

A Bayesian Framework For the Local Configuration of Retinal Junctions

Touseef Ahmad Qureshi, Andrew Hunter and Bashir Al-Diri

School of Computer Science, University of Lincoln, United Kingdom

tqureshi@lincoln.ac.uk, ahunter@lincoln.ac.uk, baldiri@lincoln.ac.uk

Abstract

Retinal images contain forests of mutually intersecting and overlapping venous and arterial vascular trees. The geometry of these trees shows adaptation to vascular diseases including diabetes, stroke and hypertension. Segmentation of the retinal vascular network is complicated by inconsistent vessel contrast, fuzzy edges, variable image quality, media opacities, complex intersections and overlaps. This paper presents a Bayesian approach to resolving the configuration of vascular junctions to correctly construct the vascular trees. A probabilistic model of vascular joints (terminals, bridges and bifurcations) and their configuration in junctions is built, and Maximum A Posteriori (MAP) estimation used to select most likely configurations. The model is built using a reference set of 3010 joints extracted from the DRIVE public domain vascular segmentation dataset, and evaluated on 3435 joints from the DRIVE test set, demonstrating an accuracy of 95.2%.

Keywords: Retinal vessels configuration, vessels connectivity, junction resolution, vessels trees reconstruction.

1. Introduction

Retinal vascular trees demonstrate adaptation in a variety of systemic diseases, including diabetic retinopathy, arteriosclerosis, and hypertension [13], and are useful in biometric applications [8]. Retinal vascular segmentation is therefore an active research topic. Prior to performing quantitative analysis, an essential step is the correct formation of retinal vessel trees.

There is relatively little published work for retinal vasculature tree formation (e.g. [2] and [14]). Segmenting human vascular trees in general fall into two classes: one segments the vascular region producing a pixel map, followed by vessel thinning to form a vessel network [11, 15]; the other tracks vessels by perceptual grouping of segments at disconnected regions to form trees. The former suffers from a range of problems, including: inconsistent vessel




Figure 1. A complex junction featuring two overlapping bifurcations.

contrast, fuzzy edges, variable image quality, media opacities, complex intersections and overlaps. The latter may use a range of graph exploratory algorithms for segments such as minimum spanning trees [4], Mahalanobis distance [7] and contour completion using fast marching algorithm [5]. These algorithms are not typically optimized to exploit the specific topological features of the retina.

Al-Diri *et al.* [2] used perceptual grouping of retinal segments at disconnected junctions using self-organizing feature map (SOFMs) that uses the cost functions for configuration. Tsai *et al.* [14] have proposed a vessel tracing algorithm that can resolve the connectivity at simpler cases such as bifurcations and crossings. Neither methodology provides a proper mechanism for correctly identifying multiple joints at complex junctions such as two overlapping bifurcations; see Figure 1. Moreover, segmentation failure may lead to broken segments, including missing junctions, and to spurious segments.

This paper introduces a Bayesian algorithm designed to support correct identification of retinal vascular trees. It does this by analyzing junctions, where vessels meet. Each junction may contain a number of joints (e.g. bifurcations); the algorithm chooses the most likely configuration based

on a Maximum a Posteriori (MAP) estimation, with the priors and likelihood densities estimated from a reference data set. The rest of the paper is organized as follows. Section 2 discusses the reference database and methodology; section 3 introduces the Bayesian model; section 4 discusses the feature extraction methods; section 5 analyzes results. Section 6 concludes the paper.

2. Methodology and Database

The retinal vasculature consists of a number of vascular trees, each consisting of vessel segments meeting at bifurcations. These trees may occlude one another, producing visual crossing points; two closely parallel occluding vessels may visually overlap, producing an apparent “shared” segment between two apparent bifurcations. Failures of segmentation may also lead to apparent gaps in a vessel (two visual segments end to end), and this may also occur, depending on the segmentation algorithm, at bifurcations or crossings. Segmentation algorithms therefore retrieve a (possibly partially discontinuous) vascular network, from which we must derive the underlying forest of trees. To solve this problem we first divide the segmented network into a set of simple segments divided by junctions, and then join the segments into trees using relationships called *joints*.

Definition 1: *segment*. A segment, s_i , is an atomic section of apparent vessel with no bifurcations, crossings or breaks; it has two segment ends, the head s_i^h (nearest the Optic Nerve Head) and the tail, s_i^t .

Definition 2: *junction*. A junction, J_k , is a clique defined within a small locality in the image where a number of segment ends either meet or come into close proximity. This includes the degenerate case of a single isolated segment end.

Definition 3: *joint*. A joint, L_k^p (the p^{th} joint in the k^{th} junction) is a collection of segment ends. There are three types of joint: a terminal contains a single segment end $\{s_i\}$, a bridge is a set of two segment ends $\{s_i, s_j\}$, a bifurcation contains one parent segment end, and a set of two child segment ends, $(\{s_i\}, \{s_j, s_k\})$.

The algorithm assumes that the network is first divided into independent segments, then assigns every segment end to a junction. The vascular trees are retrieved by connecting sets of segments via bridges and bifurcations. In this approach, bridges are used both to rejoin vessels with segmentation gaps, and to resolve intersections (which appear as four segments in a junction, resolving into two bridges).

The DRIVE (Digital Retinal Images for Vessel Extraction) public domain retinal vascular dataset [12] is used to optimize and test the algorithm. DRIVE contains 40 retinal images, resolution 565×584 , divided into 20 training and 20 test images, together with reference standard vessel segmentation maps (pixel maps). We developed an interactive computer program to further process DRIVE to pro-




Figure 2. An example of joints removed at a junction.

duce the “Drive Segment-Junction Set (DSJS),” available at <http://reviewdb.lincoln.ac.uk>. Using this program, we first removed all the junction points, including bifurcations and crossings, identified in [3]. We eliminated a circular area (radius $r = 5$ pixels) with the junction point as the center, so that all segments are isolated; see Figure 2. In 2-3% of cases – in vessel-congested areas – connections were removed manually using a drawing tool rather than a circle, avoiding excessive loss of detail. The optic nerve head area, which is very congested with vessels, was eliminated from DSJS using [9]. The segments were then identified as the morphologically thinned connected components of DSJS (representing centrelines). Junctions were created corresponding to each junction point. Segments ends within five pixels of a junction point were then assigned to that junction, and any junctions with junction points within five pixels of each other were merged. We then interactively identified the joints within each junction; in the case of bifurcation, this process includes explicitly labelling the parent segment.

Definition 4: *Segment Space*. The segment space S is the set of $2N$ vessel segment ends present in a particular DSJS image, with their segment index and end type. S is given by $S = \{s_i^e | i = 1, 2, \dots, N \wedge e \in \{h, t\}\}$, where h denotes the head and t the tail.

The following rules apply:

1. A junction J_k is a proper subset of segment space, $J_k \subset S$
2. The union of all junctions equals to segment space, i.e. $\bigcup_k J_k = S$
3. A segment s_i cannot have both its ends belonging to the same junction, i.e. $\forall i, j, s_i^h, s_j^t \in J_k \implies i \neq j$

To satisfy rule 2, any segment ends not assigned to a junction derived from DSJS must be assigned a singleton junction containing just that segment end. To satisfy rule 3, any segment abrogating the rule may be eliminated.

Definition 5: *Junction configuration*. A junction configuration, C_k^q (the q^{th} configuration of the k^{th} junction) is a set of joints, such that every segment end in the junction is assigned to exactly one joint, $C_k^q = \{L_k^p | \bigcup_p s_i \in L_k^p = J_k\}$.

Definition 6: *Feature vector*. A feature vector, \mathbf{f}_i^e , is a set of measurements extracted at a segment end and characterizing its direction and appearance.

3. Bayesian Inference

The Bayesian inference model works by enumerating all possible configurations of a junction, estimating the posterior probabilities, $P(C_k^q|J_k)$ up to a normalizing constant, and choosing the MAP estimate (highest posterior configuration). The configuration posterior is estimated by Bayes theorem using the prior times the likelihood:

$$P(C_k^q|J_k) \propto P(C_k^q)P(J_k|C_k^q) \quad (1)$$

We assume that the joints composing the configuration are statistically independent of each other, and we can therefore rewrite Equation (1) as:

$$P(C_k^q|J_k) = P(C_k^q) \prod_p P(L_k^p|\mathbf{f}_k^p) \quad (2)$$

where $P(L_k^p|\mathbf{f}_k^p)$ is the posterior probability of the joint, given the observed features $\mathbf{f}_k^p = \{\cup \mathbf{f}_i^e | s_i^e \in L_k^p\}$

In a second application of Bayes, we estimate the $P(L_k^p|\mathbf{f}_k^p)$ as

$$P(L_k^p|\mathbf{f}_k^p) = \frac{P(\mathbf{f}_k^p|L_k^p)P(L_k^p)}{P(\mathbf{f}_k^p|L_k^p)P(L_k^p) + P(\mathbf{f}_k^p|L_k^{p'})P(L_k^{p'})} \quad (3)$$

where $P(L_k^p)$ is the prior probability of the joint class, $P(\mathbf{f}_k^p|L_k^p)$ is the likelihood function for the observed features, given the (true) joint type, and $P(\mathbf{f}_k^p|L_k^{p'})$ is the likelihood function for the observed features given a “false joint” consisting of arbitrarily chosen segment ends within a junction that do not form a true joint.

The configurations of a junction with a given number of segment ends fall into several configuration classes, based on the number and type of joints in the configuration. For example, a junction J containing four segment ends ($s_1^e, s_2^e, s_3^e, s_4^e$) has four configuration classes containing 22 configurations: one with four terminal joints $\{\{s_1^e\}, \{s_2^e\}, \{s_3^e\}, \{s_4^e\}\}$, three containing two bridges $\{\{s_1^e, s_2^e\}, \{s_3^e, s_4^e\}\}, \{\{s_1^e, s_3^e\}, \{s_2^e, s_4^e\}\},$ and $\{\{s_1^e, s_4^e\}, \{s_2^e, s_3^e\}\}$, twelve with one bifurcation and one terminal (e.g. $\{\{s_1^e, \{s_2^e, s_3^e\}\}, \{s_4^e\}\}$), and six with a bridge and two terminals (e.g. $\{\{s_1^e, s_2^e\}, \{s_3^e\}, \{s_4^e\}\}$). Table 1 lists all the junction configuration classes with up to six segment ends, the frequencies within the training set, and the prior configuration probabilities, $P(C_k^q)$, estimated from these.

Some theoretically possible larger configurations never occurred; we set these to a minimal figure (0.001) to allow the system to respond if it does encounter such rare cases.

The likelihood function for a joint $P(\mathbf{f}_k^p|L_k^p)$, is modelled using the multivariate normal probability density function $\mathcal{N}(\mu_k^p, \Sigma_k^p)$, where μ_k^p and Σ_k^p are the centroid vector and

J	Configuration	Total	Frequency	$P(C_k)$
1	1T	1062	1017	0.951
2	2T	98	45	0.451
	1B		53	0.540
3	3T	1370	15	0.011
	1B, 1T		18	0.013
	1Y		1337	0.975
4	4T	711	5	0.007
	2B		632	0.888
	1B, 2T		19	0.026
	1Y, 1T		55	0.077
5	5T	51	0	0.001
	2B, 1T		2	0.039
	1B, 3T		0	0.001
	1Y, 1B		47	0.921
	1Y, 2T		2	0.039
6	6T	10	0	0.001
	3B		0	0.001
	2B, 2T		1	0.100
	1B, 4T		0	0.001
	2Y		7	0.700
	1Y, 3T		0	0.001
	1Y, 1B, 1T		2	0.200

Table 1. Prior probability estimates of all possible junction configurations. T = terminal, B = bridge, Y = bifurcation.

Number of segments at all junctions	4056
Number of terminals	1017
Prior Probability for terminals	0.250
Prior Probability for non-terminals	0.749
Groups of two segments in all junctions	2038
Number of bridges	670
Prior Probability for bridges	0.328
Prior Probability for non-bridges	0.672
Groups of three segments in all junctions	2081
Number of bifurcations	1381
Prior Probability for bifurcations	0.663
Prior Probability for non-bifurcations	0.336

Table 2. Frequency distributions and Prior probability estimates of joint classes.

covariance matrix of the joint feature vector for that class of joints; $P(\mathbf{f}_k^p|L_k^{p'})$ is similarly modelled.

The prior probabilities of the joint classes, $P(L_k^p)$, estimated from the frequency distribution in the training set, are given in Table 2.

4. Feature Extraction

To determine if a joint is valid or not a small number of features are extracted at the associated segment ends. These features bear rich information about joints. For example, terminals are small (low width), bridged segments have roughly equal widths, intensity profiles and directions, bifurcations have a consistent relationship between widths and angles. The features are extracted from an area along 7 centerline-profiles at the segment end; see Figure 3. The basic segment end features are:

- Segment direction: the direction vector \mathbf{m} represents the direction of the straight line between the first and

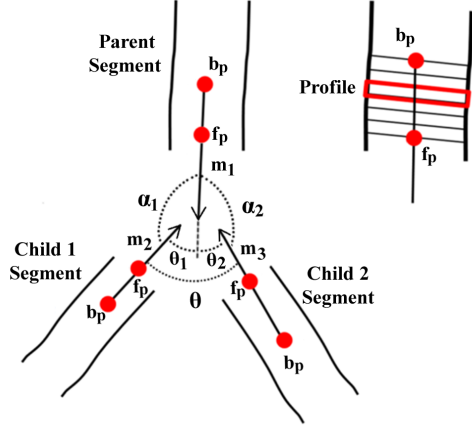


Figure 3. The geometry of a typical bifurcation, profiles at segment end.

last centreline points in the segment end f_p and b_p , towards the junction the segment end is associated with: see Figure 3.

- Segment width, w : we use the Ribbon of Twins active contour model [1] to extract the vessel width at the seven centreline pixels of the segment end, and take the mean.
- Segment intensity, μ_v, σ_v : we construct seven 1D intensity profiles orthogonal to \mathbf{m} at spacing one pixel, use Gregson's algorithm [6] to find the edges, and calculate the mean and standard deviation of the included intensities.
- Segment end class, e : as previously described, the end of a segment nearest the ONH is the head, the other end the tail.

During the width extraction process, the system sometimes generates abnormal widths, most frequently when two closely parallel segments are erroneously segmented as a single wider segment. This leads to a non-normal distribution of widths. We used the square root of the width to mitigate this outlier effect, which gives an approximately normal distribution.

4.1. Derived Features

Let s_1^e and s_2^e be two segment ends with basic features $\{w_1, \mu_{v_1}, \sigma_{v_1}, \mathbf{m}_1\}$ and $\{w_2, \mu_{v_2}, \sigma_{v_2}, \mathbf{m}_2\}$ respectively, then the deflection angle α and width ratio β are defined as:

$$\alpha = \cos^{-1} \left(\frac{\mathbf{m}_1 \cdot \mathbf{m}_2}{|\mathbf{m}_1| \cdot |\mathbf{m}_2|} \right) \quad (4)$$

$$\beta = \frac{\min(w_1, w_2)}{\max(w_1, w_2)} \quad (5)$$

We use the Bhattacharyya distance to characterize the distance between segment intensities:

$$\gamma = \frac{1}{4} \ln \left(\frac{1}{4} \left(\frac{\sigma_{v_1}^2}{\sigma_{v_2}^2} + \frac{\sigma_{v_2}^2}{\sigma_{v_1}^2} + 2 \right) \right) + \frac{1}{4} \left(\frac{(\mu_{v_1} - \mu_{v_2})^2}{\sigma_{v_1}^2 + \sigma_{v_2}^2} \right) \quad (6)$$

The parity between the end types in a pair, ϕ^2 , is defined as $\phi^2 = e_1 \oplus e_2$, where \oplus denotes the exclusive-OR operator.

The extracted features are then normalized to the range $[0, 1]$, using:

$$f_i^n = \frac{f_i - f_{min}}{f_{max} - f_{min}} \quad (7)$$

Where f_i^n denotes the standardized version of feature f_i . The minimum and maximum values are: $v_{min} = 0, v_{max} = 255, w_{min} = 0, w_{max} = 10$ (an empirical value, as the maximum width depends on the image resolution in DRIVE the maximum vessel width is 7 pixels), $\alpha_{min} = 0, \alpha_{max} = 360$.

The Shapiro-Wilk test was used to test all the features against normality; all features were found to be approximately normally distributed ($p=0.05$).

The three joint classes have distinct feature sets. For terminals, the feature set contains the width, mean intensity and end class: $\mathbf{f} = \{w_i, \mu_{v_i}, e_i\}$. The distributions of w and v were estimated from 1017 terminals in the twenty training images: $\mu_w = 0.09, \mu_v = 0.48, \sigma_w = 0.04, \sigma_v = 0.12$. The end class e is a discrete, independent variable, for which the prior probabilities are calculated using probability mass function of the Bernoulli distribution. Let, \mathcal{G}_{ter} be the probability mass function of the distribution e :

$$\mathcal{G}_{ter} = \begin{cases} 0.73 & e=t \\ 0.27 & e=h \end{cases} \quad (8)$$

where $0.73 = 742/1017$, there being 742 tail segment ends among 1017 terminals. The likelihood can then be estimated using the generalized probability density function as:

$$P(\mathbf{f}|L) = \mathcal{N}(\mu_w, \sigma_w) \mathcal{N}(\mu_v, \sigma_v) \mathcal{G}_{ter}. \quad (9)$$

For bridges, the feature set contains the deflection angle, width ratio, Battacharya distance between intensity distributions, and end type parity, $\mathbf{f} = \{\alpha, \beta, \gamma, \phi^2\}$. The distribution parameters were estimated from 564 bridges, as: $\mu_\alpha = 0.49, \sigma_\alpha = 0.056, \mu_\beta = 0.85, \sigma_\beta = 0.11, \mu_\gamma = 0.27, \sigma_\gamma = 0.15$. The end class parity has value 1 (tail meets head) in $531/564 = 0.94$ of the true bridges (the remaining 6% of cases occurring rarely where segments curve

	angle	β	γ
s_1, s_2	$\mu_{\alpha_1} = 0.36$ $\sigma_{\alpha_1} = 0.07$	$\mu_{\beta_1} = 0.71$ $\sigma_{\beta_1} = 0.32$	$\mu_{\gamma_1} = 0.31$ $\sigma_{\gamma_1} = 0.25$
s_1, s_3	$\mu_{\alpha_2} = 0.41$ $\sigma_{\alpha_2} = 0.08$	$\mu_{\beta_2} = 0.68$ $\sigma_{\beta_2} = 0.27$	$\mu_{\gamma_2} = 0.32$ $\sigma_{\gamma_2} = 0.31$
s_2, s_3	$\mu_{\theta} = 0.23$ $\sigma_{\theta} = 0.10$	$\mu_{\beta_3} = 0.72$ $\sigma_{\beta_3} = 0.24$	$\mu_{\gamma_2} = 0.27$ $\sigma_{\gamma_2} = 0.36$

Table 3. Parameters of distributions of bifurcation features.

back towards the ONH), and defines the probability mass function \mathcal{G}_{bri} . The likelihood is then given by:

$$P(\mathbf{f}|L) = \mathcal{N}(\mu_{\alpha}, \sigma_{\alpha})\mathcal{N}(\mu_{\beta}, \sigma_{\beta})\mathcal{N}(\mu_{\gamma}, \sigma_{\gamma})\mathcal{G}_{bri}. \quad (10)$$

For bifurcations, we do a pairwise analysis of the width, intensity and direction disparities between the parent segment end and each of the child segment ends in turn. We further include the angle θ between the two children, and an end class feature, ϕ^3 , which is defined as 1 when the parent is a tail and the children both heads, 0 otherwise. Let $f_{i,j}$ represent feature f taken between segments i and j , where 1 represents the parent and 2, 3 represent the children, the feature vector is $\mathbf{f} = (\alpha_{1,2}, \alpha_{1,3}, \theta_{2,3}, \beta_{1,2}, \beta_{1,3}, \beta_{2,3}, \gamma_{1,2}, \gamma_{1,3}, \gamma_{2,3}, \phi^3)$. The features were analysed at 1381 true bifurcations, giving the distribution parameters in Table 3, with $\phi^3 = 0.76 = 1050/1381$, given that 1050 of the 1381 training bifurcations had the ideal end orientation combination. As for the terminals and bridges, the likelihood is then calculated using the product of the estimated normal distribution functions for the continuous parameters and the Bernoulli mass function for the orientation.

We also estimate the parameters for the false joint distributions, using the same training set and methodology. The false terminal set consists of all non-terminal segment ends from the training set. The false bridge set is constructed by matching each bridge segment end in a crossing junction with its non-partners. The false bifurcation set is constructed by taking the true bifurcations and swapping the parent segment end with a child segment end. For brevity, the resulting distribution parameters are not shown.

4.2. Overlap Removal

An important source of errors in the system is misidentified segments. A particular problem occurs when two vessels run parallel and overlap, which causes segmentation of a single ‘‘overlap’’ segment. We detect such cases using the algorithm introduced in [2]. First, we run the algorithm as described above to perform a preliminary identification of all the joints. Second, we identify segments that have both ends identified as the parents of bifurcations in separate junctions (this occurs as the overlap segment typically has exaggerated width). Third, we merge the two junctions




Figure 4. Overlapping segments are falsely segmented as a single, wide ‘‘overlap segment’’. This is detected and eliminated, and the system detects the resulting pair of bridges.

at the ends of the overlapping segment, and discard that segment. Fourth, we run the algorithm for a second time; when the configuration is resolved it should select two bridges. See Figure 5(b).

5. Performance Evaluation and Discussion

The proposed system was evaluated on the test images of DRIVE using the vessel-classified version of DRIVE (ground truth) introduced in [10] to compare the systems output. The system was evaluated in two stages: first, using just the joint likelihood of the joints; second, the full MAP estimation taking into account the configuration priors. The system is considered successful only if it can correctly detect the configuration of all the joints at a junction: for example, if the system has correctly detected a bifurcation but wrongly identified the parent segment, it is considered a failure. Figure 6 shows the success rate at junctions with associated number of segments from one to six. The output demonstrates high stability, and illustrates that using the configuration priors to construct a MAP estimate improves performance.

Of the features used to estimate the joint likelihood functions, preliminary experiments suggest that the angle is most important, then width, then intensity (experimental results omitted due to space).

The MAP estimation significantly improves performance, although results are noticeably poorer for large junctions. This may be partly due to a relatively small number of examples.

6. Conclusion

This paper addresses the issue of determining connectivity of retinal vascular trees across complex junctions, and proposes a Bayesian inference model that enumerates the configurations of local cliques, and chooses the MAP estimate for the configuration of basic joints. A parametric model is established using a substantive custom-marked dataset. The experimental results are good, with 95.2% of junctions correctly resolved. The system is potentially applicable for other problems connecting segmented




Figure 5. Performance analysis of the junction configuration system. The square data-labels indicates the correct configuration based only on the joint likelihood of joints; and the circle data-labels shows the full MAP estimation including the configuration prior. A configuration is considered correct if all joints are correct.

branching structures. For future work we will explore other features of local joints (e.g. vessel entropy), global network connectivity constraints, and address issues where parts of the network are missing or are segmentation artefacts.

References

- [1] B. Al-Diri, A. Hunter, and D. Steel. An active contour model for segmenting and measuring retinal vessels. *IEEE Trans. Med. Imaging*, 28(9):1488–1497, 2009.
- [2] B. Al-Diri, A. Hunter, D. Steel, and M. Habib. Automated analysis of retinal vascular network connectivity. *Computerized Medical Imaging and Graphics*, 34(6):462–470, 2010.
- [3] G. Azzopardi and N. Petkov. Detection of retinal vascular bifurcations by trainable v4-like filters. *Computer Analysis of Images and Patterns, Springer*, pages 451–459, 2011.
- [4] E. Bullitt, S. Aylward, et al. 3d graph description of the intracerebral vasculature from segmented mra and tests of accuracy by comparison with x-ray angiograms. *Proc. IPMI*, 1613:308–320, 1999.
- [5] T. Deschamps and L. D. Cohen. Grouping connected components using minimal path techniques. *Proc. CVPR*, 2:102–109, 2001.
- [6] P. Gregson, Z. Shen, R. Scot, and V. Kozousek. Automated grading of venous beading. *Computers and Biomedical Research*, 28(4):291–304, 1995.
- [7] J. Jomier, V. LeDigarcher, and S. Aylward. Automatic vascular tree formation using the mahalanobis distance. *Proc. 8th MICCAI, Palm-Springs*, 8:806–812, 2005.
- [8] S. M. Lajevardi, A. Arakala, S. A. Davis, and K. J. Horadam. Retina verification system based on biometric graph matching. *IEEE Trans. Img. Processing*, 22(9):3625–3635, 2013.
- [9] T. A. Qureshi, H. Amin, et al. Automatic localization of the optic disc in retinal fundus images using multiple features. *Proc. IEEE. Conf. on Bioinformatics and Bioengineering (BIBE)*, pages 488–493, 2012.
- [10] T. A. Qureshi, M. Habib, A. Hunter, and B. Al-Diri. A manually labeled, artery/vein classified benchmark for the drive dataset. *Proc. IEEE. Int. Symp. on Computer-Based Medical Systems (CBMS)*, pages 485–488, 2013.
- [11] D. Selle, B. Preim, A. Schenk, and H.-O. Peitgen. Analysis of vasculature for liver surgical planning. *IEEE Trans. Med. Imaging*, 21(11):1344–1357, 2002.
- [12] J. Staal, M. Abramoff, et al. Ridge based vessel segmentation in color images of the retina. *IEEE Trans. Med. Imaging*, 23(4):501–509, 2004.
- [13] A. V. Stanton, B. Wasan, et al. Vascular network changes in the retina with age and hypertension. *Journal of Hypertension*, 13:1724–1728, 1995.
- [14] C. Tsai, C. Stewart, H. Tanenbaum, and B. Roysam. Model-based method for improving the accuracy and repeatability of estimating vascular bifurcations and crossovers from retinal fundus images. *IEEE Trans. on Info. Tech. in Biomedicine*, 8(2):122–130, 2004.
- [15] P. Yim, P. Choyke, and R. Summers. Gray-scale skeletonization of small vessels in magnetic resonance angiography. *IEEE Trans. Med. Imaging*, 19(1):568–576, 2000.