A Linking Framework for Pixel Classification Based Retinal Vessel Segmentation

Meindert Niemeijer[a,b], Bram van Ginneken[c], and Michael D. Abràmoff[a,b],

[a]Dept. of Electrical and Computer Engineering, and
[b]Dept. of Ophthalmology and Visual Sciences, The University of Iowa, Iowa City, IA, USA
[c]Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands

ABSTRACT

Retinal vessel segmentation is a prerequisite for the analysis of vessel parameters such as tortuosity, variation of the vessel width along the vessel and the ratio between the venous and arterial vessel width. This analysis can provide indicators for the presence of a wide range of diseases. Different types of approaches have been proposed to segment the retinal vasculature and two important groups are vessel tracking and pixel processing based methods. An advantage of tracking based methods is the guaranteed connectedness of vessel segments, in pixel processing based methods connectedness is not guaranteed. In this work an automated vessel linking framework is presented. The framework links together separate pieces of the retinal vasculature into a connected vascular tree. To determine which vessel sections should be linked together the use of a supervised cost function is proposed. Evaluation is performed on the vessel centerlines. The results show that the vessel linking framework outperforms other automated vessel linking methods especially for the narrowest vessels.

Keywords: retina, vasculature, supervised, vessel, linking

1. INTRODUCTION

Retinal vessel segmentation is an important technique as the segmented retinal vessel tree can be used for many different applications. Analysis of vessel parameters such as tortuosity, variation of the vessel width along the vessel and the ratio between the venous and arterial vessel width can provide indicators for the presence of a wide range of diseases. Examples are hypertension, retinopathy of prematurity and glaucoma. Features, such as the location of bi-furcations, extracted from the vessel tree can be useful in the registration of retinal images. In recent years the segmentation of the retinal vasculature in retinal photographs has received a lot of attention. This has been helped by the public availability of sets of retinal images with a reference standard segmentation of the vasculature such as the STARE[1] and DRIVE[2] databases.

Different types of approaches have been proposed to segment the retinal vasculature and two important groups are vessel tracking and pixel processing based methods. In vessel tracking based methods feature points that are located on the vasculature are detected in the image. From these feature points the vasculature is traced based on local image properties. Many approaches determine the vessel width and orientation locally and use these parameters to predict where the next piece of vessel can be found. A number of these types of methods have been reported in the literature.[3–6] An advantage of the vessel tracking approach is that the final vessel network is connected which facilitates further analysis of vessel parameters. For pixel processing based methods, a number of features are measured for every individual pixel in an image. Based on the values of the features each pixel is assigned a label that indicates the probability that the pixel is inside a vessel. Various supervised[2,7,8] and unsupervised[9–12] pixel labeling frameworks have been proposed. Some of the best performing, state-of-the-art, retinal vessel segmentation methods are based on this paradigm. Since each pixel is classified independently, a disadvantage of such methods is that gaps can appear in the segmented vessels (see Figure 1) resulting in an unconnected vessel network.

The detection of narrow, low contrast vessels is a problem for both types of automatic segmentation methods. Because of the way in which automated segmentation methods have been evaluated the lower performance on

Send correspondence to Meindert Niemeijer. E-mail: meindert-niemeijer@uiowa.edu

Medical Imaging 2009: Biomedical Applications in Molecular, Structural, and Functional Imaging, edited by Xiaoping P. Hu, Anne V. Clough, Proc. of SPIE Vol. 7262, 726216 · © 2009 SPIE · CCC code: 1605-7422/09/$18 · doi: 10.1117/12.811653

small vessels does not have a large influence on the final reported results. The standard evaluation method in the literature is a receiver operator characteristic (ROC) analysis. It plots the sensitivity (i.e. the true positive rate) against 1-specificity (i.e. the false positive rate) of an algorithm illustrating the ability of the algorithm to distinguish vessel from non-vessel pixels. This pixel based measure is essentially unfair as it weighs pixels from wider vessels more heavily than the harder to detect small vessels simply because there are more pixels present in the wider vessels. Often a second human observer made segmentation is used to compare the algorithm performance with that of a human observer. For a human observer it is only practical to make a single binary segmentation and its performance is thus represented by a single point in the ROC plot. The fact that the ROC curves of state-of-the-art methods very closely approach the second observer point[2,7] could lead one to believe the retinal vessel segmentation problem is nearly "solved". However, when visually examining the results it is immediately clear that automatic methods, contrary to man-made segmentations, miss a large number of the smaller vessels. This indicates a shortcoming of the currently used ROC curve analysis for retinal vessel segmentation evaluation.

In this paper we propose two novel developments. First, a dynamic programming based technique to bridge gaps in vascular structures that have been detected by a pixel classification based vessel segmentation method is proposed. The method uses a supervised cost function that learns from training examples to assign a cost to the connection of two vessel segments. Especially for smaller vessels this results in additional vessel pixels being discovered as these small vessels are linked to the main vessel network. The output of the proposed technique is a fully connected vessel network that allows for further analysis of the vascular tree. Second, we propose to use an evaluation measure based on the corresponding boundary maps framework, this avoids the unwarranted emphasis of standard pixel-based ROC analysis on thick vessels.

2. METHODS

2.1 Vessel Linking

The proposed method uses a retinal vessel segmentation based on the pixel classification technique.[13] In this method a large set of randomly sampled pixels from the training set was used to train a kNN classifier to distinguish between vessel pixels and non-vessel pixels. In this paper a new approach using two kNN classifiers, one for the narrow vessels (C_{narrow}) and one for the wide vessels (C_{wide}), is presented. An image is segmented twice, different thresholds are applied to both resulting probabilistic segmentations and a novel vessel linking framework is applied to link resulting vessel pieces together.

The method starts with the training set reference standard. Each binary reference standard image I_{ref} is thinned. For each resulting vessel centerline pixel i, the vessel width w_i is determined by probing from i along a line perpendicular to the vessel in both directions. The distance between the two points where the linearly interpolated pixel value falls below 0.5 defines w_i. Vessels with $w_i \geq 2.5$ pixels are considered "wide" while vessels with a width below that are "narrow". The threshold of 2.5 pixels was determined empirically. Each vessel pixel in I_{ref} needs to be assigned a width based on the width values associated with the vessel centerline pixels. For each reference standard vessel pixel the closest centerline pixel is found and the width value of that pixel is assigned to the vessel pixel. To obtain separate reference standard images for the wide as well as the narrow vessels we simply apply a width threshold to the image.

Using these two sets of training reference standard images the kNN classifiers C_{narrow} and C_{wide} are trained. Features are extracted from the green plane of each of the original color fundus images in the training set. The features consist of Gaussian filterbank outputs up to and including second order derivatives at different standard deviations (scales) $\sigma = 0.75, 1.20, 1.93, 4.98, 8$. Then we randomly sample a number of pixels and their associated feature vectors from each of the training images. Those training samples that are labeled non-vessel are used in the training of both classifiers while positive samples from wide vessels are only used in the training of C_{wide} and positive samples from narrow vessels are only used for C_{narrow}.

Once both classifiers are trained they can be applied to any image and generate probabilistic vessel segmentations. "Gaps" in the vessels are present, especially in the small vessels. These gaps are caused by errors of the pixel classification or the fact that the contrast for the smallest vessels is not constant over the length of the vessel. Often the posterior probability in the gaps is slightly above 0 but lower than the surrounding vessel

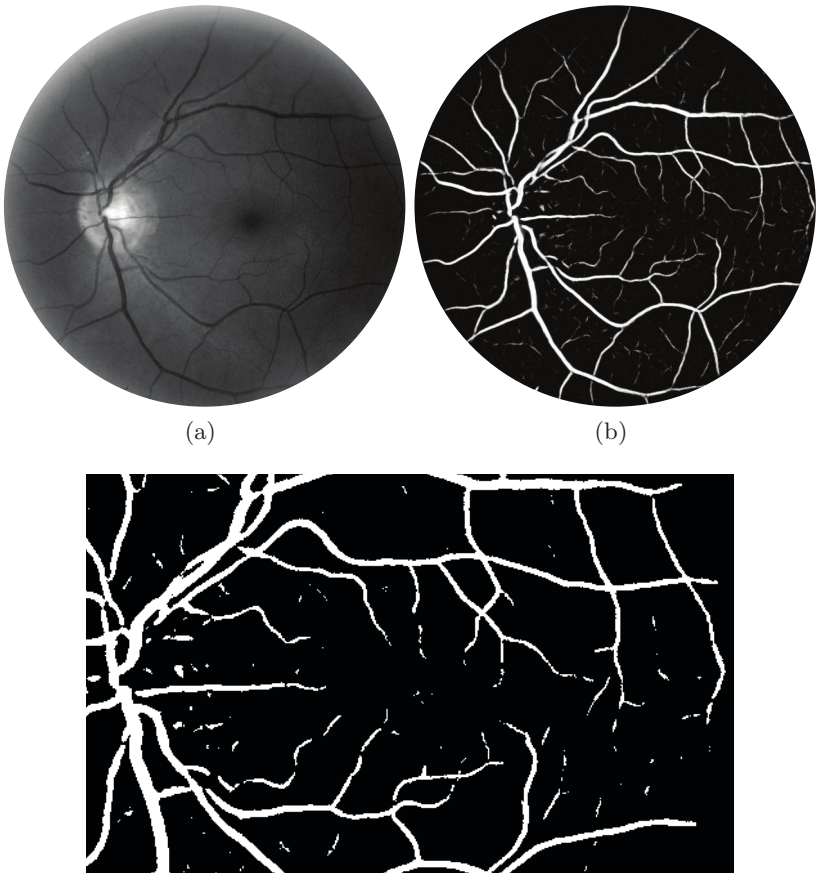

Figure 1. Typical pixel classification based vessel segmentation showcasing the problem with disconnected vessels. (a) Green plan of a digital color fundus photograph. (b) Posterior probability map, the pixel intensity corresponds to the likelyhood the pixel is within a vessel. (c) A detail from the vessel probability map after a threshold has been applied. Note the many disconnected vessels.

pixels. To generate a binary segmentation a threshold is set and thus these pixels are not included. As the wide vessels usually are easier to detect they tend to have a higher posterior probability and as such a higher threshold $T = 0.6$ can be used to make a binary segmentation. For the narrow vessels a lower threshold of $T = 0.15$ is used. Both thresholds were determined empirically, the final algorithm outcome does not depend heavily on these values as long as the narrow vessel segmentation threshold is chosen low enough to include enough parts of narrow vessels.

After generating the binary segmentations a thinning operation is applied to obtain the vessel centerlines. Any overlapping vessel centerlines that are present in both the wide and narrow segmentation result are consolidated into one centerline image $I_{centerline}$. Then the endpoints of all vessel pieces are determined by scanning over all centerline pixels and storing in set S those that have only one neighbor. Each of these vessel piece endpoints $s \in S$ also has an angle s_α associated with them that represent the local angle of the vessel piece with respect to the x-axis. The $s \in S$ will form the starting points for the vessel linking procedure.

1. For each $s \in S$ do:

a. Find all centerline pixels i in $I_{centerline}$ closer than distance D and at an angle not greater than $|\frac{1}{4}\pi|$ from s_α.

b. Use dynamic programming to determine the optimal paths in I_{narrow} between each $s \in S$ and the centerline pixels found in **a**. Store the found paths p in set P.

Nr.	Feature description
1	Avg. value of I_{narrow} under p.
2	Standard deviation of the value of I_{narrow} under p.
3	Avg. value of I_{narrow} under the originating vessel piece.
4	Avg. value of I_{narrow} under the connecting vessel piece.
5	Length of p in pixels.
6	Length of the originating vessel piece.
7	Length of the connecting vessel piece.
8	Euclidian distance between the start and endpoints of p.
9	Average curvature of p.
10	Difference in orientation between the start of p and the originating vessel.
11	Difference in orientation between the endpoint of p and the connecting vessel.
12	Orientation of a straight line between the start and endpoint of p.

Table 1. The features used by the proposed supervised cost function.

c. Evaluate for each $p \in P$ a cost function F, find the path with the lowest cost. If this cost is lower than threshold T_{path} accept the path and add it to $I_{centerline}$. For an accepted path, remove any pixels from S present in p. In case two paths have the same minimal cost, the path with the smallest angular difference with s_α is chosen. If this provides no resolution, the shortest path is chosen. If multiple paths remain, a random one is chosen.

d. Jump back to 1 unless all $s \in S$ were examined and no path was added.

2. Increase distance D with ΔD and then jump back to 1. However, if a certain maximum distance D_{max} is reached, continue to 3.

3. After connected component analysis, the largest component in $I_{centerline}$ is retained as the final result.

The free parameters in this scheme are the maximum path cost threshold T_{path}, the initial distance D, the value of ΔD and the maximum distance D_{max}.

2.2 Cost Function

One critically important part of the scheme outlined above is the cost function that determines the cost of adding a certain path to the vessel network. We propose to use a supervised cost function that is trained using example paths in a training set. The training procedure needs to be completed once, then the trained classifier can be used to determine the cost of paths in previously unseen vessel segmentation images. The training procedure is very similar to the vessel linking scheme as described in Section 2.1. The major difference is that, since the reference standard is available, the cost function F used in step **c.** is calculated directly using the reference standard. For each of the paths p the cost function can vary between 0 and 1 and is given by $F_{training}(p) = 1 - \frac{n_{overlap}}{n_p}$ where n_p is the total number of pixels in the path and $n_{overlap}$ is the number of pixels in the path that overlap with the reference standard I_{ref}. For each path p we calculate a set of 12 features as shown in Table 1 and together with the value given by $F_{training}$ these features form a training sample that is added to the training set. The training path with the lowest cost is added to $I_{centerline}$, however if no valid path can be found, no action is undertaken and the method continues.

A kNN classifier trained using the training samples will be used to assign costs to previously unseen paths. Given a previously unseen path p_u, and the 12 feature values, an estimate for the path cost can be determined using kNN regression. We estimate the cost function value by looking up the k nearest neighbors of p_u in the 12 dimensional feature space. The final estimated path cost is then determined by $F(p_u) = \frac{1}{k} \sum_{i=1}^{k} cost_i$ where $cost_i$ is the cost associated with nearest neighbor training sample i.

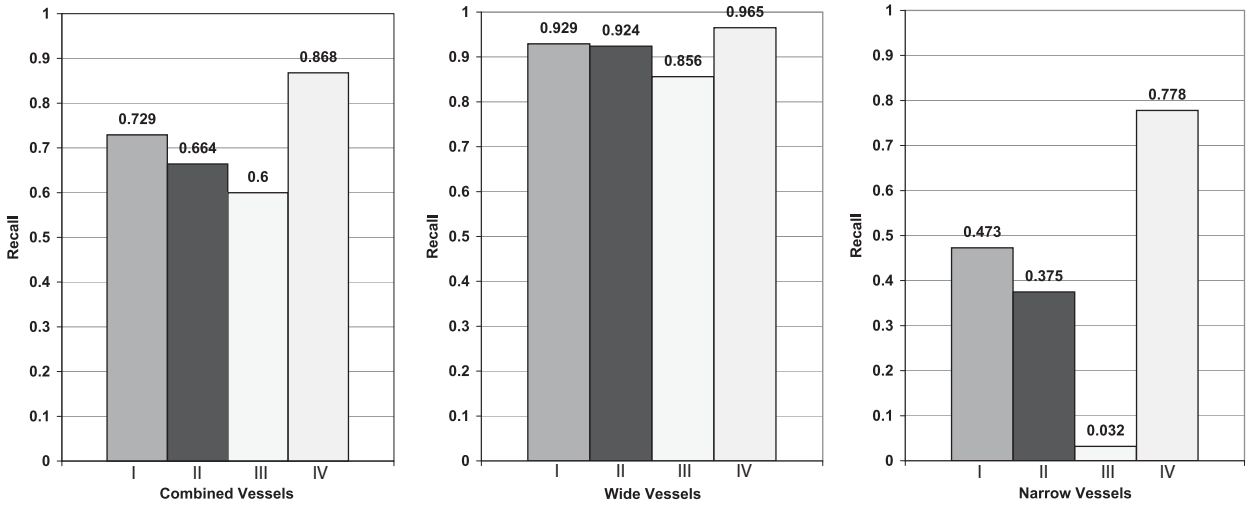

(a) (b) (c)

Figure 2. Each of the above recall graphs shows, from left to right, the vessel linking framework with supervised cost function, the vessel linking framework without supervised cost function, the pixel classification method and the second human observer. (a-c) Evaluation on the complete vascular network, the wide vessels and the narrow vessels. Precision for each graph has been set at the same level as that of the second human observer: 0.88, 0.67 and 0.79 respectively.

2.3 Centerline Based Evaluation

For our experiments, data from the publicly available DRIVE database[2] was used. This dataset consists of 40 color fundus photographs, 20 for algorithm training and development and 20 for testing. One reference standard manual vessel segmentation is available for both the training and test sets, for the test set an additional second observer segmentation is also available. Evaluation was performed by comparing the centerlines from the reference standard of the DRIVE database with the centerlines as produced by the proposed method. The advantage of evaluating centerlines instead of the complete vessels is that most of the bias towards wider vessels is removed when compared with a pixel based evaluation. A disadvantage is that the evaluation is more difficult as vessel centerlines do not necessarily precisely overlap due to small differences in the segmentations. We therefore propose to use the corresponding boundary maps framework as proposed by Martin et al.[14] This algorithm offers an elegant way to match boundary pixels, centerline pixels in our application, of multiple segmentations with each other. The only free parameter of the algorithm is the distance threshold $T_{distance} = 2$ that determines the maximum distance that two matched pixels can be apart from each other.

As the number of negative pixels, i.e. pixels inside the field of view but not part of a vessel centerline, in this application is large compared to the number of positive pixels, specificity is always high. A precision/recall analysis is used which does not include true negative (TN) pixels and is a useful analysis methodology for two class problems when there is a large skew in the class distributions. Precision and recall are defined a follows:

$$precision = \frac{TP}{TP + FP}, recall = \frac{TP}{TP + FN} \tag{1}$$

Where TP =true positives, FP =false positives and FN =false negatives. Following the definition, recall is the same as sensitivity indicating the fraction of all vessel centerline pixels that were classified as vessel pixel in the vessel linking result. The precision indicates the fraction of all centerline pixels in the result that truly were vessel centerline pixels according to the reference standard.

3. EXPERIMENTS AND RESULTS

After training, the described system was applied to the 20 images in the test set. The goals of the proposed method are two-fold; obtain performance as close to that of a human observer as possible and find more connected

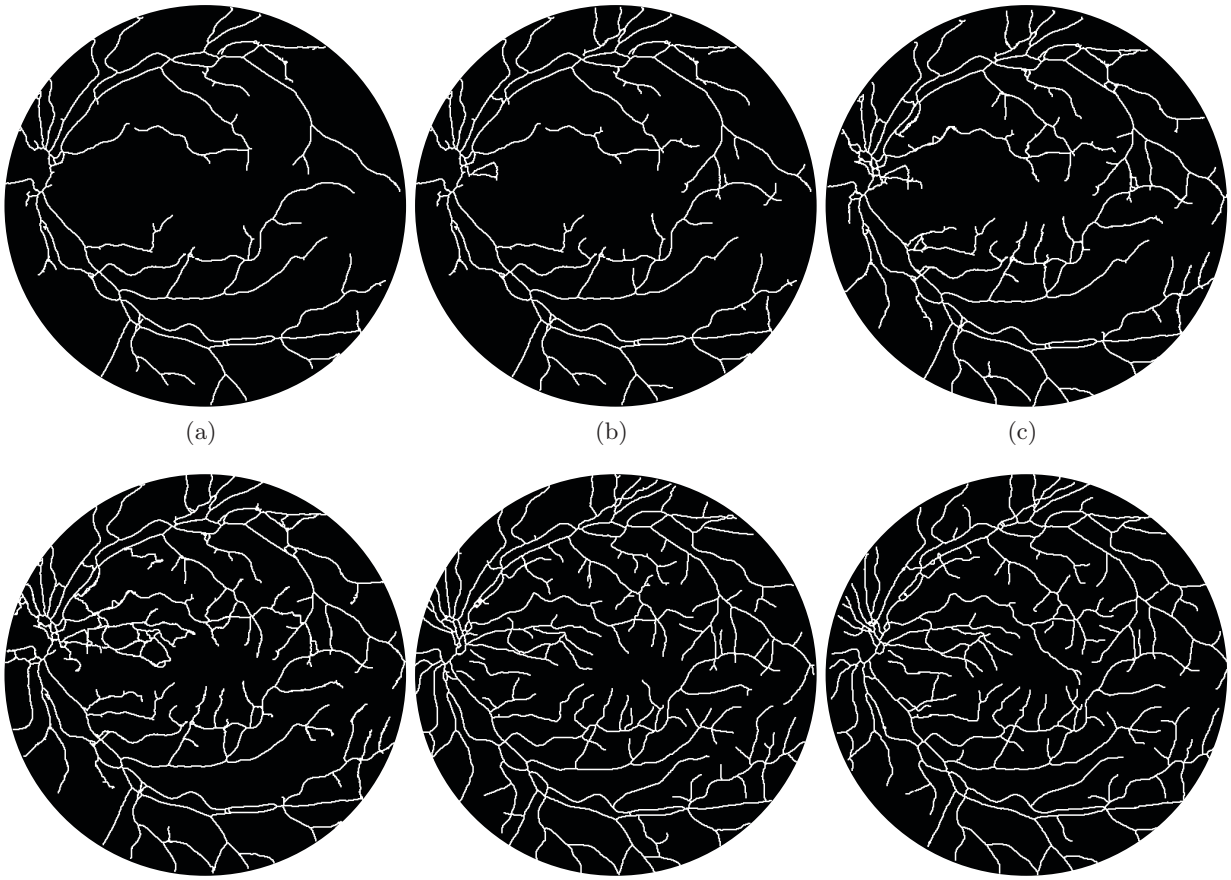

Figure 3. (a-d) Centerline segmentation results obtained by increasing the value of T_{path}. (e) The segmentation result of the second human observer. (f) The reference standard.

vessels than regular pixel classification. Figure 2 shows the results of the proposed method, the proposed method using a simple cost function (i.e. unsupervised), regular pixel classification and a human observer. In order to enable direct comparison the precision for all methods was set equal to that of the second human observer. To show that the framework performance benefits from using a supervised cost function we also compared with a using simple cost function based on averaging the posterior probabilities under p in I_{narrow}. For the comparison with regular pixel classification we applied our vessel segmentation[13] directly to the test images of the DRIVE database and used only the centerlines of the largest connected component for the evaluation.

The above evaluation was performed on both the complete image as well as separately on just the wide (≥ 2.5 pixels) and just the narrow vessels (Figure 2a-c). To visually illustrate the kind of results the proposed method produces, an example segmentation with reference standard at different levels of T_{path} is shown in Figure 3.

4. CONCLUSION AND DISCUSSION

The results of an automatic vessel linking method have been presented. Vessel linking is important as unconnected vessel pieces complicate analysis of the vascular network. The results show that the proposed method with supervised cost function outperforms both regular pixel classification as well as a version of the vessel linking framework using a simple, unsupervised cost function. The second human observer performs better than all automated methods.

The difference in performance between the second observer and the automated methods could be because less vessels are found due to the vessel linking framework failing to link certain vessels as well as the way in which

we evaluated. The performance of the pixel classification based vessel segmentation method we used approaches that of the second human observer when employing regular ROC analysis.[13] But, using centerlines we remove the bias associated with normal ROC analysis. In our test set the wide vessels contained 88340 centerline pixels while the narrow vessels contained 85168 centerline pixels. The performance of the automated systems on the wide vessels was much better and closer to that of the second observer. For the narrow vessels (i.e. nearly half the centerline pixels) the performance difference was much larger. Especially for these, our proposed method showed better results than the other automated methods.

One of the more important issues for future work are vessels that leave the field of view and re-enter it elsewhere. These vessels seem "unconnected" and are not linked by our current algorithm. One possible solution would be to allow linking with the field of view border. Using the connected vessel networks we hope to develop an automatic analysis system that will also detect vessel crossings and bifurcations which will enable separation of the venous and arterial vessel trees. An issue for the presented method is that spurious detections of the vessel segmentation algorithm tend to be linked to the main vessel network at low thresholds. At low thresholds links may be formed between separate vessels while there is no link there. As some of the smaller vessels have real gaps in them the supervised system learns that under certain circumstances such as when linking two vessel segments with high probability it is allowed to assign a relatively high probability to links that have low support in the posterior probability image. Changing the features extracted from the vessel segments or adding additional features may alleviate this issue.

To summarize, an automatic vessel linking system has been presented. The linking of vessel segments into a complete vascular network is an essential step in the automated analysis of the retinal vascular network.

REFERENCES

1. A. Hoover, V. Kouznetsova, and M. Goldbaum, "Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response," *IEEE Transactions on Medical Imaging* **19**(3), pp. 203–210, 2000.
2. J. Staal, M. Abràmoff, M. Niemeijer, M. Viergever, and B. Van Ginneken, "Ridge based vessel segmentation in color image of the retina," *IEEE Transactions on Medical Imaging* **23**(4), pp. 501–509, 2004.
3. A. Y. Tolias and S. Panas, "A fuzzy vessel tracking algorithm for retinal images based on fuzzy clustering," *IEEE Transactions on Medical Imaging* **17**(2), pp. 263–273, 1998.
4. A. Can, H. Shen, J. Turner, H. Tanenbaum, and B. Roysam, "Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms," *IEEE Transactions on Information Technology in Biomedicine* **3**(2), pp. 125–138, 1999.
5. E. Grisan, E. Grisan, A. Pesce, A. Giani, M. Foracchia, and A. Ruggeri, "A new tracking system for the robust extraction of retinal vessel structure," in *Proc. 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society IEMBS '04*, A. Pesce, ed., **1**, pp. 1620–1623 Vol.3, 2004.
6. M. Sofka and C. V. Stewart, "Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures," *IEEE Transactions on Medical Imaging* **25**(12), pp. 1531–1546, 2006.
7. J. V. B. Soares, J. J. G. Leandro, R. M. C. Jnior, H. F. Jelinek, and M. J. Cree, "Retinal vessel segmentation using the 2-d gabor wavelet and supervised classification.," *IEEE Trans Med Imaging* **25**(9), pp. 1214–1222, 2006.
8. E. Ricci and R. Perfetti, "Retinal blood vessel segmentation using line operators and support vector classification.," *IEEE Trans Med Imaging* **26**(10), pp. 1357–1365, 2007.
9. A. Mendonca, A. Mendonca, and A. Campilho, "Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction," *IEEE Transactions on Medical Imaging* **25**(9), pp. 1200–1213, 2006.
10. M. E. Martinez-Perez, A. D. Hughes, S. A. Thom, A. A. Bharath, and K. H. Parker, "Segmentation of blood vessels from red-free and fluorescein retinal images.," *Med Image Anal* **11**(1), pp. 47–61, 2007.
11. L. Wang, A. Bhalerao, and R. Wilson, "Analysis of retinal vasculature using a multiresolution hermite model.," *IEEE Trans Med Imaging* **26**(2), pp. 137–152, 2007.
12. D. Adjeroh, U. Kandaswamy, and J. Odom, "Texton-based segmentation of retinal vessels," *Journal of the Optical Society of America A* **24**(5), pp. 1384–1394, 2007.

13. M. Niemeijer, J. Staal, B. van Ginneken, M. Loog, and M. Abràmoff, "Comparative study of retinal vessel segmentation methods on a new publicly available database," in *Proceedings of SPIE: Medical Imaging*, **5370**, pp. 648–656, 2004.
14. D. Martin, D. Martin, C. Fowlkes, and J. Malik, "Learning to detect natural image boundaries using local brightness, color, and texture cues," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**(5), pp. 530–549, 2004.