Detection of Retinal Vascular Bifurcations by Rotation- and Scale-Invariant COSFIRE Filters

George Azzopardi and Nicolai Petkov

Johann Bernoulli Institute for Mathematics and Computer Science
University of Groningen, The Netherlands
{g.azzopardi,n.petkov}@rug.nl

Abstract. The analysis of the vascular tree in retinal fundus images is important for identifying risks of various cardiovascular diseases. We propose trainable COSFIRE (Combination Of Shifted FIlter REsponses) filters to detect vascular bifurcations. A COSFIRE filter is automatically configured to be selective for a bifurcation that is specified by a user from a training image. The configuration selects given channels of a bank of Gabor filters and determines certain blur and shift parameters. A COSFIRE filter response is computed as the product of the blurred and shifted responses of the selected Gabor filters. The filter responds to bifurcations that are similar to the one used for its configuration. The proposed filters achieve invariance to rotation and scale. With only five COSFIRE filters we achieve a recall of 98.77% at a precision of 95.32% on a data set of 40 binary fundus images (from DRIVE), containing more than 5000 bifurcations.

Keywords: DRIVE, Gabor filters, retinal fundus, trainable filters, vessel bifurcation.

1 Introduction

Retinal fundus images provide a unique possibility to analyse the state of the vascular system of a person in a non-invasive way. The vascular tree of the retina is known to conform to structural principles that are related to certain physical properties [10, 14]. The analysis of the geometrical structure is important as deviations from the optimal principles may indicate some cardiovascular diseases, such as hypertension [17] and atherosclerosis [5]; a review is given in [12]. The identification of vascular bifurcations is one of the basic steps in this analysis which is currently being performed by a tedious manual process [5]. Automating the identification of such features is an essential step in the description of the vascular tree which is needed for further analysis.

The existing attempts to automate the detection of vascular bifurcations can be categorized into two classes usually referred to as geometrical-feature based and model based approaches. The former involve extensive preprocessing such as segmentation and skeletonization followed by local pixel processing and branch point analysis. These techniques are known for their robustness in bifurcation

A. Campilho and M. Kamel (Eds.): ICIAR 2012, Part II, LNCS 7325, pp. 363–371, 2012.

localization [3, 4, 6, 9]. On the other hand, model based approaches are usually more adaptive and have lower computational complexity which makes them more appropriate for real-time applications [1, 16]. However, model based approaches are known to suffer from insufficient generalization ability and consequently, they may fail to detect some relevant features.

We propose trainable filters which we call COSFIRE (Combination Of Shifted FIlter REsponses) and use for the detection of vascular bifurcations in retinal images. The proposed COSFIRE filters are trainable as they are configured to be selective for vascular bifurcations that are specified by a user. In a single-step training phase a user selects a typical bifurcation by a point of interest in a training image. The automatic configuration process extracts information about the properties of the contour segments in the concerned feature and their mutual geometrical arrangement. This training procedure can be repeated as many times as required in order to configure COSFIRE filters that are selective for different bifurcations. We demonstrate that the COSFIRE filters can then be applied in a rotation- and scale-invariant mode to detect bifurcations that are similar to the ones that were used to configure the filters.

The rest of the paper is organized as follows: in Section 2 we present the COSFIRE filter and demonstrate how it can be trained and used to detect retinal vascular bifurcations. In Section 3, we apply the COSFIRE filters on retinal images from the DRIVE data set [15]. We provide a discussion about some aspects of the proposed approach and draw conclusions in Section 4.

2 Proposed Method

2.1 Overview

Fig. 1a illustrates a typical bifurcation encircled in a binarized retinal image[1]. We use this feature, which is shown enlarged in Fig.1b, to automatically configure a COSFIRE filter that will respond to the same and similar bifurcations.

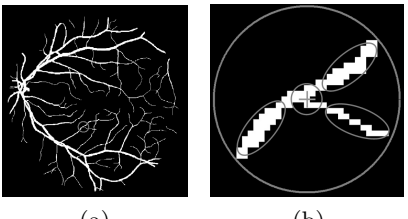

(a) (b)

Fig. 1. (a) The circle indicates a vascular bifurcation that is selected by a user. (b) Enlargement of the selected feature. The ellipses represent the dominant orientations around the specified point of interest that is shown by a '+' marker.

[1] The image is named 21_manual1.gif in the DRIVE data set [15].

The three ellipses shown in Fig.1b illustrate the dominant orientations in the surroundings of the point of interest. We use 2D Gabor filters to detect these orientations. The central circle represents the overlapping supports of a group of such Gabor filters. The response of a COSFIRE filter is computed by combining the responses of the concerned Gabor filters by multiplication.

The preferred orientations of these Gabor filters and the locations at which we take their responses are determined by a configuration process which analyses the specified pattern automatically. Taking the responses of Gabor filters at different locations around a point can be implemented by shifting the responses of these Gabor filters by different vectors before using them for the pixel-wise evaluation of a function which gives the COSFIRE filter output.

In the following, we explain the automatic configuration process of a COSFIRE filter that determines which responses of which Gabor filters in which locations need to be multiplied in order to obtain the output of the filter.

2.2 Detection of Dominant Orientations by 2D Gabor Filters

We use 2D symmetric Gabor filters to detect the presence of vessels in retinal images. We denote by $g_{\lambda,\theta}(x, y)$ the half-wave rectified response of a Gabor filter of preferred wavelength λ and orientation θ to a given input image. Such a filter has a number of other parameters, including spatial aspect ratio, bandwidth and phase offset, which we set as proposed in [13].

We re-normalize all Gabor functions that we use such that all positive values sum up to 1 and all negative values sum up to -1. This ensures that the response to a line of width w will be largest for a symmetrical Gabor filter of preferred wavelength $\lambda = 2w$ and that the response to an image of constant intensity is 0. Without re-normalization, this is true only for antisymmetrical filters.

For this application, we apply a bank of Gabor filters with five wavelengths equidistantly spaced on a logarithmic scale ($\lambda \in \{4, 4\sqrt{2}, 8, 8\sqrt{2}, 16\}$) and eight equidistant orientations ($\theta \in \{0, \frac{\pi}{8}, \ldots, \frac{7\pi}{8}\}$) on retinal images of size 565×584 pixels. We threshold the responses of Gabor filters at a given fraction $t_1 = 0.2$ of the maximum response of $g_{\lambda,\theta}(x, y)$ across all combinations of values (λ, θ) used and all positions (x, y) in the image, and denote these thresholded responses by $|g_{\lambda,\theta}(x, y)|_{t_1}$. This threshold value is sufficient to preserve all junction regions and suppress the undesirable responses of Gabor filters.

2.3 Configuration of a COSFIRE Filter

A COSFIRE filter uses as input the responses of a number of Gabor filters, each characterized by parameter values (λ, θ), around certain positions (ρ, ϕ) with respect to the center of the COSFIRE filter. A set of four parameter values $(\lambda, \theta, \rho, \phi)$ characterizes the properties of a contour part that is present in the specified pattern of interest: $\lambda/2$ represents the width, θ represents the orientation and (ρ, ϕ) represents the location. In the following we explain how we obtain the parameter values of such contour parts around a given point of interest.

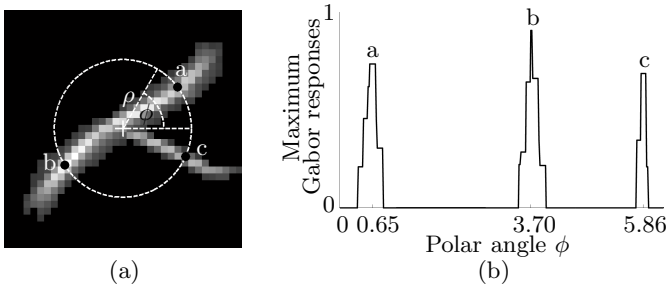

Fig. 2. (a) The gray-level of a pixel is the maximum value superposition of the thresholded ($t_1 = 0.2$) responses of a bank of Gabor filters. The white cross indicates the location of the specified point of interest, while the bright circle of a given radius ρ (here $\rho = 10$) indicates the considered locations in the analysis phase. (b) Values of the maximum value superposition of Gabor filter responses along the concerned circle. The labeled black dots in (a) mark the positions (relative to the center of the filter) at which the respective strongest Gabor filter responses are used. These three positions correspond to the three local maxima in the plot in (b).

We consider the responses of the bank of Gabor filters along a circle of a given radius ρ around the point of interest, Fig.2. In each position along that circle we take the maximum of all responses across the possible values of (λ, θ). The positions that have values greater than the values of the neighbouring positions along an arc of angle $\pi/8$ are chosen as the points that characterize the dominant orientations. We determine the polar coordinates (ρ_i, ϕ_i) for each such point with respect to the center of the filter. The pair of values (λ_i, θ_i) for which the concerned local maximum is reached are the preferred wavelength and orientation of a Gabor filter. In this way we obtain a tuple $(\lambda_i, \theta_i, \rho_i, \phi_i)$.

In our experiments we configure COSFIRE filters using multiple values of the parameter ρ. For non-zero values of ρ we determine a group of locations with the method mentioned above. For $\rho = 0$, we consider the responses of the bank of Gabor filters used at the specified point of interest where we consider all combinations of (λ, θ) for which the corresponding responses $g_{\lambda,\theta}(x, y)$ are greater than a fraction[2] $t_2 = 0.75$ of the maximum of $g_{\lambda,\theta}(x, y)$ across the different combinations of values (λ, θ) used. For each value θ that satisfies this condition, we consider a single value of λ, the one for which $g_{\lambda,\theta}(x, y)$ is the maximum of all responses across all values of λ. At this central location, $(\rho = 0, \phi = 0)$, multiple responses of Gabor filters characterized by different values of the parameters (λ, θ) can thus be used as inputs to the concerned COSFIRE filter.

We denote by $S_f = \{(\lambda_i, \theta_i, \rho_i, \phi_i) \,|\, i = 1 \dots n_f\}$ the set of parameter value combinations, which fulfill the above conditions. The subscript f stands for the local pattern around the selected point of interest. Every tuple in the set S_f specifies the parameters of some contour part in f.

[2] At the central location, which is the intersection of multiple vessels, a response is achieved by most of the Gabor filters that are tuned for different wavelengths and orientations. We use the threshold t_2 to consider only the dominant orientations.

For the point of interest shown in Fig.2a and two values of the radius parameter ρ ($\{0, 10\}$), the automatic selection method described above results in five contour parts with parameter values specified by the tuples in the following set: $S_f = \{(\lambda_1 = 4, \theta_1 = \frac{3\pi}{8}, \rho_1 = 0, \phi_1 = 0), (\lambda_2 = 8, \theta_2 = \frac{5\pi}{8}, \rho_2 = 0, \phi_2 = 0), (\lambda_3 = 4\sqrt{2}, \theta_3 = \frac{3\pi}{4}, \rho_3 = 10, \phi_3 = 0.65), (\lambda_4 = 4\sqrt{2}, \theta_4 = \frac{3\pi}{4}, \rho_4 = 10, \phi_4 = 3.70), (\lambda_5 = 4, \theta_5 = \frac{3\pi}{8}, \rho_5 = 10, \phi_5 = 5.86)\}$. The last tuple in this set, $(\lambda_5 = 4, \theta_5 = \frac{3\pi}{8}, \rho_5 = 10, \phi_5 = 5.86)$, for instance, describes a contour part with a width of ($\lambda_5/2 =$) 2 pixels and an orientation ($\theta_5 =$) $\frac{3\pi}{8}$ that can be detected by a Gabor filter with preferred wavelength ($\lambda =$) 4 and orientation ($\theta =$) $\frac{3\pi}{8}$, at a position of ($\rho_5 =$) 10 pixels to the south east ($\phi_5 = 5.86$) of the point of interest; this location is marked by the label 'c' in Fig. 2. This selection is the result of the presence of a diagonal vessel to the south east of the center of the bifurcation pattern that is used for the configuration of the COSFIRE filter.

2.4 Blurring and Shifting Gabor Filter Responses

The above analysis of the bifurcation feature f produces five strong responses $g_{\lambda_i,\theta_i}(x, y)$ of Gabor filters in the positions with polar coordinates (ρ_i, ϕ_i) with respect to the filter center. Next we use these responses to compute the output of the COSFIRE filter. We first shift these responses appropriately so that they meet at the filter center. The COSFIRE filter output can then be evaluated as a pixel-wise multivariate function of the shifted Gabor filter responses.

Before these shift operations, we blur the Gabor filter responses in order to allow for some tolerance in the position of the respective contour parts. We define the blurring operation as the computation of maximum value of the weighted responses of a Gabor filter. For weighting we use a Gaussian function $G_\sigma(x, y)$, the standard deviation σ of which is a linear function of the distance ρ from the center of the COSFIRE filter: $\sigma = 0.67 + 0.1\rho$. In practice, we implement this computation as: $\max_{x',y'}\{g_{\lambda,\theta}(x-x', y-y')G_\sigma(x', y')\}$ where $-3\sigma \leq x', y' \leq 3\sigma$.

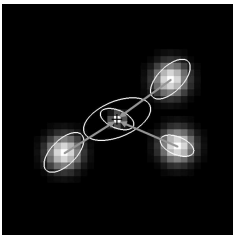

Fig. 3. Structure of a COSFIRE filter that is configured for the bifurcation shown in Fig. 2a. This filter is trained to detect the local geometrical arrangement of five contour parts. The ellipses illustrate the orientations and wavelengths of the selected Gabor filters, and the bright blobs are intensity maps of 2D Gaussian functions that are used to blur the responses of the corresponding Gabor filters. The blurred responses are then shifted by the corresponding vectors.

Next we shift the blurred responses of each selected Gabor filter (λ_i, θ_i) by the shift vector $(\rho_i, \phi_i + \pi)$, Fig.3. We denote by $s_{\lambda_i,\theta_i,\rho_i,\phi_i}(x, y)$ the blurred and shifted response of the Gabor filter that is specified by the i-th tuple in the set S_f. In practice, the computation of one blurred and shifted response (for the same values of the parameters λ, θ and ρ), for instance with $s_{\lambda,\theta,\rho,\phi=0}(x, y)$, is sufficient: the result of $s_{\lambda,\theta,\rho,\phi}(x, y)$ for any value of ϕ can be obtained from the result of the output of $s_{\lambda,\theta,\rho,\phi=0}(x, y)$ by appropriate shifting.

2.5 COSFIRE Filter (Rotation- and Scale-Invariant) Response

We define the response $r_{S_f}(x, y)$ of a COSFIRE filter as the geometric mean of all the blurred and shifted Gabor filter responses $s_{\lambda_i,\theta_i,\rho_i,\phi_i}(x, y)$ that correspond to the properties of the contour parts described by S_f.

$$r_{S_f}(x, y) = \left| \left(\prod_{i=1}^{|S_f|} s_{\lambda_i,\theta_i,\rho_i,\phi_i}(x, y) \right)^{\frac{1}{|S_f|}} \right|_{t_3} \quad (1)$$

where $|.|_{t_3}$ stands for thresholding the response at a fraction t_3 of its maximum.

Rotation and scale invariance are achieved by manipulating the set of parameter values in S_f, rather than by computing them from the responses to rotated and/or thicker versions of the original pattern. Using the set S_f that defines the concerned filter, we first form a new set $\Re_{\psi,\nu}(S_f) = \{(\nu\lambda_i, \theta_i + \psi, \nu\rho_i, \phi_i + \psi) | \forall\ (\lambda_i, \theta_i, \rho_i, \phi_i) \in S_f\}$. We then define the rotation- and scale-invariant response as $\widehat{r}_{S_f}(x, y) = \max_{\psi,\nu}(r_{\Re_{\psi,\nu}(S_f)}(x, y))$ across all combinations of values (ψ, ν).

3 Experimental Results

We use the bifurcation illustrated in Fig.1 to configure a filter denoted by S_{f_1} with three values of the parameter ρ ($\rho \in \{0, 4, 10\}$). Fig.4a shows the result ($t_3 = 0.2$) of filter S_{f_1} applied without rotation- and without scale-invariance ($\psi = 0,\ \nu = 1$) to the retinal image shown in Fig.1a. The encircled regions are centered on the local maxima of the filter response and if two such regions overlap by 75%, only the one with the stronger response is shown. Besides the original bifurcation the filter detects another feature with similar vessel orientations.

If the same filter S_{f_1} is applied in a scale-invariant mode ($\psi = 0,\ \nu \in \{2^{-\frac{1}{2}}, 1, 2^{\frac{1}{2}}\}$), it detects another bifurcation in a similar orientation but with thicker vessels, Fig. 4b. If we apply rotation-invariance but omit scale-invariance ($\psi \in \{\frac{\pi i}{8} \mid i = 0 \dots 7\}$, $\nu = 1$), filter S_{f_1} detects 38 similar bifurcations, Fig. 4c. Finally, when we use both rotation- and scale-invariance ($\psi \in \{\frac{\pi i}{8} \mid i = 0 \dots 7\}$, $\nu \in \{2^{-\frac{1}{2}}, 1, 2^{\frac{1}{2}}\}$), a total of 52 correct features are detected. This illustrates the strong generalization capability of this approach because 48.60% (52 out of 107) of the features of interest are detected by a single filter.

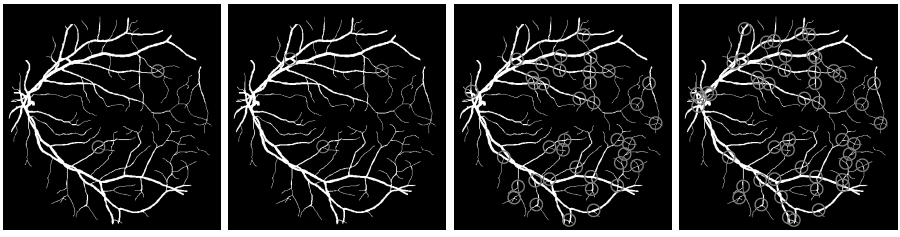

(a) (b) (c) (d)

Fig. 4. (a) Result of applying the filter S_{f_1} (for $t_3 = 0.2$) in rotation- and scale-non-invariant mode, (b) rotation-non-invariant but scale-invariant mode, (c) rotation-invariant but scale-non-invariant mode and (d) rotation- and scale-invariant mode

As to the features that are not detected by this filter, we proceed as follows: we use one of these features that we denote by f_2 to train a second filter S_{f_2}, Fig. 5. With the new filter we detect 42 features of interest ($t_3 = 0.33$) of which 13 coincide with features detected by filter S_{f_1} and 29 are newly detected features. Merging the responses of both filters results in the detection of 81 features. By configuring another filter S_{f_3} (Fig. 5) and use it ($t_3 = 0.14$) along with the other two filters we achieve 100% recall and precision[3] for the concerned image.

We apply the three COSFIRE filters on 40 binary retinal images[4] taken from the DRIVE data set [15]. The ground truth data[5], which comprises the coordinates of bifurcations and cross-overs in the selected images, is defined by the authors of this paper. For this data set we achieve a total recall of 97.23% and a total precision of 94.44%. We configure two more COSFIRE filters, S_{f_4} and S_{f_5} (Fig. 5) for two patterns[6] that are not detected by the other three filters. With only five COSFIRE filters, the maximum of the harmonic mean[7] is reached for a recall of 98.77% and precision of 95.32% that are achieved with the following set of threshold values: $t_3(S_{f_1})=0.21,\ t_3(S_{f_2})=0.33,\ t_3(S_{f_3})=0.18,\ t_3(S_{f_4})=0.32,\ t_3(S_{f_5})=0.24$.

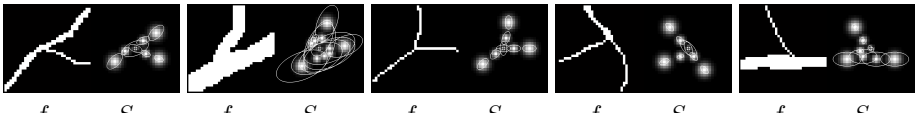

f_1 S_{f_1} f_2 S_{f_2} f_3 S_{f_3} f_4 S_{f_4} f_5 S_{f_5}

Fig. 5. Five bifurcations and the structures of the corresponding five COSFIRE filters

[3] Recall is the percentage of true bifurcations that are successfully detected. Precision is the percentage of correct bifurcations from all detected features.
[4] Named in DRIVE 01_manual1.gif, 02_manual1.gif, ..., 40_manual1.gif.
[5] Ground truth data is online: www.cs.rug.nl/~imaging/databases/retina_database
[6] Patterns f_4 and f_5 are cropped from 01_manual1.gif and 32_manual1.gif, respectively.
[7] Harmonic mean of precision P and recall R is defined as $2PR/(P+R)$.

4 Discussion and Conclusion

In this work we enrich the method proposed in [2] by adding scale-invariance and use individual t_3 threshold values. With only five COSFIRE filters we achieve 98.77% recall and 95.32% precision on a data set of 40 binary retinal images (containing more than 5000 bifurcations). These results outperform the ones (98.52% recall, 95.19% precision) reported in [2] that were achieved with 25 trainable filters, the responses of which were thresholded at the same value. In [3] a recall of 95.82% is reported on a small data set of five retinal images.

A COSFIRE filter is trainable, in that the structure of the filter is determined by a feature that is specified by a user. The way this is achieved is not by template matching, but rather by the extraction of information about the dominant orientations in the concerned feature and their mutual geometrical arrangement. This versatility allows the user to improve the performance by configuring new filters selective for specific features that are not detected by the existing filters.

Notable is the fact that we achieve the above results without any optimization in the selection of bifurcations that we used to configure the five COSFIRE filters. In our experiments, we configured these filters in an iterative way where in each iteration we chose a random bifurcation from the set of bifurcations that were not detected by the filters that were configured up to that iteration. For the 40 binary images (from DRIVE) it turned out that five iterations were sufficient to achieve the above high precision and recall rates. We only tune the individual t_3 threshold parameters to values that contribute to the best performance results. On a standard computer (3GHz processor) a retinal image of 565×584 pixels took approximately 30 seconds to be processed by the five COSFIRE filters.

The proposed COSFIRE filters share similar properties with some neurons in area V4 of visual cortex [11]. The function that we use to compute the filter response is a geometric mean and it is based on psychophysical evidence that human visual processing of shape is likely performed by multiplication [7].

A COSFIRE filter produces a response only when all constituent parts of a pattern of interest are present. This is in contrast to other feature detection approaches, such as SIFT [8], in which two patterns are considered similar if the Euclidean distance between their respective local descriptors satisfies certain criteria. Euclidean-distance based approaches, however, suffer from insufficient selectivity regarding the shape properties of features.

The COSFIRE filters that we propose are conceptually simple and easy to implement: the filter output is computed as the product of blurred and shifted Gabor filter responses. They are versatile keypoint detectors as they can be configured with any specified local contour patterns. In future work we will evaluate the proposed COSFIRE filters on grayscale fundus images and on binary fundus images that are obtained by automatic segmentation methods.

In principle, all vascular bifurcations can be detected if a sufficient number of filters are configured and used. The performance that we obtain is sufficient to the needs of the medical application at hand. We conclude that the proposed rotation- and scale-invariant COSFIRE filters are an effective means to automatically detect bifurcations in retinal fundus images.

References

1. Ali, C., Hong, S., Turner, J., Tanenbaum, H., Roysam, B.: Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms. IEEE Transactions on Information Technology in Biomedicine 3, 125–138 (1999)
2. Azzopardi, G., Petkov, N.: Detection of Retinal Vascular Bifurcations by Trainable V4-Like Filters. In: Real, P., Diaz-Pernil, D., Molina-Abril, H., Berciano, A., Kropatsch, W., et al. (eds.) CAIP 2011, Part I. LNCS, vol. 6854, pp. 451–459. Springer, Heidelberg (2011)
3. Bhuiyan, A., Nath, B., Chua, J., Ramamohanarao, K.: Automatic detection of vascular bifurcations and crossovers from color retinal fundus images. In: Third International IEEE Conference on Signal-Image Technologies and Internet-Based System (SITIS), pp. 711–718 (2007)
4. Chanwimaluang, T., Guoliang, F.: An efficient blood vessel detection algorithm for retinal images using local entropy thresholding. In: Proceedings of the 2003 IEEE International Symposium on Circuits and Systems (Cat. No.03CH37430) (2003)
5. Chapman, N., Dell'omo, G., Sartini, M., Witt, N., Hughes, A., Thom, S., Pedrinelli, R.: Peripheral vascular disease is associated with abnormal arteriolar diameter relationships at bifurcations in the human retina. Clinical Science, 103
6. Eunhwa, J., Kyungho, H.: Automatic retinal vasculature structure tracing and vascular landmark extraction from human eye image. In: International Conference on Hybrid Information Technology, 7 (2006)
7. Gheorghiu, E., Kingdom, F.: Multiplication in curvature processing. Journal of Vision 9 (2009)
8. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60, 91–110 (2004)
9. Martinez-Perez, M., Hughes, A., Stanton, A., Thom, S., Chapman, N., Bharath, A., Parker, K.: Retinal vascular tree morphology: A semi-automatic quantification. IEEE Transactions on Biomedical Engineering 49, 912–917 (2002)
10. Murray, C.D.: The physiological principle of minimum work applied to the angle of branching of arteries 9, 835–841 (1926)
11. Pasupathy, A., Connor, C.E.: Responses to contour features in macaque area v4. Journal of Neurophysiology 82, 2490–2502 (1999)
12. Patton, N., Aslam, T., MacGillivray, T., Deary, I., Dhillon, B., Eikelboom, R., Yogesan, K., Constable, I.: Retinal image analysis: Concepts, applications and potential. Progress in Retinal and Eye Research 25, 99–127 (2006)
13. Petkov, N.: Biologically motivated computationally intensive approaches to image pattern-recognition. Future Generation Computer Systems 11, 451–465 (1995)
14. Sherman, T.F.: On connecting large vessels to small - the meaning of murray law. Journal of General Physiology 78, 431–453 (1981)
15. Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., van Ginneken, B.: Ridge-based vessel segmentation in color images of the retina. IEEE Transactions on Medical Imaging 23, 501–509 (2004)
16. Tsai, C., Stewart, C., Tanenbaum, H., Roysam, B.: Model-based method for improving the accuracy and repeatability of estimating vascular bifurcations and crossovers from retinal fundus images. IEEE Transactions on Information Technology in Biomedicine 8, 122–130 (2004)
17. Tso, M., Jampol, L.: Path-physiology of hypertensive retinopathy. Opthalmology, 89