Retinal Artery/Vein Classification
via Graph Cut Optimization

Koen Eppenhof[1], Erik Bekkers[1], Tos T.J.M. Berendschot[2],
Josien P.W. Pluim[1], and Bart M. ter Haar Romeny[1,3]

[1] Biomedical Engineering, Eindhoven University of Technology, The Netherlands
[2] University Eye Clinic Maastricht, The Netherlands
[3] Northeastern University, Shenyang, China
{k.a.j.eppenhof, e.j.bekkers, j.pluim, b.m.terhaarromeny}@tue.nl
t.berendschot@maastrichtuniversity.nl

Abstract. In many diseases with a cardiovascular component, the geometry of microvascular blood vessels changes. These changes are specific to arteries and veins, and can be studied in the microvasculature of the retina using retinal photography. To facilitate large-scale studies of artery/vein-specific changes in the retinal vasculature, automated classification of the vessels is required. Here we present a novel method for artery/vein classification based on local and contextual feature analysis of retinal vessels. For each vessel, local information in the form of a transverse intensity profile is extracted. Crossings and bifurcations of vessels provide contextual information. The local and contextual features are integrated into a non-submodular energy function, which is optimized exactly using graph cuts. The method was validated on a ground truth data set of 150 retinal fundus images, achieving an accuracy of 88.0% for all vessels and 94.0% for the six arteries and six veins with highest caliber in the image.

Keywords: artery/vein classification, retinal image analysis, graph cut optimization, contextual feature analysis

1 Introduction

The retinal vasculature can change function or geometry in a variety of ocular and systemic diseases, including glaucoma, macular degeneration, atherosclerosis, and hypertensive or diabetic retinopathy [1]. One of the earliest signs of these diseases is generalized arteriolar narrowing, in which the calibers of arteries decrease relatively to the calibers of veins. This phenomenon can be quantified using the ratio of the arteriolar and venular diameters, which is summarized in a value called the arteriovenous ratio (AVR). The AVR is commonly computed from the six arteries and veins with highest caliber in a zone around the optic disc [2]. Automated methods that measure the AVR directly from a retinal image require localization of the vessels and classification of vessels into arteries and veins. Additionally, artery/vein classification of retinal vessels is necessary for other artery/vein specific vessel features, such as vessel tortuosity [3].

Most existing automated artery/vein classification methods have focused on classification of vessels via local intensity-based features [4–6]. The difference in appearance of arteries and veins in white light retinal images is primarily determined by the oxygen-content of the blood, causing arteries to appear bright red, whereas veins look darker. An additional difference in intensity is sometimes present in the central light reflex on the vessels caused by the white flash of the camera. In arteries this reflection is often more pronounced, and can be used to distinguish them from veins [5, 6]. Recent methods have also focused on contextual information present in bifurcations and crossings of vessels. These methods make use of the fact that vessels that join in a bifurcation must be of the same type, and two crossing vessels must be of opposite type [7–9].

We propose a novel and generic graph-based method to combine local and contextual features of the retinal vasculature for artery/vein classification. In contrast to previous graph-based methods such as [7–9], here the classification is defined as an optimization problem, based on a non-submodular energy function that is minimized exactly and efficiently using the Quadratic Pseudo-Boolean Optimization (QPBO) graph cut algorithm. Optimization of this energy generates a classification of all detected vessels based on local features, while the contextual features limit the number of possible configurations of vessel labels. This results in a flexible approach in which the influence of contextual and local information can be weighted based on the their confidence levels.

2   Methods

Models of the retinal vasculature form the basis of our approach. In our work, we rely on models of the vessels around the optic disc, made using the Edge Tracking in Orientation Scores (ETOS) algorithm [10], but in principle any model that provides the edges of vessels and detects their crossings and bifurcations can be used, for example [11, 12, 9]. The models are used to define an energy function (Section 2.1) that uses features extracted from both the image and the model (Section 2.2). Optimization of the energy results in artery/vein labeling of the vessels in the model (Section 2.3).

2.1   Energy function definition

The energy function has the form of a binary first-order Markov Random Field (MRF)

$$U(\mathbf{y}) = \underbrace{\sum_{u \in V} \theta_u(y_u, \mathbf{x}_u)}_{\text{unary term}} + \lambda \underbrace{\sum_{(u,v) \in E} \theta_{uv}(y_u, y_v)}_{\text{pairwise term}}. \tag{1}$$

In this formulation, the first term is a sum of unary potential functions θ_u, taking as input a binary label $y_u \in \{0, 1\}$ and an evidence vector for this label $\mathbf{x}_u$. The second term consists of pairwise potentials θ_{uv}, each a function of a pair of labels. The binary MRF variables can be described as a set of vertices V in

an undirected graph $G = (V, E)$. The set of edges E in this graph contains the pairwise potentials. The λ parameter weights the pairwise term relative to the unary term.

In our artery/vein classification approach, each of the binary variables y_u corresponds to a vessel u that needs to be labeled as artery ($y_u = 1$) or vein ($y_u = 0$). Minimization of the energy in Equation (1) should result in an optimal labeling of the vessels

$$\mathbf{y}^* = \arg\min_{\mathbf{y}} U(\mathbf{y}). \tag{2}$$

The unary potentials optimize the labels based on local feature vectors $\mathbf{x}_u$ that have been extracted for each vessel. We define the unary potential as $\theta_u(y_u, \mathbf{x}_u) := -\ln p(y_u|\mathbf{x}_u)$, where $p(y_u|\mathbf{x}_u)$ is a posterior PDF which can be derived using Bayes' theorem

$$p(y_u|\mathbf{x}_u) = \frac{p(\mathbf{x}_u|y_u)p(y_u)}{p(\mathbf{x}_u|y_u = 0)p(y_u = 0) + p(\mathbf{x}_u|y_u = 1)p(y_u = 1)}. \tag{3}$$

where $p(\mathbf{x}_u|y_u)$ is the likelihood of a feature vector $\mathbf{x}_u$ for a given label y_u, and $p(y_u)$ is the prior PDF of the label y_u. We assume that the likelihood of the feature vectors $\mathbf{x}_u$ follows a multivariate Gaussian distribution $p(y_u|\mathbf{x}_u) = \mathcal{N}(\boldsymbol{\mu}_{y_u}|\boldsymbol{\Sigma}_{y_u})$ for both classes $y_u \in \{0, 1\}$. This means a training step is required to obtain the mean vectors $\boldsymbol{\mu}_{y_u}$, covariance matrices $\boldsymbol{\Sigma}_{y_u}$, and prior probabilities $p(y_u)$ for $y_u \in \{0, 1\}$.

The pairwise potentials penalize crossings of vessels with the same label, or bifurcations of vessels with different labels. We define $\theta_{uv}(y_u, y_v) = \gamma_{uv}(1 - \delta(y_u, y_v))$, where $\gamma_{uv} := -1$ if vessels u and v cross, $\gamma_{uv} := 1$ if vessel u bifurcates from v or vice versa, and $\delta(y_u, y_v) := 1$ if $y_u = y_v$ and $\delta(y_u, y_v) := 0$ otherwise. This results in a utility of $-\lambda$ for correct crossings and a penalty of λ for incorrect bifurcations. Note that if one is able to determine a confidence measure for each of the bifurcations and crossings, the method can be extended by using individual weighting factors γ_{uv} that are a function of the confidence of a bifurcation's or crossing's detection.

We arrive at the following problem-specific cost function

$$U(\mathbf{y}) = \underbrace{\sum_{u \in V} -\ln p(y_u|\mathbf{x}_u)}_{\text{local features}} + \lambda \underbrace{\sum_{(u,v) \in E} \gamma_{uv}(1 - \delta(y_u, y_v))}_{\text{contextual features}}. \tag{4}$$

2.2   Feature extraction and training

To find the unary and pairwise potential functions, two types of features are required. To determine the posterior PDF in the unary potentials, a feature vector $\mathbf{x}_u$, representative of the differences between arteries and veins, is needed. To determine the pairwise potentials, the method requires contextual information from the tracking models.
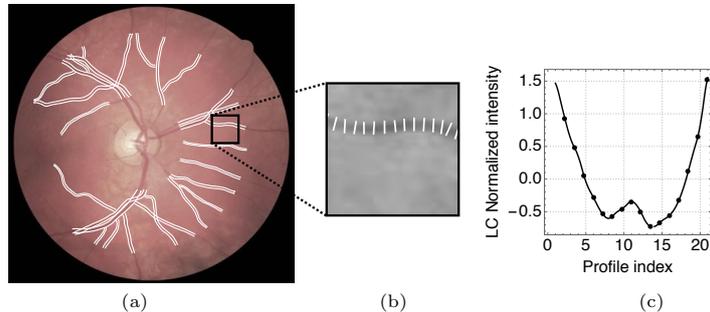
Fig. 1. (a) Example of a retinal vasculature tracking model. (b) Detail of a tracked vessel and its profile measurement locations. (c) Average vessel intensity profile.

Local features The differences in color and central light reflex of arteries and veins can be captured in transverse intensity profiles [4, 5, 13]. Like in previous work [9], we found that these differences are largest in the red color channel.

For each pair of edge points on the tracked vessels, the intensity values on 21 points on a line between the edge points were extracted via cubic Hermite spline interpolation of the red channel of the image. The number of points is constant, independent of vessel width or resolution. This resulted in a number of intensity profiles for each vessel the size of which depends on the length of the tracked vessel. Prior to this, the red channel had been normalized using a lightness and contrast (LC) normalization described in [14]. Additionally, each of the intensity profiles was normalized to zero mean and unit standard deviation. The intensity profiles of each vessel were averaged pointwise, resulting in one profile $\mathbf{x}_u$ per vessel (Figure 1). The intensity profiles of a training set were used to determine $p(y_u|\mathbf{x}_u)$ by calculating the mean vector $\boldsymbol{\mu}_{y_u}$, covariance matrix $\boldsymbol{\Sigma}_{y_u}$, and prior PDF $p(y_u)$ for arteries and veins separately.

Contextual features The contextual features γ_{uv} can be directly extracted from the tracking models made with the ETOS algorithm. The actual crossings were detected by determining intersections of the tracked center lines of the vessels. Bifurcations are detected during tracking itself [10].

2.3   Energy minimization

The energy function $U(\mathbf{y})$ in (4) is a supermodular function of binary variables. This supermodularity is caused by the fact that $\theta_{uv}(0,0) + \theta_{uv}(1,1) \leq \theta_{uv}(1,0) + \theta_{uv}(0,1)$ does not hold in case of crossings. This means that the energy function cannot be optimized using a standard graph cut algorithm. Therefore, we use the Quadratic Pseudo-Boolean Optimization algorithm (QPBO) [15]. This algorithm transforms supermodular problems into submodular problems that have a partial solution. For every vertex u in the original problem, the QPBO problem has two
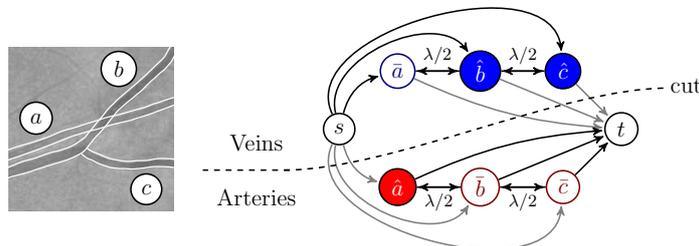
Fig. 2. Example of a set of vessels and resulting QPBO graph. Vessels a and b cross and need to be labeled oppositely, whereas vessels b and c belong to the same branch and thus must receive the same label. Hence, in the QPBO graph, $\hat{a}$ and $\bar{a}$ are connected to $\bar{b}$ and $\hat{b}$, while $\hat{b}$ and $\bar{b}$ are connected to $\hat{c}$ and $\bar{c}$ respectively. These edges all have weight $\lambda/2$. The edge weights of the edges to s and t correspond to the unary terms in equation (4). The cut shown is the result if the sum of weights of the black edges is higher than the sum of the weights of the gray edges, and results in a labeling where a is an artery and b and c are veins.

vertices $\hat{u}$ and $\bar{u}$. Each of these vertices is connected to a source s and a sink t in a flow graph. The $\hat{u}$ vertices connect to s by an edge with weight $\theta_u(1, \mathbf{x}_u)$ and to t by an edge with weight $\theta_u(0, \mathbf{x}_u)$. These edge weights correspond to the unary terms in equation (4). The $\bar{u}$ vertices connect to s by an edge with weight $\theta_u(0, \mathbf{x}_u)$ and to t by an edge with weight $\theta_u(1, \mathbf{x}_u)$. The pairwise terms are modeled as edges between vertices. For two crossing vessels u and v the weights are $\theta_{\hat{u}\bar{v}} = \lambda/2$ and $\theta_{\hat{v}\bar{u}} = \lambda/2$, for bifurcating vessels u and v the weights are $\theta_{\bar{u}\bar{v}} = \lambda/2$ and $\theta_{\hat{u}\hat{v}} = \lambda/2$. All other pairwise weights are 0.

To find the minimum cut in the QPBO-version of the flow graph, we used the Ford-Fulkerson algorithm [16]. After computing a minimum $s-t$ cut (S, T), the final labels of the vertices are determined as follows: if $\hat{u}$ and $\bar{u}$ are in opposite cut sets, the original vertex u receives the label of $\hat{u}$. If $\hat{u}$ and $\bar{u}$ are in the same cut set, an 'unknown' label $y_u = \emptyset$ is attached to u. See Figure 2 for an illustration.

3   Experiments

3.1   Ground truth data

We validated the method on a ground truth data set of 150 images. All images were taken with a digital non-mydriatic white light fundus camera of the Nidek brand, model AFC-230, with a 45° field of view and a resolution of 3744×3744. For each of the images, a model of the vasculature was generated using the ETOS algorithm [10], and corrected manually when necessary[4]. The vessels were labeled as arteries and veins through consensus of two graders. In total, the dataset contains 3186 vessels, divided in 1693 arteries and 1493 veins.

---

[4] Data on the performance of the ETOS algorithm can be found in [10].
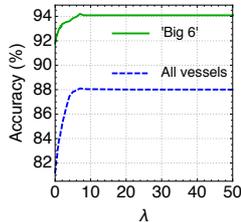
Fig. 3. Classification accuracy against the value of λ for all vessels, and the six biggest arteries and veins found after classification for all images ('Big-6'). For $\lambda > 8$ accuracy remains constant.
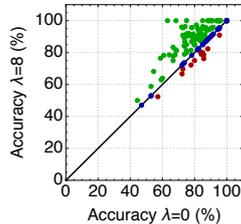
Fig. 4. Accuracies for each image for $\lambda = 8$ against $\lambda = 0$. A majority of 105 images (green) show an increase in accuracy, with 12 images showing a decrease (red) and 33 images showing no change (blue).
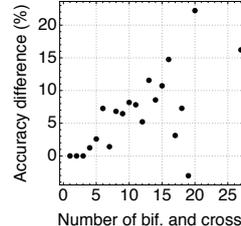
Fig. 5. Accuracy difference between using $\lambda = 8$ and $\lambda = 0$ against total number of bifurcations and crossings. Each point shows the average for images with that number of bifurcations and crossings.

3.2   Method parameterization and validation

The energy function in Equation (4) contains one parameter λ which weights the influence of unary and pairwise terms, and therefore of local and contextual features. To determine the dependence of our method on this parameter, the accuracies for classification for increasing values of λ was evaluated.

For the classification of the vessels in each of the images we computed a confusion matrix for each value of λ. For each of the images, the other 149 images were used for training of the unary potential functions in a leave-one-out cross validation. By summing the confusion matrices for all images, we achieved a confusion matrix of the whole data set, and thus accuracy and sensitivity, for each of the λ values.

By setting λ to zero, no contextual information is included at all. This setting is used as a baseline, achieving a classification accuracy of 81.2% for all vessels, and 91.8% for the six arteries and veins with highest caliber (the 'Big-6' arteries and veins, see Table 1). Accuracy increased with increasing λ, to 94.0% for the 'Big-6' arteries and veins, and to 88.0% for all vessels. This compares well to accuracy levels acquired with methods in literature that were validated on different data sets (87.6% [4], 85.5% [5], 88.3% [9], 88.8% [13]).

We evaluated the accuracy for values of λ up to 50, which is sufficient to let the pairwise terms dominate the unary terms completely. At $\lambda = 8$, accuracy stopped increasing, which is why this value is used in subsequent experiments (Figure 3). We confirmed the significance of the improvement of using $\lambda = 8$ compared to using $\lambda = 0$ using McNemar's test ($\chi^2 = 150, p < 0.001$).

In 105 out of 150 images the accuracy was higher for $\lambda = 8$ than for $\lambda = 0$ (Figure 4). In twelve images the accuracy did not improve. This is caused by incorrect high posterior probabilities of a vessel that bifurcates from, or has a crossing with, another vessel. As an example, consider the case of a crossing of an artery v_A and vein v_V for which the posteriors $p(\mathbf{x}_u|\text{Artery})$ are 0.6 and 0.7

Table 1. Comparison of accuracy and sensitivity for arteries and veins for $\lambda = 0$ and $\lambda = 8$ (maximal contextual information), for vessels of different calibers, all vessels, and the six biggest arteries and veins found in the classification. N is the number of vessels in the category. For the 'Big 6' category the number N reflects how many of the six required arteries and veins were found in the classification.

| | | $\lambda = 0$ | | | | $\lambda = 8$ | | |
| | | | Sensitivity | | | | Sensitivity | |
Caliber	N	Accuracy	Arteries	Veins	N	Accuracy	Arteries	Veins
0 - 20 pixels	960	65.7%	64.4%	74.5%	960	78.2%	75.3%	86.9%
20 - 25 pixels	985	83.2%	82.2%	86.1%	985	89.1%	88.6%	90.4%
25 - 30 pixels	664	90.1%	89.3%	90.9%	664	94.1%	93.5%	94.9%
$\geq$ 30 pixels	577	93.4%	85.0%	96.6%	577	95.3%	88.7%	97.8%
All vessels	3186	81.2%	76.4%	90.1%	3186	88.0%	84.5%	93.2%
'Big 6'	1704	91.8%	88.8%	95.1%	1797	94.0%	92.4%	95.8%

respectively. Based on the posterior probabilities, both would be classified as arteries, but by the fact that they cross, they are classified as a vein and artery respectively, since the penalties for classifying v_V as a vein or classifying both as arteries is higher than classifying v_A as a vein.

Other errors are caused by incorrect local classifications. Table 1 shows that overall veins are classified more accurately than arteries, which was seen before by other authors [5, 9], and is likely caused by the fact that the smallest arteries often look like veins and consequently are classified incorrectly.

To investigate which vessels benefit most from the contextual information, we separated the vessels into groups based on caliber. We found that the accuracy improves most for the smallest vessels (0-25 pixels in caliber) and less for the larger vessels (above 25 pixels, see Table 1). Furthermore, we clearly see a positive correlation of the accuracy with the total number of crossings and bifurcations, which demonstrates the value of contextual features for classification (Figure 5).

4   Conclusion

In this work we propose a graph-based method to combine local and contextual features in artery/vein classification of vessels in retinal images, which we validated on a clinical dataset of 150 retinal images. As local features of the vessels we use intensity profiles, and as contextual features we include restrictions imposed by bifurcations and crossings. The inclusion of contextual features significantly improves the classification accuracy compared with using local features alone. The classification accuracy of our method of 88.0% compares well with previous methods. For the six arteries and veins with highest caliber we obtain an accuracy of 94.0%. The classification of vessels with a lower caliber benefits most from the contextual information. Because local information is less reliable in these vessels, contextual information is essential.

References

1. Ikram, M.K., Ong, Y.T., Cheung, C.Y., Wong, T.Y.: Retinal vascular caliber measurements: clinical significance, current knowledge and future perspectives. Ophtalmologica **229**(3) (2013) 125–36
2. Knudtson, M.D., Lee, K.E., Hubbard, L.D., Wong, T.Y., Klein, R., Klein, B.E.K.: Revised formulas for summarizing retinal vessel diameters. Current eye research **27**(3) (2003) 143–9
3. Kalitzeos, A.A., Lip, G.Y., Heitmar, R.: Retinal vessel tortuosity measures and their applications. Experimental Eye Research **106** (2013) 40 – 46
4. Grisan, E., Ruggeri, A.: A divide et impera strategy for automatic classification of retinal vessels into arteries and veins. In: IEEE EMBS. (2003) 890–893
5. Li, H., Hsu, W., Lee, M., Wang, H.: A piecewise Gaussian model for profiling and differentiating retinal vessels. In: IEEE ICIP. (2003) I–1069–72
6. Niemeijer, M., Xu, X., Dumitrescu, A.V., Gupta, P., van Ginneken, B., Folk, J.C., Abramoff, M.D.: Automated measurement of the arteriolar-to-venular width ratio in digital color fundus photographs. IEEE TMI **30**(11) (2011) 1941–50
7. Rothaus, K., Jiang, X., Rhiem, P.: Separation of the retinal vascular graph in arteries and veins based upon structural knowledge. Image and Vision Computing **27**(7) (2009) 864–875
8. Joshi, V.S., Reinhardt, J.M., Garvin, M.K., Abramoff, M.D.: Automated method for identification and artery-venous classification of vessel trees in retinal vessel networks. PLoS ONE **9**(2) (2014) e88061
9. Dashtbozorg, B., Mendonça, A.M., Campilho, A.: An automatic graph-based approach for artery/vein classification in retinal images. IEEE TIP **23**(3) (2014) 1073–83
10. Bekkers, E., Duits, R., Berendschot, T., ter Haar Romeny, B.: A Multi-Orientation Analysis Approach to Retinal Vessel Tracking. Journal of Mathematical Imaging and Vision **49**(3) (2014) 583–610
11. Perez, M., Highes, A., Stanton, A., Thorn, S., Chapman, N., Bharath, A., Parker, K.: Retinal vascular tree morphology: a semi-automatic quantification. IEEE TBME **49**(8) (2002) 912–917
12. Al-Diri, B., Hunter, A., Steel, D., Habib, M.: Automated analysis of retinal vascular network connectivity. Computerized Medical Imaging and Graphics **34**(6) (2010) 462 – 470
13. Vazquez, S.G., Cancela, B., Barreira, N., Penedo, M.G., Saez, M.: On the Automatic Computation of the Arterio-Venous Ratio in Retinal Images: Using Minimal Paths for the Artery/Vein Classification. In: IEEE DICTA. (2010) 599–604
14. Foracchia, M., Grisan, E., Ruggeri, A.: Luminosity and contrast normalization in retinal images. Medical Image Analysis **9**(3) (2005) 179–90
15. Kolmogorov, V., Rother, C.: Minimizing nonsubmodular functions with graph cuts - a review. IEEE PAMI **29**(7) (July 2007) 1274–9
16. Ford, L.R., Fulkerson, D.R.: Maximal flow through a network. Canadian Journal of Mathematics **8** (1956) 399–404