Tracking of Blood Vessels in Retinal Images Using Kalman Filter

Tamir Yedidya and Richard Hartley
The Australian National University and National ICT Australia *
{tamir.yedidya, richard.hartley}@rsise.anu.edu.au

Abstract

We present an automatic method to segment the blood vessels in retinal images. Our method is based on tracking the center of the vessels using the Kalman filter. We define a linear model to track the blood vessels, suitable for both the detection of wide and thin vessels in noisy images. The estimation of the next state is computed by using gradient information, histogram of the orientations and the expected structure of a vessel. Seed points are detected by a set of matched filters in different widths and orientations. Tracking is carried out for all detected seed points, however we retrace the segmentation for seeds with small confidence. Our algorithm also handles branching points by proceeding in the previous moving direction when no dominant gradient information is available. The method is tested on the public DRIVE database [10] and shows good results with a low false positive rate.

1 Introduction

Detection of blood vessels in retinal images offers a non-invasive method to diagnose various diseases in the retina. Segmenting the blood vessels can be used as the first step for subsequent image processing as they cover big areas of the retinal image. For example, the automatic detection of other lesions such as microaneurysms, which are the first symptoms to appear in diabetic retinopathy can rely on vessel segmentation [6, 12]. Another usage is for registering a number of retinal images from different perspectives and creating a 3D image of the retinal surface [11, 14]. The registered image can be used to measure the depression in the optic disc for the detection of glaucoma.

Methods for segmenting blood vessels can be roughly divided into those that trace the center of the vessels [2, 4, 8], methods based on learning and classification of feature vectors [9, 10, 13] and segmenting the vessel boundaries by using some set of filters or thresholds [1, 3, 5, 7, 9]. The tracking methods usually find a set of seed points and then use them to trace the retinal vasculature. The algorithm described in [4] is most closely related to our work, as both methods use the Kalman filter for tracking and matched Gaussian filters for detecting the blood vessels. In [4], the seed points are found in the circumference of the optic disc and the center of the vessel is tracked using an extended Kalman filter. Tracing the vessel is made by using matched filters and detecting branching points. Tracing stops when the response from the Gaussian filter is low. However, our method differs in three important key points: 1) We find seed points all over the image and do not rely on the hard task of tracing the vessel for the whole of its length; 2) The next state in the tracking is based not only on the response from the matching filters, but also on the distribution of the gradients in a window around the vessel. It is helpful, for instance, when tracking hollow vessels (where the vessel is brighter than the background); 3) The tracking method is tolerant to vessels with areas of low response to the filters (for example, thin vessels with changing contrast) by ending the tracing only after a few consecutive bad responses. However, segmentations which have little resemblance to a vessel after a few steps are retraced.

2 Proposed Method

An overview of our proposed method is as follows:

1. Find a set of seed points by convolving the image with a group of matched filters in different widths and orientations. Unlike many other algorithms, we try to find as many as possible correct seed points, and do not limit them to certain areas (such as only the optic disc area). The aim is to have at least one seed point at every vessel, therefore removing the need to follow all branches.

*National ICT Australia is funded by the Australian Government's Backing Australia's Ability initiative, in part through the Australia Research Council

2. Track the center of the blood vessel by using the Kalman filter and starting from the set of seed points found in the previous step:
   (a) Estimate the next location by using gradient information from both edges of the blood vessel.
   (b) Refine the estimation by correlating a cross-section of the estimated vessel with the shape of a vessel.
   (c) Estimate the vessel width using the correlation results.
3. Stop the tracking if the likelihood of tracing a vessel is small for a number of consecutive moves or when we hit an already segmented vessel.
4. Retrace the segmentation, if the tracking failed in less than a minimum number of steps.

2.1 Detection of seed points

Matching filters have been used before [3] for the detection of blood vessels, and we use them as well, mainly for the detection of the seed points. A cross-section of a blood vessel will usually have the darkest value in the middle and brighter values as approaching the edges (see Fig. 2(b)). We therefore build the matching filters with a similar shape. The kernel is approximated by using a second order derivative Gaussian filter, $k(x;\sigma)$ as described in [4], where the value of σ is half of the estimated vessel width. Because the vessels are considered to be piecewise linear segments, we create a 2D kernel by using a number of cross-sections. We define a set of 2D kernels with different orientations and widths:

$$k_{ij}(x,y;\theta_i,\sigma_j), \qquad 0 \le \theta_i < 180 \qquad 2\sigma_j = 2,4,6,8, \tag{1}$$

in order to cater for the different vessels' widths and orientations. We use 6 different orientations for θ spaced in 30 degrees from each other, therefore having a total of 24 kernels. The kernel's size is defined as 11×11 when $2\sigma_i < 8$ and 17×17 when $2\sigma_i = 8$, as we expect to have a longer continuous vessel structure at the same orientation for wide vessels. The kernels' size is based on measuring the minimum and maximum widths of the blood vessels in the given dataset. The kernel is normalized so the sum of its elements is 0 and variance is 1, therefore higher intensity pixels would not dominate the convolution response. We then convolve each of the 6 oriented filters for a given width with the image:

$$C(x,y;\theta_i) = k(x,y;\theta_i,\sigma) * f(x,y), \qquad 1 \le i \le 6 \tag{2}$$

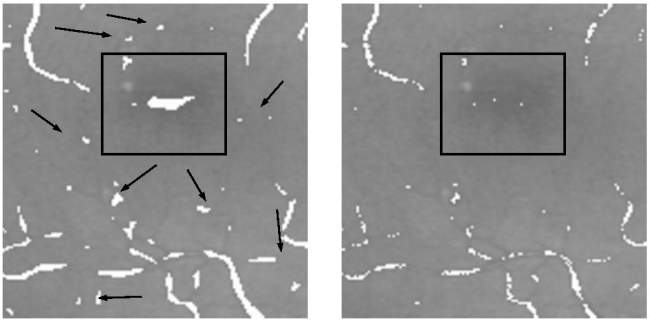

(a) (b)

Figure 1. *The detected seed points around the macular area (in white): (a) after using Eq. 3 (b) after further using gradient information. The macula is framed by the black square and the arrows point to incorrect seeds. Most of the erroneous seeds are removed, and the rest should not be developed at the tracking step. Some legitimate seeds have been removed as well, but they are still surrounded by seed points. The final segmentation of this image appears in the third column of Fig. 5.*

and then only the responses that are above a threshold in a certain direction, but have weaker response in the orthogonal direction are taken:

$$C(x,y;\theta_i) > T_s \text{ and } C(x,y;\theta_j) < T_s, \quad |\theta_i - \theta_j| = 90. \tag{3}$$

However, the threshold cannot be too high, because then no seed points are found in thin, low contrast vessels. It is generally preferable to use a mild threshold and then remove the false seed points during the tracking (see section 2.3). For each pixel found by a kernel $k_{ij}(x,y;\theta_i,\sigma_j)$, a cross-section in the direction normal to θ_i and of length $2\sigma_j$ is built. As we expect the pixel to be in the vessel's center, the gradients' angles at each end should be roughly in opposite directions. Finally, a small percentage of the brightest seed points is removed to avoid taking seeds that are between bright lesions or on the optic disc. The whole process is repeated for the 4 different vessel widths given by σ. For each seed point we record the vessel's estimated width and orientation to be used for the tracking. This step is usually quite fast and takes only about 10 percent of the running time. An example of the seed points that are found around the macular area (which does not contain vessels) after using (3) is shown in Fig. 1(a). The remaining seed points in this area after applying further validation as described before are shown in Fig. 1(b). The rest of the erroneous seeds are expected to

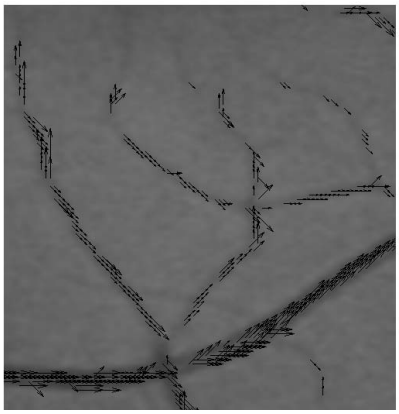

(a)

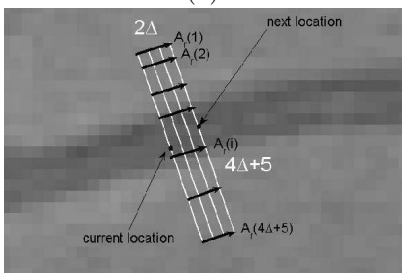

(b)

Figure 2. *(a) A zoomed area of the vasculature from Fig. 5(a) middle row, showing the detected seed points (as vectors) and their orientation. The vector's length is related to the estimated vessel's width. (b) A blood vessel and its cross-section painted by the white lines. The black lines show the direction the pixels are averaged to create the vector A_r as explained in Sec. 2.3. The next state found by the tracker after the validation step is pointed out.*

be removed during the tracking step, so they do not develop to vessel segments. In Fig. 2(a), a magnified area around the main vessel is shown with the seed points found and their orientation and length.

2.2 The Kalman Filter

Starting from a detected seed point with an estimated vessel width and orientation, the aim is to predict the trajectory of the center of the vessel. By considering the vessel tracking as a time series, where each state is a location in the vessel, it can be modelled by a state-space approach, such as the Kalman filter. The Kalman filter has been widely used to solve computer-vision related time series problems, such as road tracking, where the center of the road has to be detected [15]. Our method uses the traditional Kalman filter as the basis for the tracking, but we add additional error models that fit better to the problem in hand.

At each step, the measurement is computed by examining the vessel structure (see Sec. 2.3), and deciding the next location and vessel orientation by combining the prediction and the measurement. The evolution of the tracking process is defined by:

$$\mathbf{x_k} = f_k(\mathbf{x_{k-1}}) + \mathbf{w_{k-1}},$$

where $\mathbf{x_k}$ is the state vector at time k and $\mathbf{w_k}$ the system noise and $\mathbf{f_k}$ is a function of the state. Given a measurement $\mathbf{z_k}$, the relationship between the state and the measurement is given by:

$$\mathbf{z_k} = h_k(\mathbf{x_k}) + \mathbf{v_k},$$

where $\mathbf{v_k}$ is the measurement noise. The system noise and the measurement noise are assumed to be indpendent with covariance matrices Q_k and R_k.

For a given seed point with estimated vessel width 2σ and orientation θ, the current state $\mathbf{x_k}$ and the transition matrix are given by:

$$\mathbf{x_k} = [x, y, \sin(\alpha), \cos(\alpha)]^T, \qquad A = \begin{bmatrix} 1 & 0 & \Delta & 0 \\ 0 & 1 & 0 & \Delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

so the prediction for the next state is $\mathbf{x_k} = A\mathbf{x_{k-1}}$. The current tracking direction is given by α, and Δ represents the move size in each step of the filter, and is initialized to $\Delta = \sigma$. Therefore the prediction for the next step is to move Δ pixels in the same direction as we moved before. The use of 4 elements in $\mathbf{x_k}$ (and not 3), is to be able to use the linear version of the Kalman filter. The measurement vector $\mathbf{z_k}$ which is the observation for the next vessel center, and H, the measurement model, are given by:

$$\mathbf{z_k} = (x, y)^T, \qquad H = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

The matrix H relates the state $\mathbf{x_k}$ to the measurement $\mathbf{z_k}$. We use a simple model with $H\mathbf{x_k}$ just returning the location in the blood vessel given by the state $\mathbf{x_k}$.

The system noise covariance matrix Q, is dependent on the vessel width, where for thinner vessels the noise in the moving direction is defined to be higher, because small thin vessels can be winding. The measurement noise covariance matrix R is set to 1 pixel in each direction. Both matrixes are fixed during the tracking and are given as follows:

$$Q = \begin{bmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.1 & 0 & 0 \\ 0 & 0 & e_w & 0 \\ 0 & 0 & 0 & e_w \end{bmatrix}$$

$$e_w = \begin{cases} 0.6 & \text{if } 2\sigma \leq 4\text{pixels} \\ 0.2 & \text{otherwise} \end{cases}$$

$$R = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

At each step the measurement returns an error value, which is 0 if it seems that the tracking still follows a vessel and 1 otherwise. The tracing stops only after 3 consecutive errors, therefore letting the tracing continue in locations of uncertainty such as low gradient information, branching or fast change in direction. This empirical value proved to be suitable for our application, because by continuing an additional two tracing steps, the problematic section can be passed.

Finally, we define a value of the minimum number of steps that have to be completed by the tracker, in order to keep the segmentation. The aim is to remove segmentations from seed points which are not within a vessel at all and after a few moves no vessel texture is found. This can happen, for example, when the seed points are on the macula, on the optic disc or between two close vessels. We define this value to equal to 9 steps, which is three times the number of accepted errors before terminating. We found out that this value will remove most of the erroneous vessels, but sometimes can also retrace short sections of valid vessels. When a decision to retrace the segmentation from a seed point is made, the whole segment is removed. This is equivalent to treating the seed point as a false one (or not detecting it at first place).

2.3 Vessel Tracking

In order to predict the next state, the Kalman filter calls a measurement function that estimates the next location, given the current location. The function performs a number of correctness checks and if any one fails, an error status is returned to make a decision whether to terminate the tracking.

2.3.1 Gradient computation.

A window centered around the current location x_k in the blood vessel is created. The window is of size $(2w+1)\times(2w+1)$, assuming the vessel estimated width is w. The gradients G_x,G_y in the x and y directions respectively are computed in the window. We are interested in finding the orientation which has the strongest magnitude response and assume the vessel orientation is perpendicular to it. Therefore, the gradient's angle and magnitude are given by:

$$\theta = \arctan(G_y/G_x), \qquad \text{mag}(x,y) = \sqrt{G_x^2 + G_y^2} \quad (4)$$

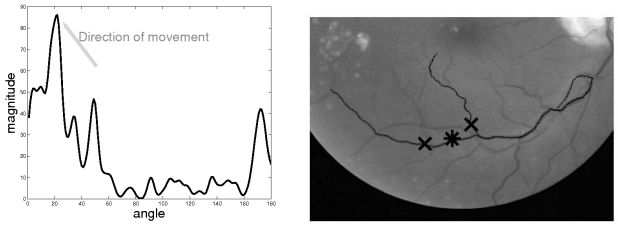

(a) (b)

Figure 3. *(a) Histogram of the magnitudes of the double angle for the point marked with a star in part (b). In this histogram the double angle is divided by two, therefore retrieving the original angle from 0 to 180. The highest peak is the estimated vessel's orientation. (b) Example of tracking the vessels using two seed points marked by X. The left seed has an initial width estimation of 7 pixels and the right one of 3 pixels. Only the center of the vessel has being segmented, so it is possible to see how the segmentations follow the center of the vessels and merge into one line and then separate again. Also the hollow vessel on the upper right hand side has been traced.*

However, in order to take advantage of the gradient information from both edges of the vessel, we compute the double gradient angle, 2θ, for each pixel. The reason is that gradients of points on opposite edges of the vessel will differ by 180 degrees, so doubling the angle will make them equal and reinforce each other.

2.3.2 Histogram computation.

In order to find the best orientation for the next state, we use a similar method to the Parzen's window method. Our random variable is the double angle ranging from 0 to 360 degrees. The sample points are the double angles computed inside the window. For each angle we sum the corresponding pixel's magnitude as given by (4) (and not just the number of occurrences of the angle). The histogram is smoothed by convolving with a Gaussian filter with a small variance. As we are looking at a section of a vessel we expect the histogram to have only one peak with a probability distribution exceeding a threshold, as most vessels have stronger gradient information than the surroundings (see Fig. 3(a)).

2.3.3 Validation by cross-section.

The angle θ_m of the strongest peak in the histogram is taken. This angle is perpendicular to the vessel's ori-

entation, so is used to build a cross-section of the blood vessel at the current location x_k to validate and refine our estimation for the orientation. We use a similar idea to the one used to built the matching filters: A single cross-section is built as a vector of intensities of length $4\Delta + 5$ and orientation θ_m. Utilizing the piecewise linearity of the vessel, and looking at a set of 2Δ parallel cross-sections, a matrix of size $2\Delta \times (4\Delta + 5)$ is created. For each column the average pixel value $A_r(x)$ is calculated, as shown in Fig. 2(b) by the direction of the black arrows. The length of the cross-section is chosen so it will always span the whole width of the vessel (2Δ), even when the current state is on the vessel's end. The vector A_r is correlated with a Gaussian to find the center of the vessel and the degree of correlation. If the correlation is above a threshold the new orientation is taken as pointing towards the point of maximum correlation (shown in Fig. 2(b) as the next location). Otherwise, the orientation stays the same one as the previous one. An example of tracing the vessels from two seed points until the ending condition is fulfilled is shown in Fig. 3(b).

2.3.4 Estimating the vessel's width.

The averaged cross-section vector A_r is scanned towards both ends of the vector starting from the pixel with the highest correlation. A pixel will be segmented only if:

$$[A_r(x) > \max(A_r)T_3] \text{ and } [A_r(x) > T_4], \qquad (5)$$

where the first term puts a boundary on the highest intensity that is still a vessel and the second one puts a minimum threshold on the difference in intensity from the background. Once a pixel does not comply with (5) the search in that direction is stopped. The thresholds T_3 and T_4 have strong influence on the result as they define the sensitivity to low contrast vessels and to pixels on the vessel's boundary. We found out that different thresholds give better results with different manual segmentations. We also incorporate our initial estimation of the vessel's width to control the number of pixels segmented.

3 Results

We tested our method on the public database of DRIVE [10], which includes 40 images: 20 training images and 20 test images. For the test images, 2 manual segmentations are given. Previously published results use the 20 testing images and compare their results to the first manual segmentation. We used the same images and manual segmentations and also the provided

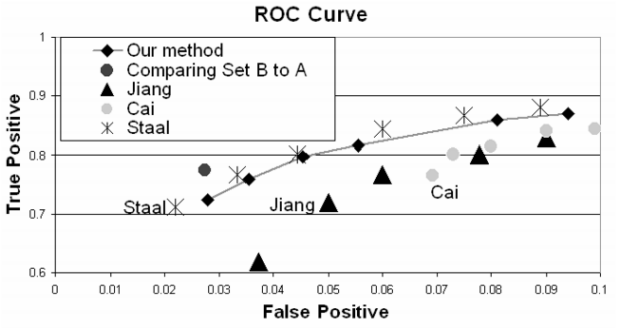

Figure 4. *ROC curve of our algorithm. Our results are drawn by the continuous line. Other methods are plotted with unconnected points.*

mask image for the retinal area to calculate the false positives (FP). Fig. 4 shows the ROC curve for our algorithm. In order to induce the graph the parameters T_s, T_3 and T_4 were set and changed in (3) and (5). Their values were tested by running experiments on all images and estimating the range which provides valid results. As the manual segmentations are subjective, and there is only a 78% agreement between the two given manual segmentations, it makes sense to have parameters that will be finalized according to the given database. The threshold in (3) determines the number of seed points that will be detected and whether low contrast vessels will have seed points in them. The thresholds in (5) set the sensitivity when segmenting the vessel's width and mainly affect the number of boundary pixels that will be segmented.

In the ROC curve we only show the part where the false positives are lower than 10 percent, since for high FP values the results become almost useless, since the segmented widths of the vessels are very inaccurate and also large areas of non-vessels are segmented. This is especially true in the case of retinal images as the number of non-vessels is much higher than the number of vessels (up to 8:1). We managed to produce almost 80% of true positives (TP) with 4.5% of false negatives and 85.9% with 8.1% using a different set of parameters. Results of four images with their manual segmentations are shown in Fig. 5. Most of the main and some small vessels are segmented and there are almost no segmented lesions or other retinal structures. The rightmost image in the figure is probably one of the most challenging images in the dataset, as it contains lesions and hollow vessels, which are hard to segment, since they are not darker than the background as normal vessels are. The last row shows the TP, FP and False Negatives (FN) in different colors over a white

background. Most of the FP (green) are around actual vessels, therefore the vessels have been segmented, but do not share the exact border with the manual segmentation. Only small areas of FP are other retinal structures. The main difficulty is in segmenting the very thin vessels (red pixels). Some of these vessels might have seed points in them, but they do not develop due to poor gradient information. Also not all of these vessels appear in the second manual segmentation. It takes between 20sec to 40sec to segment an image, depending on the threshold used in (3) and the number of seed points found for the specific image.

We compare our results with recent publications that use the DRIVE database. All the results from the competing approaches were copied or estimated from the curves provided in the original publications, and are drawn as unconnected points in the ROC graph. Cai [1] computed the distance between the ROC curve to the ideal point (0,1) and compared their results to Jiang [7] and Staal [10]. Since their graph is plotted for FP < 0.5, it is not necessary that the best distance is found for small value of false positives. However, for the area where FP < 0.1 our results are superior to both the methods of Jiang and Cai, as can be seen in Fig. 4. The results presented by Staal [10] are slightly superior to ours. However, their method is based on machine learning and requires a lengthy training time for accumulating data and hand segmenting the images by clinicians. Furthermore, retraining may be necessary if the imaging method or technology changes. This is not needed by our method. Finally, we compared our results with the second manual segmentation and found out that with all sets of parameters the TP rate was 2-2.5% higher with almost the same rate of FP. This result agrees with the subjectiveness of the manual segmentation.

4 Conclusion

We showed in this paper that combining the tracking method of the Kalman filter with multiple seeds spread out over all the image and retracing those that seem to be outside a vessel provides good segmentation results. We managed to segment thin vessels and avoid segmenting other retinal structures. It will be interesting to improve our method for segmenting the vessel's width to be more accurate as well as testing on other databases.

References

[1] W. Cai and A. Chung. Multi-resolution vessel segmentation using normalized cuts in retinal images. *MICCAI*, pages 928–935, 2006.

[2] A. Can, H. Shen, J. Taylor, H. Tanenbaum, and Roysam. Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms. *IEEE Trans. Info. Tech. in Biomedicine*, 3(2):125–138, 1999.

[3] S. Chaudhuri, S. Chatterjee, N. Katz, N. Nelson, and M. Goldbaum. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Trans. Med. Imag.*, 8(3):263–269, 1989.

[4] O. Chutatape, L. Zheng, and S. Krishnan. Retinal blood vessel detection and tracking by matched gaussian and kalman filters. *IEEE Eng. in Med. and Bio. Soc.*, 20:3144–3149, 1998.

[5] A. Hoover, V. Kouznetsova, and M. Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Trans. Med. Imag.*, 19(3):203–210, 2000.

[6] K. Huang and M. Yan. A local adaptive algorithm for microaneurysms detection in digital fundus images. In *CVBIA*, pages 103–113, 2005.

[7] X. Jiang and D. Mojon. Adaptive local thresholding by verification-based multithreshold probing with application to vessel detection in retinal images. *PAMI*, 25:131–133, 2003.

[8] P. Quelhas and J. Boyce. Vessel Segmentation and Branching Detection using an Adaptive Profile Kalman Filter in Retinal Blood Vessel Structure Analysis. In *IbPRIA 2003, LNCS*, volume 2652, pages 802 – 809, 2003. IDIAP-RR 03-58.

[9] M. Sofka and C. V. Stewart. Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures. *IEEE Trans. Med. Imag.*, 25(12):1531–1546, 2006.

[10] J. Staal, M. Abramoff, M. Niemeijer, M. Viergever, and B. van Ginneken. Ridge based vessel segmentation in color images of the retina. *IEEE Tran. Med. Imag.*, 23(4):501–509, 2004.

[11] C. Stewart, C. Tsai, and B. Roysam. The dual-bootstrap iterative closest point algorithm with application to retinal image registration. *IEEE Trans. Med. Imag*, 22:2003, 2003.

[12] T. Walter and J. Klein. Automatic detection of microaneurysms in color fundus images of the human retina by means of the bounding box closing. In *ISMDA '02: Proceedings of the Third International Symposium on Medical Data Analysis*, pages 210–220, London, UK, 2002. Springer-Verlag.

[13] L. Wang, A. Bhalerao, and R. Wilson. Analysis of retinal vasculature using a multiresultion hermite model. *IEEE Trans. Med. Imag.*, 26(2):137–152, 2007.

[14] F. Zana and J. Klein. A multimodal registration algorithm of eye fundus images usingvessels detection and hough transform. *IEEE Trans. Med. Imag*, 18:419–428, 1999.

[15] J. Zhou, W. Bischof, and T. Caelli. Road tracking in aerial images based on human-computer interaction and bayesian filtering. *SPRS journal of photogrammetry and remote sensing*, 61(2):108–124, 2006.

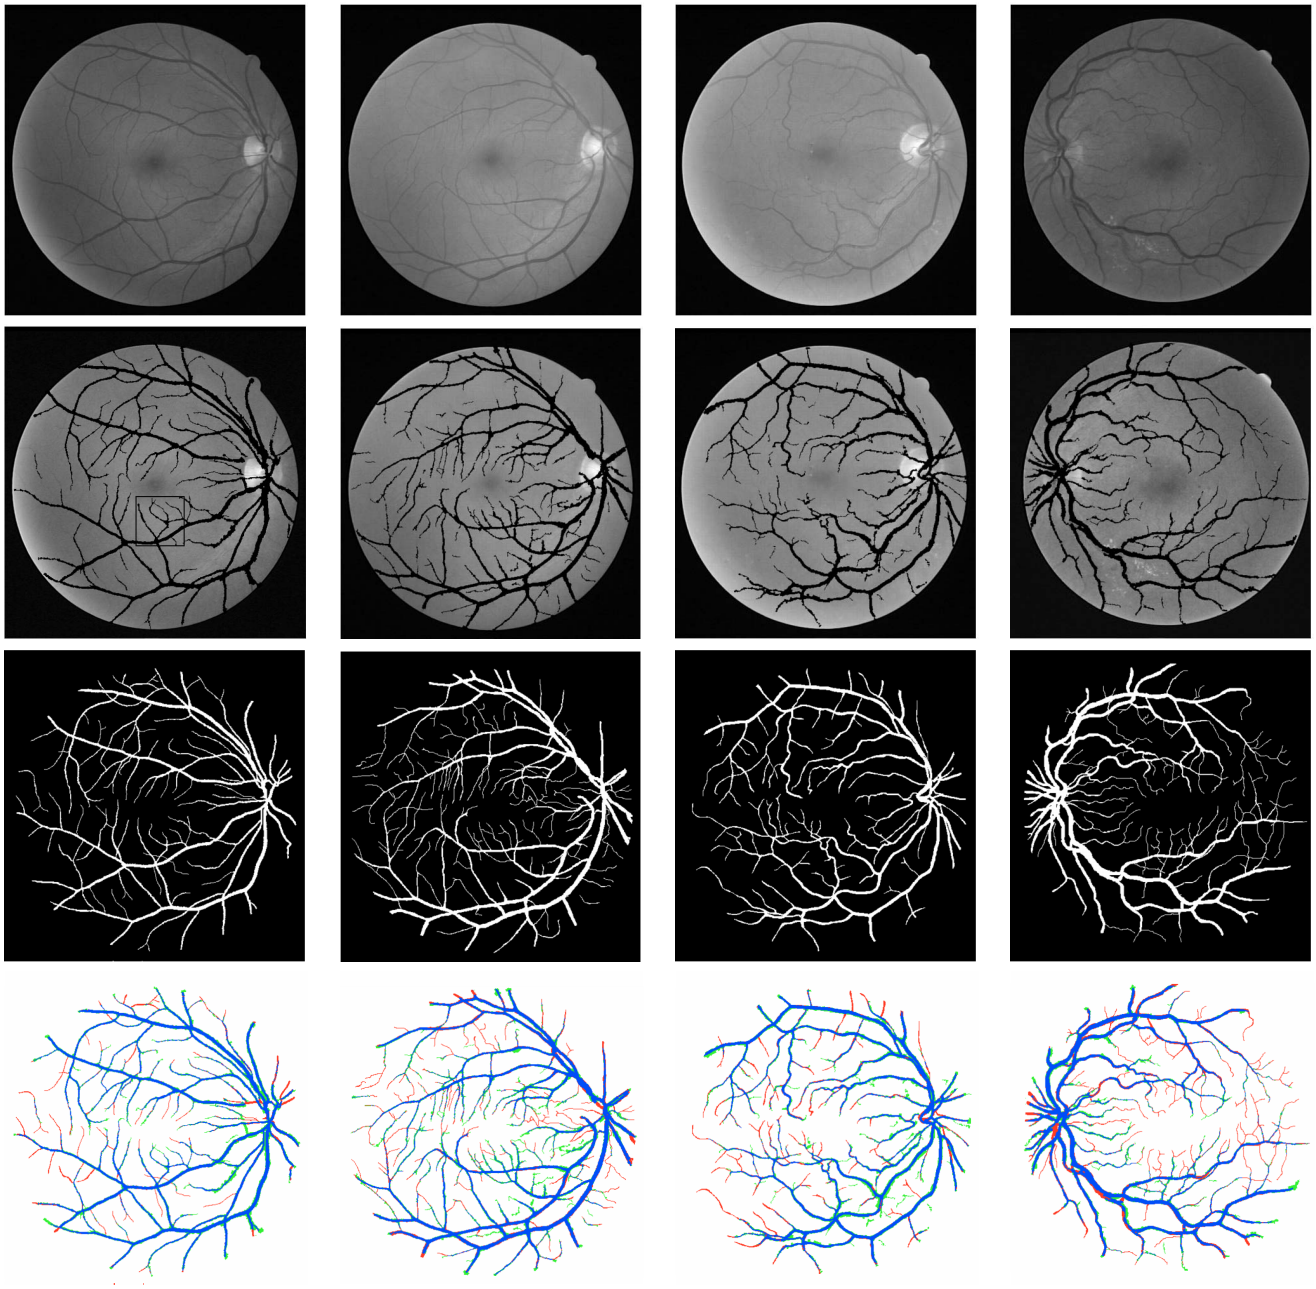

Figure 5. *Some of our results (middle row) for 80% TP and 4.5% FP and the manual segmentations (bottom row). Our segmentation appears as black pixels on top of the original image (top row). In columns (a) and (b) the images are of healthy retinas. In column (c) the image has hollow vessels, which are hard to segment. In column (d) the image contains lesions, hollow vessels and many thin vessels. The last row shows the TP in blue, the FN in red and FP in green. Most of the green pixels appear next to the manually segmented vessels and only few are other structures. The red pixels are mainly very thin vessels, which are hard to segment with a low FP rate. These vessels probably require a different segmentation approach, which is more sensitive.*