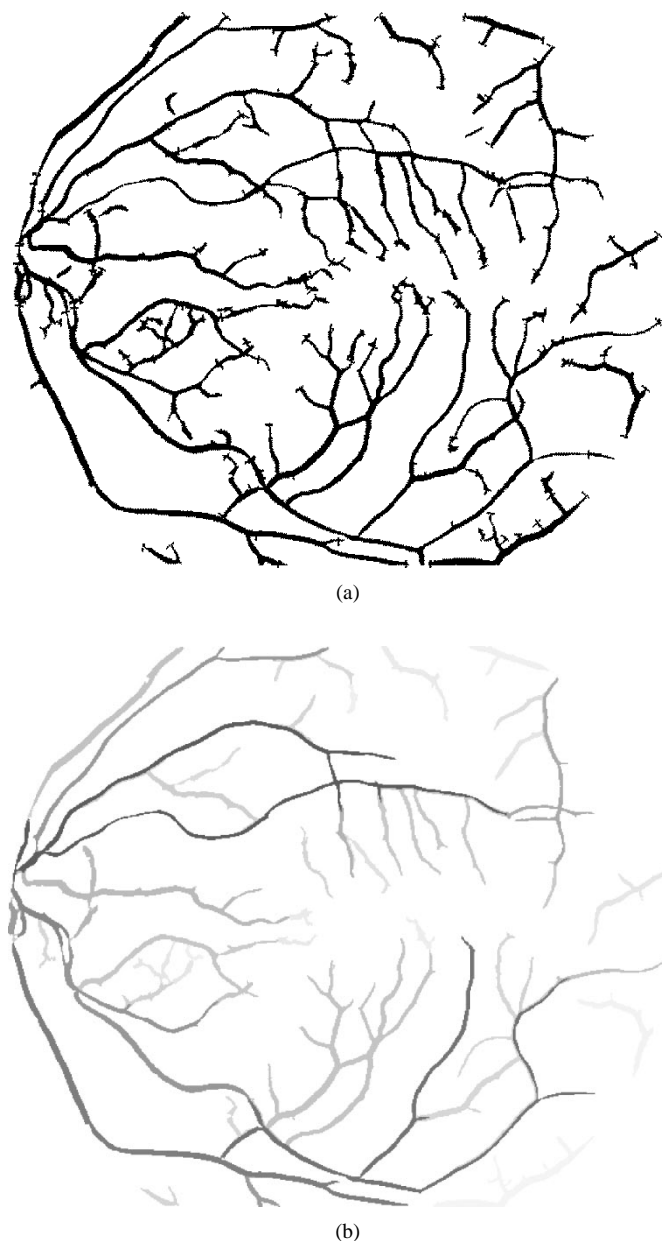
(a)

(b)

Fig. 11.   (a) An example result (normal case) from threshold probing. (b) The final local thresholds used in probing are shown coded in greyscale (the darker the pixel, the higher the threshold).
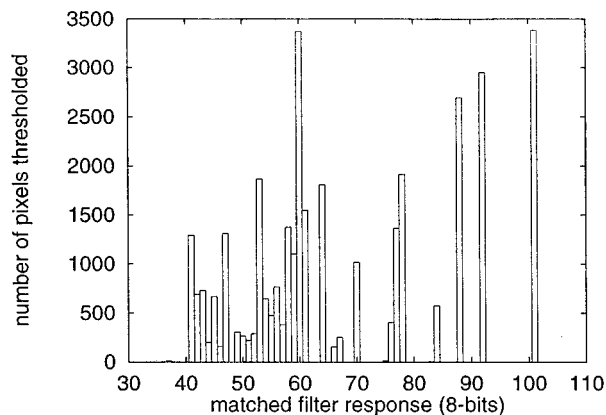

Fig. 12.   Distribution of final thresholds used in local probing for result shown in Fig. 11.

REFERENCES

[1] K. Bowyer and J. Phillips, *Empirical Evaluation Techniques in Computer Vision*: IEEE Computer Society, 1998.

[2] J. Capowski, J. Kylstra, and S. Freedman, "A numeric index based on spatial-frequency for the tortuosity of retinal-vessels and its application to plus disease in retinopathy of prematurity," *Retina*, vol. 15, no. 6, pp. 490–500, 1995.

[3] 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. Med. Imag.*, vol. 8, pp. 263–269, Sept. 1989.

[4] A. Clark and P. Courtney, "Performance characterization and benchmarking of vision systems," workshop organizers of, 1998.

[5] B. Cote, W. Hart, M. Goldbaum, P. Kube, and M. Nelson, "Classification of blood vessels in ocular fundus images," Computer Science and Engineering Dept., Univ. of California, San Diego, Tech Rep., 1994.

[6] P. Eichel, E. Delp, K. Koral, and A. Buda, "A method for a fully automatic definition of coronary arterial edges from cineangiograms," *IEEE Trans. Med. Imag.*, vol. 7, pp. 313–320, Dec. 1988.

[7] J. Fessler and A. Macovski, "Object-based 3-D reconstruction of arterial trees from magnetic resonance angiograms," *IEEE Trans. Med. Imag.*, vol. 10, pp. 25–39, 1991.

[8] M. Fischer, J. Tenenbaum, and H. Wolf, "Detection of roads and linear structures in low resolution aerial imagery using a multisource knowledge integration technique," *Comput. Graph. Image Processing*, vol. 15, no. 3, pp. 201–223, 1981.

[9] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*.   Reading, MA: Addison-Wesley , 1992.

[10] P. Hall, M. Ngan, and P. Andreae, "Reconstruction of vascular networks using three-dimensional models," *IEEE Trans. Med. Imag.*, vol. 16, pp. 919–929, Dec. 1997.

[11] W. Hart, M. Goldbaum, B. Cote, P. Kube, and M. Nelson, "Automated measurement of retinal vascular tortuosity," in *AMIA Annu. Symp.*, Nashville, TN, 1997.

[12] A. Hoover, G. Jean-Baptiste, D. Goldgof, and K. Bowyer, "A methodology for evaluating range image segmentation techniques," in *Second IEEE Workshop Applications Computer Vision*, Sarasota, FL, 1994, pp. 264–271.

[13] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. Goldgof, K. Bowyer, D. Eggert, A. Fitzgibbon, and R. Fisher, "An experimental comparison of range image segmentation algorithms," *IEEE Trans. Pattern Anal. Machine Intell.*, pp. 673–689, July 1996.

[14] A. Hoover, V. Kouznetsova, and M. Goldbaum, "Locating blood vessels in retinal images by piece-wise threshold probing of a matched filter response," in *Proc. AMIA Annu. Symp.*, Orlando, FL, Nov. 1998.

[15] R. Jain, R. Kasturi, and B. G. Schunck, *Machine Vision*.   New York: McGraw-Hill, 1995.

[16] R. Nekovei and Y. Sun, "Back-propagation network and its configuration for blood vessel detection in angiograms," *IEEE Trans. Neural Networks*, vol. 6, pp. 64–72, Jan. 1995.

[17] T. Pappas and J. Lim, "A new method for estimation of coronary artery dimensions in angiograms," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 36, pp. 1501–1513, Sept. 1988.

[18] C. Pellot, A. Herment, M. Sigelle, P. Horain, H. Maitre, and P. Peronneau, "A 3-D reconstruction of vascular structures from two X-ray angiograms using an adapted simulated annealing algorithm," *IEEE Trans. Med. Imag.*, vol. 13, pp. 48–60, Mar. 1994.

[19] A. Pinz, S. Bernogger, P. Datlinger, and A. Kruger, "Mapping the human retina," *IEEE Trans. Med. Imag.*, vol. 17, pp. 606–619, Aug. 1998.

[20] Y. Sun, "Automated identification of vessel contours in coronary arteriograms by an adaptive tracking algorithm," *IEEE Trans. Med. Imag.*, vol. 8, pp. 78–88, Mar. 1989.

[21] S. Tamura, K. Tanaka, S. Ohmori, K. Okazaki, A. Okada, and M. Hoshi, "Semiautomatic leakage analyzing system for time series fluorescein ocular fundus angiography," *Pattern Recognit.*, vol. 16, no. 2, pp. 149–162, 1983.

[22] S. Tamura, Y. Okamoto, and K. Yanashima, "Zero-crossing interval correction in tracing eye-fundus blood vessels," *Pattern Recognit.*, vol. 21, no. 3, pp. 227–233, 1988.

[23] Y. Tolias and S. Panas, "A fuzzy vessel tracking algorithm for retinal images based on fuzzy clustering," *IEEE Trans. Med. Imag.*, vol. 17, pp. 263–273, Apr. 1998.