Automatic classification of retinal vessels into arteries and veins

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

[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
[d]Veteran's Administration MC, Iowa City, IA, USA

ABSTRACT

Separating the retinal vascular tree into arteries and veins is important for quantifying vessel changes that preferentially affect either the veins or the arteries. For example the ratio of arterial to venous diameter, the retinal a/v ratio, is well established to be predictive of stroke and other cardiovascular events in adults, as well as the staging of retinopathy of prematurity in premature infants. This work presents a supervised, automatic method that can determine whether a vessel is an artery or a vein based on intensity and derivative information. After thinning of the vessel segmentation, vessel crossing and bifurcation points are removed leaving a set of vessel segments containing centerline pixels. A set of features is extracted from each centerline pixel and using these each is assigned a soft label indicating the likelihood that it is part of a vein. As all centerline pixels in a connected segment should be the same type we average the soft labels and assign this average label to each centerline pixel in the segment. We train and test the algorithm using the data (40 color fundus photographs) from the DRIVE database[1] with an enhanced reference standard. In the enhanced reference standard a fellowship trained retinal specialist (MDA) labeled all vessels for which it was possible to visually determine whether it was a vein or an artery. After applying the proposed method to the 20 images of the DRIVE test set we obtained an area under the receiver operator characteristic (ROC) curve of 0.88 for correctly assigning centerline pixels to either the vein or artery classes.

Keywords: retina, vasculature, artery, vein, artery vein separation, artery vein ratio

1. INTRODUCTION

The eye is unique in the sense it is the only place in the human body where the vasculature can be directly observed using visible light. There are many parameters one can measure from the retinal vasculature that can give an indication of the state of the vasculature in general or indicate the presence of abnormalities elsewhere in the body. Amongst these are variations in the vessel width, the tortuosity of the vasculature and the width ratio between the arteries and the veins of the retinal vasculature (a/v ratio). This work is focussed on the a/v ratio as it is well established to be predictive of stroke and other cardiovascular events in adults.

A large number of works on general vessel segmentation have appeared in the literature.[1–12] However, the automated classification of the segmented vasculature in arteries and veins has received limited attention. A semi-automatic method for the analysis of retinal vascular trees in which the venous and arterial trees were analyzed separately was presented by Martinez-Perez et al.[13] A more recent work by Rothaus et al.[14] shows a method to label all vessels as either artery or vein using an existing vessel segmentation and some manually labeled starting vessel segments. The work closest to this one is an automated classification method by Grisan et al.[15] in which the vasculature is segmented using a vessel tracking procedure and the vessel centerlines are detected. After defining an area of interest around the optic disc and dividing this area into four quadrants, color based features are extracted from the vessel segments that are then classified into arteries and veins using an unsupervised clustering method. Results are reported on 24 images in which the total error rate was 12.4%. Recently this work was extended by Ruggeri et al.[16] with an evaluation based on the actual artery and vein

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

Medical Imaging 2009: Computer-Aided Diagnosis, edited by Nico Karssemeijer, Maryellen L. Giger
Proc. of SPIE Vol. 7260, 72601F · © 2009 SPIE · CCC code: 1605-7422/09/$18 · doi: 10.1117/12.813826

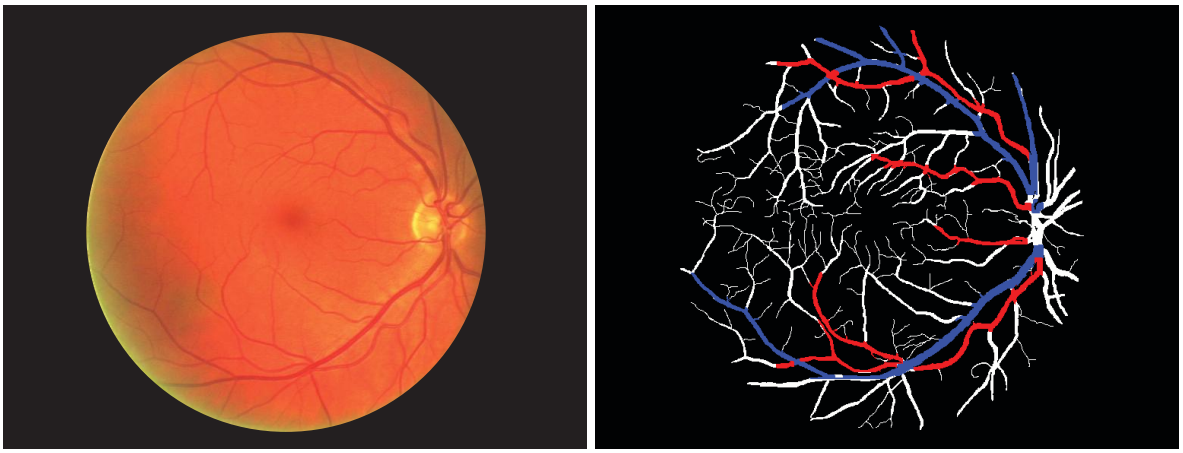

(a) (b)

Figure 1. An example of a retinal image with the annotated reference standard image. a) The original retinal image. b) The reference standard image. All white vessels are vessels for which the expert observer could not see whether they were arteries or veins without determining their connectivity. Red vessels are arteries while blue vessels represent veins. All black pixels represent background.

ratio. The authors report a correlation with a manual reference standard in 14 images that varies between 0.73 and 0.83 depending on how the a/v ratio is calculated. Tramontan et al.[17] further extended the algorithm with enhanced vessel tracking. The reported correlation with the manual reference standard increased to 0.88 (20 images).

The focus of this work is on the classification of vessel segments into arteries and veins. Initially we focus on the larger vessels as they tend to be easier to classify and they don't require the use of vessel connectivity information. Smaller vessels are difficult to classify without tracing them back to determine to which vessel they connect. We do an extensive evaluation of a large set of possibly useful features and attempt classification of vessel segments everywhere in the image where a human observer was able to visually distinguish arteries from veins. The main purpose of the presented work is to serve as a basis for a tool that will automatically measure the a/v ratio as well as a system that uses a completely connected model of the retinal vasculature to determine the artery and vein labels for all vessels in the image.

This paper is structured as follows. In Section 2 the used data is discussed. The different steps of the registration process are presented in Section 3. Section 4 discusses the performed experiments and shows the obtained results. The paper is concluded in Section 5.

2. MATERIALS

The proposed system was developed and evaluated using a set of 40 color fundus photographs, the publicly available DRIVE[1,18] database. It consists of a training set of 20 images that were used for training the system and an independent test set of 20 images used for the evaluation. Each image in DRIVE has been JPEG compressed. The images were acquired using a Canon CR5 non-mydriatic 3CCD camera with a 45 degree field of view (FOV). Each image is captured using 8 bits per color plane at 768 $\times$ 584 pixels. The FOV of each image is circular with a diameter of approximately 540 pixels.

Three observers manually segmented the vasculature in the training and test set. All observers were trained by an experienced ophthalmologist (MDA). The first observer segmented 14 images of the train set while the second observer segmented the other 6 images. The test set was segmented twice resulting in a set X and Y. Set X was segmented by both the first and second observer (13 and 7 images respectively) while set Y was completely segmented by the third observer. The performance of the vessel segmentation algorithms is measured on the test set. In set X the observers marked 577,649 pixels as vessel and 3,960,494 as background (12.7% vessel). In set Y 556,532 pixels are marked as vessel and 3,981,611 as background (12.3% vessel).

To be able to use the DRIVE database data the provided manual vessel segmentations for all 40 images in the database needed to be refined. To enable the training and evaluation of our method an expert (MDA) was asked to assign all vessel pixels that could be visually identified as artery or vein to those classes, using TruthseekerJ. The expert specifically avoided including vessels that could be assigned a label based on their connectivity. Preliminary experiments showed that this was the case mostly for the smaller vessels near the fovea, these vessels are visually indistinguishable from each other especially at the resolution of the images in the DRIVE database. For an example of an annotated reference standard image see Figure 1.

3. METHODS

3.1 Pre-processing

The method assumes a vessel segmentation is already available.[18] First, the centerlines of the vessels are found through thinning[19,20] of the binary vessel segmentation. This binary vessel segmentation can be computer generated but in this case we used the manual segmentation already available for the DRIVE database. After thinning, cross-over and bifurcation points were removed by counting the number of neighbors for all centerline pixels and removing those with more than two neighbors. This is necessary because the vessel width and angle in bifurcations is not well defined or difficult to measure in the case of cross over points. This operation also subdivides the vascular network into a collection of vessel segments. All segments with a length below four pixels are removed.

For each centerline pixel i in each of the segments the width w_i and angle θ_i with the x-axis is calculated. The angle of the vessel for a particular centerline pixel is given by the direction of the largest eigenvector of the covariance matrix of the coordinate of the centerline pixel along with the coordinates of its three connected neighbors to both sides (i.e. 7 coordinates in total). As it is unknown where the vessel begins or ends the range of θ_i is $[0 \ldots \pi]$. Near the end of the vessel segment fewer coordinates are used, 4 for the end pixel. After all θ_i have been determined, the width of the vessel is measured for all i by probing to the left and right of the vessel, perpendicular to θ_i, to find the vessel edge. The vessel edge is defined as the point in the image where the linearly interpolated pixel intensity falls below half the intensity of the centerline pixel i. The distance between the edges of the vessel defines w_i. All i for which no ground truth is available, those vessels that have not been indicated by the expert as either artery or vein, are removed before further processing. For many of the smaller retinal vessels it is impossible to differentiate locally whether they are a vein or an artery.

3.2 Training phase

The proposed method is supervised, i.e. trained with examples. After a one-time training procedure the method can be used to classify previously unseen centerline pixels. The pre-processing procedure as detailed above is applied to all 20 images in the training set. For each training i a set of 24 features is extracted. Table 1 shows the list of extracted features. All features that do not use specific color information were calculated using the green color plane of the original image.

The first five features are normalized with respect to all the other values found found for this feature in a single image. To determine the average Hue, Saturation and Intensity for each i we averaged these values along a line perpendicular to θ_i and inside the vessel. To calculate the steered Gaussian derivatives[21] filter response, a set of Gaussian basis filters are applied to the image. Using the basis filters (e.g. $L_x x$, $L_x y$ and $L_y y$ for the second order Gaussian derivative) responses the filter response at any arbitrary angle can be calculated. This allows us to make these features invariant for the vessel orientation by orienting the filters so they are always perpendicular to the vessel. This also allows for the use of higher order derivatives without prohibitively expanding the total number of features. Higher order derivatives may be useful in this task due to the presence of central reflexes on the vasculature. The central reflexes tend to be much more pronounced in arteries then they are in veins. See Figure 2 for some examples. After sampling the features for each i the appropriate ground truth labels are assigned and all training samples are stored in a training set. We manually excluded the centerline pixels inside the optic disc. By sampling every other centerline pixel in the training set the total training set contains 54046 samples. All extracted features were normalized to zero mean and unit standard deviation.

The training set data is used to train a statistical classifier to distinguish centerline pixels inside an artery from those inside a vein. We tested several different classifiers (i.e. linear discriminant analysis, quadratic

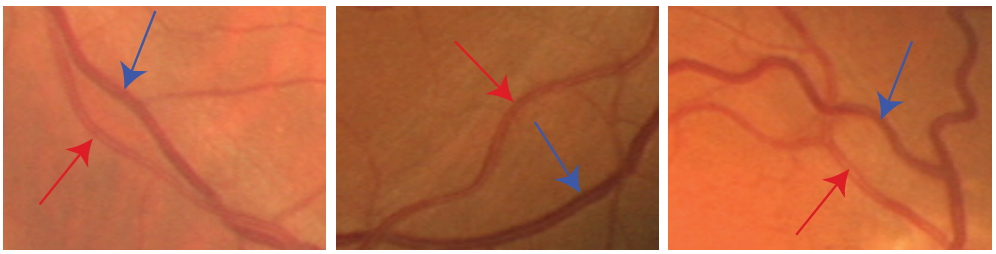

(a) (b) (c)

Figure 2. Three examples of artery vein pairs with central reflex. The artery is indicated with the red arrow while the vein is indicated with the blue arrow. Note the prominent presence of the central reflex in the arteries.

Nr.	**Feature description**
1	Normalized w_i.
2	Normalized vessel contrast.
3	Normalized average hue inside the vessel at i.
4	Normalized average saturation inside the vessel at i.
5	Normalized average intensity inside the vessel at i.
6	Hue of i.
7	Saturation of i.
8	Intensity of i.
9-12	Steered second Gaussian derivative centered on i at angle $\theta_i + \frac{\pi}{2}$ with $\sigma = 1, 2, 4, 8$.
13-16	Steered third order Gaussian derivative centered on i at angle $\theta_i + \frac{\pi}{2}$ with $\sigma = 1, 2, 4, 8$.
17-20	Steered fourth order Gaussian derivative centered on i at angle $\theta_i + \frac{\pi}{2}$ with $\sigma = 1, 2, 4, 8$.
21-24	Value of i in the Gaussian blurred green plane of the original image with $\sigma = 1, 2, 4, 8$.

Table 1. The complete set of features extracted for each centerline pixel for the proposed artery/vein classification method. After feature selection features 1-5, 8, 12-15, 17 and 19 were selected.

discriminant analysis, support vector machine and k-Nearest Neighbor (kNN) classifier) on the training data and found that the kNN classifier provided the best overall performance measured in area under the ROC curve. The optimal value of parameter k was determined to be 286 on the set of the training data.

Feature selection is often used in pattern recognition to reduce the dimensionality of the feature space and select salient features from a larger set. We split the 20 training images into a feature selection training and testing set. To select the best features, a wrapper based feature selection method, sequential forward floating selection,[22] was used. This method starts with an empty feature set, and adds or removes features when this improves the performance of the classifier. The area under the Receiver Operator Characteristic (ROC) curve was used as the performance criterion. After application of this method a set of 12 features, 1-5, 8, 12-15, 17 and 19 were selected. The classifier that was to be used for the final experiments on the test data was retrained on the complete training data set using only these 12 selected features.

3.3 Application phase

After the training phase the trained classifier can be applied to previously unseen images. As with the training images we assume a vessel segmentation is given and we manually excluded those centerline pixels inside the optic disc. To obtain the vessel segmentation we used the reference standard vessel segmentation provided with the DRIVE database. These vessel segmentations were pre-processed as described in Section 3.1. For each centerline pixel i we extract the set of 12 features that were selected during the training process. The resulting feature vector is assigned a soft label by the classifier. The soft label indicates the probability the centerline pixel belongs to a vein. The label is assigned by counting the number of vein pixels amongst the $k = 286$ nearest neighbors in the feature space and dividing this number by k.

After all centerline pixels have been assigned a soft label we use the fact that connected centerline pixels in a vessel segment are all of the same type by definition. The final soft label assigned to a centerline pixel thus depends on the labels assigned to the other centerline pixels in the vessel segment. We compared taking

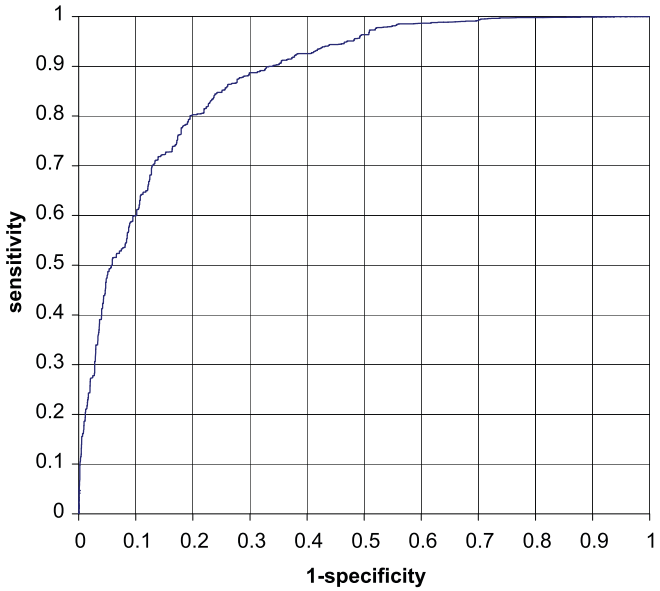

Figure 3. ROC curve of the proposed system for assigning centerline pixels to either artery or vein class. Area under the curve is 0.88.

the median as well as the average label for the entire vessel segment on the training data. These preliminary experiments showed that taking the average label resulted in slightly higher performance. The final output of the proposed method is thus a set of centerline pixels, each with a soft label.

4. EXPERIMENTS AND RESULTS

To evaluate the output of our method we used ROC analysis.[23] The ROC curve plot the sensitivity versus the 1-specificity. The entire curve can be compressed into a single number by taking the area under the curve. An area of 0.5 indicates the method is no better than random guessing and an area of 1.0 indicates the method is perfectly able to separate artery from vein centerline pixels. The final ROC of the method applied to the 20 images in the DRIVE test set is shown in Figure 3. The area under the ROC curve is 0.88. To visually asses the results we assigned a color varying between blue (i.e. vein) through purple to red (i.e. artery) to each of the vessel pixels for which a reference standard label was available. To determine the appropriate label the nearest centerline pixel for each of the vessel pixels was found. The label of that centerline pixel was then assigned to the vessel pixel. Some typical results are shown in Figure 4.

5. DISCUSSION AND CONCLUSION

The proposed artery/vein classification method shows good results in automatically distinguishing between arteries and veins as shown by the ROC analysis with an area under the ROC of 0.88. Even though the results in Figure 4 may not always appear to be correct due to the assigned colors it is important to realize that it is more important that the arteries in an image locally have a lower label value than the nearby veins than to have all vein pixel assigned a higher label than all artery pixels. To separate arteries from veins a localized approach is most likely beneficial as variation in the lighting and exposure over the image can have an influence on the locally assigned label values.

We hypothesized that the use of higher order derivatives could be useful in the separation of arteries and veins due to the presence of central reflexes. The feature selection process showed that this was indeed the case with several third and fourth order derivatives being selected. The feature selection also showed that unnormalized green-plane intensity features (20-24) were not useful which makes sense because of the considerable variation between the local intensity within an image as well as between images.

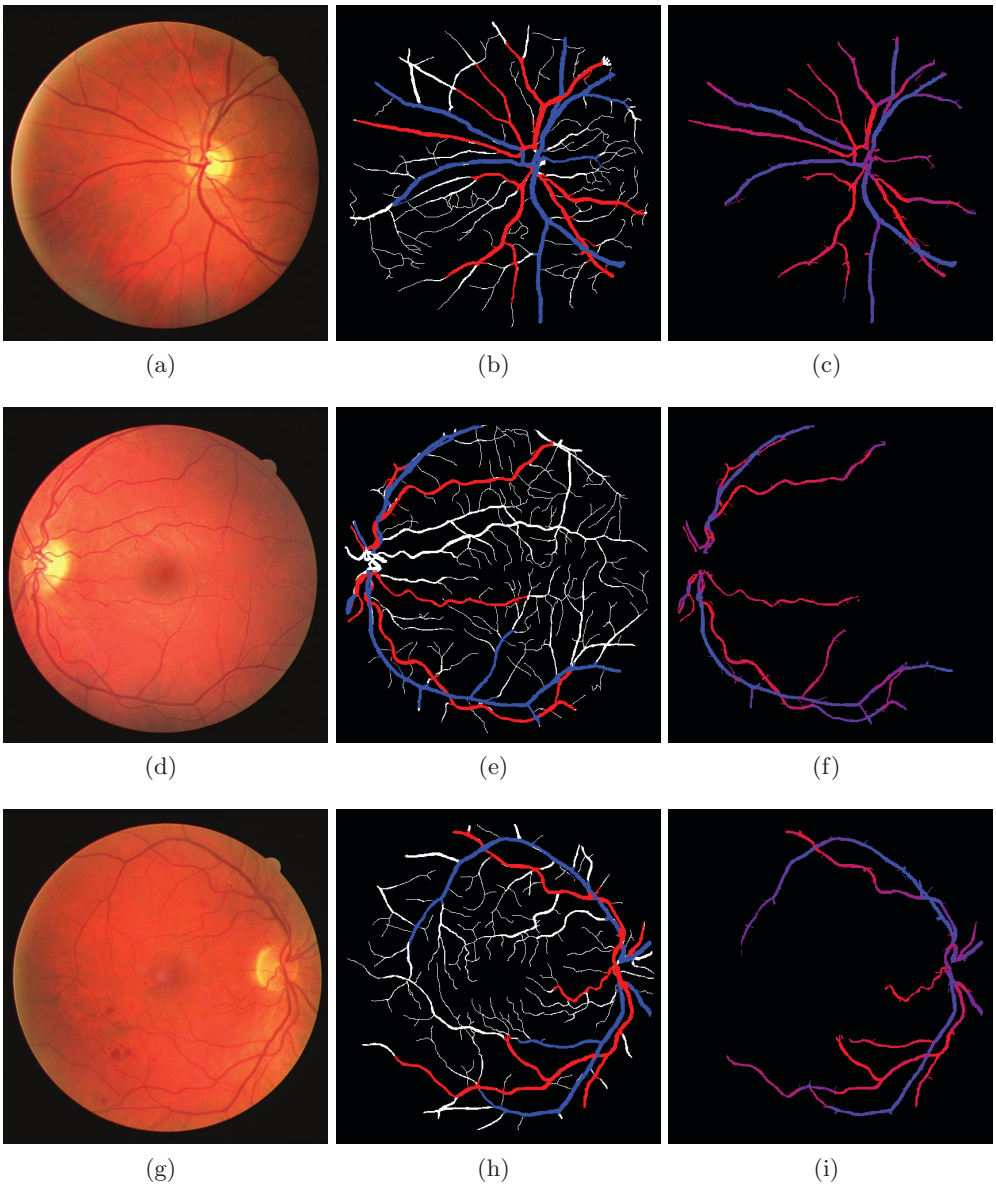

Figure 4. Some typical results with the color fundus photograph (a,d,g), the reference standard with the labeled veins in blue and the arteries in red (b,e,h) and the algorithm output (c,f,i). The algorithm only classifies centerline pixels so to generate these output images the centerline labels were propagated to the other vessel pixels for which a reference standard label was available. Based on the soft label a color was assigned to the vessel, bluer vessels have a higher probability of being veins and vice versa for redder vessels. The closer the vessel color is to purple the more ambiguous the algorithm output is (i.e. the soft label is close to 0.5).

The reference standard we used was set by a single expert observer, ideally one would like an additional observer to repeat the annotation to compare the two observers. Another issue with the reference standard is that it is likely the expert observer was given an unfair advantage by the fact that he may have unconsciously used the connectedness of the vessel network in setting the reference standard. A fairer way would be to devise an experiment where only small pieces of the vessels are shown to the observer who then has to indicate either artery or vein.

Future integration of vessel connectedness would likely result in better performance because assigned labels can be propagated from vessel segment to vessel segment. This would potentially also allow us to label small

vessels for which no label can be assigned now.

REFERENCES

1. J. Staal, M. D. 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.
2. 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.
3. 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.
4. 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.
5. 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.
6. 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 Transactions on Medical Imaging* **25**(9), pp. 1214–1222, 2006.
7. E. Ricci and R. Perfetti, "Retinal blood vessel segmentation using line operators and support vector classification.," *IEEE Transactions on Medical Imaging* **26**(10), pp. 1357–1365, 2007.
8. 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.
9. 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.
10. L. Wang, A. Bhalerao, and R. Wilson, "Analysis of retinal vasculature using a multiresolution hermite model.," *IEEE Trans Med Imaging* **26**, pp. 137–152, Feb 2007.
11. 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.
12. B. Y. Lam and H. Yan, "A novel vessel segmentation algorithm for pathological retina images based on the divergence of vector fields.," *IEEE Trans Med Imaging* **27**, pp. 237–246, Feb 2008.
13. M. Martínez-Pérez, A. Hughes, A. Stanton, S. Thom, N. Chapman, A. Bharath, and K. Parker, "Retinal vascular tree morphology: A semi-automatic quantification," *IEEE Transactions on Biomedical Engineering* **49**(8), pp. 912–917, 2002.
14. K. Rothaus, P. Rhiem, and X. Jiang, "Separation of the retinal vascular graph in arteries and veins," in *LNCS*, **4538**, p. 251262, 2007.
15. E. Grisan and A. Ruggeri, "A divide et impera strategy for automatic classification of retinal vessels into arteries and veins," in *Proceedings of the 25th Annual International Conference of the IEEE EMBS*, pp. 890–893, 2003.
16. A. Ruggeri, E. Grisan, and M. De Luca, "An automatic system for the estimation of generalized arteriolar narrowing in retinal images," in *Proc. 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society EMBS 2007*, pp. 6463–6466, 22–26 Aug. 2007.
17. L. Tramontan, E. Grisan, and A. Ruggeri, "An improved system for the automatic estimation of the arteriolar-to-venular diameter ratio (avr) in retinal images," in *Proc. 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society EMBS 2008*, pp. 3550–3553, 20–25 Aug. 2008.
18. M. Niemeijer, J. Staal, B. van Ginneken, M. Loog, and M. D. 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.
19. M. Ahmed and R. Ward, "A rotation invariant rule-based thinning algorithm for character recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**(12), pp. 1672–1678, 2002.

20. P. Rockett, "An improved rotation-invariant thinning algorithm," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **27**(10), pp. 1671–1674, 2005.
21. W. Freeman and E. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(9), pp. 891–906, 1991.
22. P. Pudil, J. Novovicova, and J. Kittler, "Floating search methods in feature selection," *Pattern Recognition Letters* **15**(11), pp. 1119–1125, 1994.
23. C. Metz, "ROC methodology in radiologic imaging," *Investigative Radiology* **21**(9), pp. 720–733, 1986.