Asociación Argentina
de Mecánica Computacional

Mecánica Computacional Vol XXXIII, págs. 2729-2743 (artículo completo)
Graciela Bertolino, Mariano Cantero, Mario Storti y Federico Teruel (Eds.)
San Carlos de Bariloche, 23-26 Setiembre 2014

REVIEWING PREPROCESSING AND FEATURE EXTRACTION TECHNIQUES FOR RETINAL BLOOD VESSELS SEGMENTATION IN FUNDUS IMAGES

José Ignacio Orlando[a,b] **and Mariana del Fresno**[a,c]

[a]*Instituto Pladema, Universidad Nacional del Centro de la Provincia de Buenos Aires, UNCPBA Campus Universitario, Paraje Arroyo Seco, (B7000), Tandil, Buenos Aires, Argentina, pladema@exa.unicen.edu.ar, http://www.pladema.net*

[b]*Consejo Nacional de Investigaciones Científicas y Técnicas, CONICET Av. Rivadavia 1917, (C1033AAJ), Ciudad Autónoma de Buenos Aires, Argentina, info@conicet.gov.ar, http://www.conicet.gov.ar*

[c]*Comisión de Investigaciones Científicas de la Provincia de Buenos Aires, CIC-PBA Calle 526 entre 10 y 11, (B1900), La Plata, Buenos Aires, Argentina info@cic.gba.gov.ar, http://cic.gba.gov.ar/*

Keywords: Fundus imaging, Blood vessels segmentation, Feature engineering, Support Vector Machines.

Abstract. Several ophthalmological and cardiovascular diseases–such as diabetic and hypertensive retinopathies, choroidal neovascularization, arteriosclerosis, among others–can be diagnosed by analysing the structure of the retinal vasculature. Such analysis require to count with precise segmentation of blood vessels, being manual delineation tedious and time-consuming. Various algorithms for automatic blood vessel segmentation have been proposed in the last years, most based on supervised methods. These approaches deal with the automatic detection of retinal blood vessel features and non-vessel features by learning on the basis of a training set of manually segmented reference images. Performance of such methods is usually determined by the features capability of discriminating vessels from other anatomical or pathological structures. In this work, we present a review of different preprocessing and feature extraction techniques for blood vessel segmentation in retinal images. Using a linear Support Vector Machine as the segmentation approach, we study the behaviour of several state-of-the-art preprocessing and feature extraction techniques in the detection of retinal vasculature, summarizing their computation and results. Finally, we propose a standard methodology to evaluate and compare blood vessel segmentation algorithms. A publicly available data set of fundus images is employed for evaluation, and our results are compared against other state-of-the-art approaches.

1 INTRODUCTION

Optical fundus images–also known as retinal images or retinographies–are projective color photographs of the inner surface of the human eye (Figure 1 (a)). Such images allow ophthalmologists and other physicians to directly observe the retina and their internal parts, including the vascular tree, the optic disc and the fovea (Abràmoff et al., 2010). Quantitative measurements of variations in the retinal vasculature are frequently used in diagnosing several ophthalmological and vascular diseases, including diabetic retinopathy, hypertension, occlusion and glaucoma (Fraz et al., 2012a). Vessel occlusion, for instance, makes vessels longer; hypertension reduces arteries, while diabetes creates new blood vessels (Fathi and Naghsh-Nilchi, 2013). To compute morphological features–such as length, width, tortuosity or branching patterns and angles–precise segmentations are needed. However, manual delineation is an intensive and time consuming labour, specially in screening programs where large databases of retinal images need to be analysed (Abràmoff et al., 2010). Thus, automatic blood vessel segmentation in retinal images is extremely valuable.

Several works have been introduced in the last years to solve this problem; yet, none of the available methods have been assumed as a standard by the medical community due to their lack of precision (Abràmoff et al., 2010). These methods have to detect both narrow and wide vessels in the presence of illumination artifacts and/or pathological structures, such as red-lesions, cotton wool spots or exudates. In general, these segmentation methods can be classified into two categories: supervised and unsupervised. Supervised methods require a set of labeled samples and their features to learn a model–usually a classifier–which is then applied over new and non-annotated images to obtain the final segmentation. Most of the effort in supervised segmentation involves finding new features for training, or better classifiers to perform the pixel classification task (Orlando and Blaschko, 2014). For instance, Marín et al. (2011) propose a supervised method based on a combination of gray-level and moment invariants-based features with neural networks. Lupascu et al. (2010) utilize an AdaBoost classifier trained with several state-of-the-art features for blood vessel segmentation. The work by Osareh and Shadgar (2009), on the contrary, employs Gabor filters to train a generative Gaussian mixture model. Unsupervised methods are based on models which are able to adjust themselves to obtain the segmentation. This type of methods involve systems based on thresholding (Chaudhuri et al., 1989), region-oriented approaches (Martinez-Perez et al., 2007), clustering techniques (Salem et al., 2007) or morphological operators Zana and Klein (2001), among others.

Support Vector Machines (SVMs) are supervised learning models widely utilized in pattern recognition (Abe, 2010). Given a set of training samples, each one identified as belonging to one of two classes, a SVM training algorithm finds a model that assigns new examples into one class or the other. Samples in the training set are represented as points in a multidimensional feature space mapped so that the examples of the separate classes are divided by a clear gap which is as wide as possible. New examples are mapped into the same space and predicted to belong to a class based on which side of the gap they fall on (Hastie et al., 2009). This definition exactly matches the problem of blood vessel segmentation, where a pixel has to be classified as part of a vessel or not according to its features. Several authors have considered SVMs for retinal vasculature segmentation, introducing novel features to train the classifier. Ricci and Perfetti (2007) propose to combine line detectors and linear SVMs. Osareh and Shadgar (2009) prove that SVMs trained with Gabor filters responses outperforms generative Gaussian mixture models. Recently, a more sophisticated version of SVMs, namely structured SVM, have been used by Orlando and Blaschko (2014) to learn fully-connected conditional random fields for

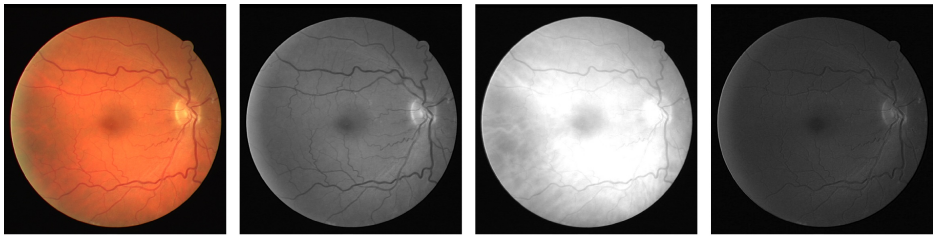

Figure 1: RGB color bands of a retinal image. (a) Original color image. (b) Green, (c) red and (d) blue bands.

retinal blood vessels segmentation.

The performance of SVMs for vessel segmentation, like in other supervised methods, strongly depends on the ability of the selected features to distinguish vessels from non-vessel structures (Becker et al., 2013). In this paper, we quantitatively review the current best practices in preprocessing and feature extraction techniques for supervised blood vessel segmentation in retinal images. Using a publicly available data set of fundus photographs, we analyse different combinations of features and preprocessing methods to study their ability to train a linear SVM. We also propose to consider different vessel-enhancement techniques utilized in other existing unsupervised methods as features to train the SVM classifier. Feature selection using a forward selection approach is performed to find the best combination of features. Results prove that by preprocessing the image with different techniques according to the feature to be extracted it is possible to improve the final segmentations. We also show that our results outperform SVMs based on features extracted from non-preprocessed images. Finally, we demonstrate that our approach achieve results comparable to other state-of-the-art methods. An evaluation strategy adapted to the characteristics of this segmentation problem is additionally proposed to standardize the way this methods are compared.

The remainder of this paper is organized as follows. In Section 2 we review some of the most utilized state-of-the-art features, including details about their implementations. Section 3 describes some of the existing preprocessing techniques to reduce their influence. Details about SVMs are analyzed in Section 4. Section 5 summarizes our experiments and the results we obtained, compared with other state-of-the-art works. Finally, Section 6 concludes the paper.

2 FEATURES FOR BLOOD VESSEL SEGMENTATION

Several features have been introduced for blood vessel segmentation in retinal images in the last years. In general, all of them are extracted from the green channel of the original color image, since it exhibits the best vessel/background contrast while the red and blue ones tend to be very noisy (Figure 1). Current approaches can be classified according to the information they utilize to characterize blood vessels in the following categories: based on image intensities (Chaudhuri et al., 1989; Sinthanayothin et al., 1999; Marín et al., 2011), ridge detectors based on second-order derivatives (Frangi et al., 1998; Martinez-Perez et al., 2007; Staal et al., 2004), line detectors (Ricci and Perfetti, 2007; Nguyen et al., 2012), wavelets (Soares et al., 2006) and features based on morphological operations (Zana and Klein, 2001; Marín et al., 2011).

In the following subsections we describe all the features analysed in this work. Although more complex features have been also proposed in the state-of-the-art, in this paper we summarize the most widely applied for supervised segmentation (Fraz et al., 2012a).

2.1 Intensity based features

Intensity-based features (Figure 2) are based on the application of different filters to enhance blood vessels. Those filters can be typical from the image processing domain, such as mean or standard deviation filters, or can be designed according to local properties of the vessels, as in the case of matched filter responses.

Matched filter responses (F1.1) were introduced by Chaudhuri et al. (1989), based on the idea that vessel cross sectional profiles traditionally look like an inverse Gaussian. Since vessels may be considered as piecewise linear segments, matching a number of cross section (of identical profile) along its length simultaneously result in a vessel enhanced image that can be utilized as a feature. Such a kernel may be mathematically expressed as:

$$K(x,y) = -\exp\left(-\frac{x^2}{2\sigma^2}\right) \qquad |y| \leq \frac{L}{2} \tag{1}$$

where L is the length of the segment for which the vessel is assumed to have a fixed orientation, and σ is the standard deviation of the Gaussian kernel. Image is convolved applying K at different orientations (here, angles from 0º to 165º with a step of 15º), and the maximum response over orientations is taken as a feature. By changing the value of σ it is possible to capture responses of different width vessels. In our experiments, we fix $L = 9$ and $\sigma^2 = \{1, 2, 4, 8\}$. Responses at each scale are considered as individual features.

Sinthanayothin et al. (1999) propose a preprocessing technique that we suggest to include as a feature, named **Sinthanayothin *et al.* feature (F1.2)**. Authors work with an intensity-hue-saturation representation of the original RGB image, but we prefer to use the green band of the image instead because of its contrast. The feature consists on a vessel enhanced image obtained by applying the locally adaptive transformation:

$$I'(i,j) = \left(1 + \exp\left(\frac{\underset{(s,t)\in N_{i,j}^{49}}{\text{mean}}\{I\} - I}{\underset{(s,t)\in N_{i,j}^{49}}{\text{std}}\{I\}}\right)\right)^{-1} \tag{2}$$

Additionally intensity-based features were introduced by Marín et al. (2011). They are obtained after applying a combination of morphological and filtering operations. First, morphological opening using a three-pixel diameter disc with 8-connectivity is applied to the green band of the image to remove the brighter strip in the center of some vessels, resulting in a new image I_γ. Afterwards, a mean filter using square windows of side 3 is applied to remove occasional salt&pepper noise. Further noise is smoothed with an additional filtering based on a Gaussian filter with 9×9 neighbourhoods, mean $\mu = 0$ and variance $\sigma^2 = 1.8^2$. The background of the image I_γ is then homogenized by subtracting a background estimated from the noise smoothed image using a 69×69 mean filter. Intensities of the resulting image–namely I_{SC} by Marín et al. (2011)–are then linearly mapped in the interval $[0, 255]$ to cover the entire range. Finally, an homogenized image I_H is obtained by applying the following transformation:

$$I_H = \begin{cases} 0 & \text{if } \; g < 0 \\ 255 & \text{if } \; g > 255 \\ g & \text{otherwise} \end{cases} \tag{3}$$

where $g = I_{SC} + 128 - \text{mode}(I_{SC})$.

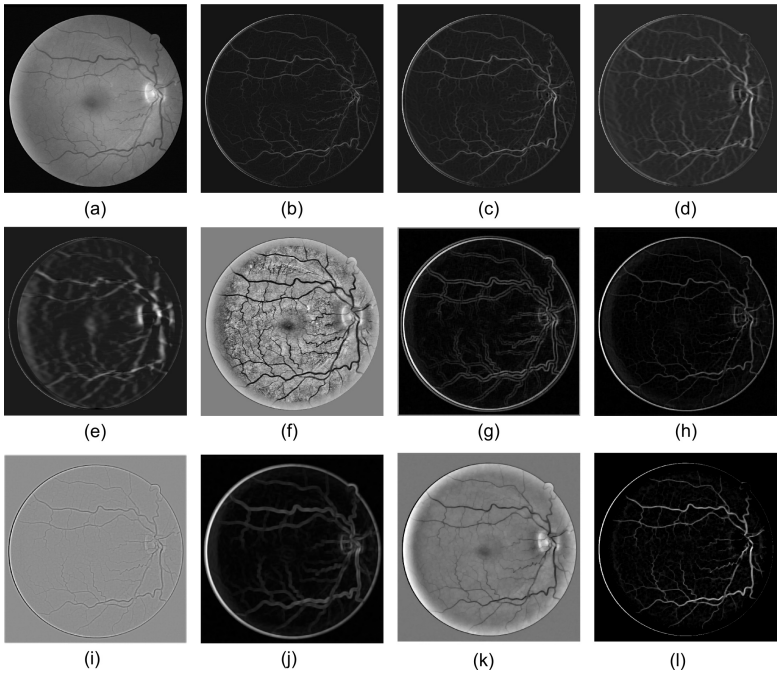

Figure 2: Intensity based features. (a) Green band of a retinal image. (b-e) Matched filter responses (F1.1) over scales $\sigma^2 = \{1, 2, 3, 4\}$, (f) Sinthanayothin *et al.* (F1.2), (g-k) Marin *et al.* intensity based features (F1.3), (k) Marin *et al.* homogenized image (F1.4) and (l) Marin *et al.* enhanced image (F1.5).

Five features, named **Marin *et al.* intensity based features (F1.3)**, are extracted from I_H:

$$f_1(i,j) = I_H(i,j) - \min_{(s,t)\in N^9_{i,j}} \{I_H(s,t)\}, \quad f_2(i,j) = \max_{(s,t)\in N^9_{x,y}} \{I_H(s,t)\} - I_H(i,j) \tag{4}$$

$$f_3(i,j) = I_H(i,j) - \operatorname*{mean}_{(s,t)\in N^9_{i,j}} \{I_H(s,t)\}, \quad f_4(i,j) = \operatorname*{std}_{(s,t)\in N^9_{i,j}} \{I_H(s,t)\}$$

$$f_5(i,j) = I_H(i,j)$$

We consider I_H as an additional feature, called **Marin *et al.* homogenized image (F1.4)**. Our last intensity-based feature is the vessel enhanced image I_{VE}, obtained by applying a top-hat transformation with a disc of eight pixels of diameter over I_H. This image is utilized in (Marín et al., 2011) to obtain moment invariant features. However, we suggest to include them as a feature since it allows to distinguish the vessels properly. We name it **Marin *et al.* enhanced image (F1.5)**.

2.2 Ridge detectors based on second order derivatives

Ridge detectors based on second order derivatives are utilized as features since blood vessels appear as ridge-like structures in the images (Figure 3) (Martinez-Perez et al., 2007). Following this prior, we can identify vessels by looking for pixels where the intensity image has a local maximum in the direction for which the gradient of the image undergoes the largest change. First, images are convolved with Gaussian filters with different values of σ in order to capture vessels at different scales. Afterwards, information about second order derivative of the image is derived from the Hessian matrix, which is symmetrical with real eigenvalues and orthogonal and rotation invariant eigenvectors. The eigenvalues λ_1 and λ_2, with $|\lambda_1| \geq |\lambda_2|$, measure convexity

and concavity in the corresponding eigendirections, respectively (Martinez-Perez et al., 2007). Mathematical operations between eigenvalues are usually employed as features.

Frangi et al. (1998) introduce a popular feature for blood vessel enhancement in 2D and 3D images based on second order derivatives. First, image at different scales σ_i is obtained. A ridgeness score $\mathcal{R}_\mathcal{B} = \lambda_1/\lambda_2$ is computed, and this value is combined with the Frobenius norm of the Hessian, $\mathcal{S}$, which measures overall strength, into a single vesselness measure, given by the following expression:

$$\mathcal{V}_0 = \begin{cases} 0 & \lambda_2 > 0 \\ \exp\left(-\frac{\mathcal{R}_\mathcal{B}^2}{2\beta^2}\right)\left(1-\exp\left(-\frac{\mathcal{S}^2}{2c^2}\right)\right) & \text{otherwise} \end{cases} \tag{5}$$

where $\beta = 0.5$ and c is equal to half of the maximum Frobenius norm of the Hessian. This vesselness measure is applied over each scale of the image, and the maximum over scales is utilized to characterize retinal vessels. We also include two additional features, the scale where $\mathcal{V}_0$ is maximum and the angle of the minor eigenvector, resulting in a three-dimensional feature named **Frangi *et al.* vesselness measure (F2.1)**. In our experiments, we consider values of $\sigma = \{1, 2, 3\}$.

In (Martinez-Perez et al., 2007), authors propose to use the value of λ_2, named **Maximum principal curvature (F2.2)**, as a feature. This value is evaluated over different scales of the image in order to capture responses for both narrow and wide vessels. Instead of grouping all the responses into a single feature, we suggest to utilize each individual response as a feature, learning the weights to combine them by using the linear SVM.

Another feature based on second order derivatives can be derived from Staal et al. (2004), where authors propose to describe blood vessels by composing primitives in the form of line elements extracted from image ridges. In this paper we propose to utilize those line elements as features. Let I_σ be the original image I in a given scale σ. A scalar field $\rho(\mathbf{p}, \sigma)$ can be defined over I_σ according to the eigenvalue λ_2 of the Hessian matrix, by taking -1 for ridges of local minima ($\lambda_2 > 0$), 1 for ridges of local maxima ($\lambda_2 < 0$) and 0 elsewhere:

$$\rho(\mathbf{p}, \sigma) = -\frac{1}{2}\text{sign}(\lambda_2)|\text{sign}(\nabla I_\sigma(\mathbf{p} + \epsilon\hat{\mathbf{v}}) \cdot \hat{\mathbf{v}}) - \text{sign}(\nabla I_\sigma(\mathbf{p} - \epsilon\hat{\mathbf{v}}) \cdot \hat{\mathbf{v}})| \tag{6}$$

where $\mathbf{p}$ is the position vector of a given pixel, ∇ is the gradient vector defined over I_σ and $\hat{\mathbf{v}}$ is the eigenvector corresponding to the eigenvalue λ_2. The parameter ϵ is the spatial accuracy with which the point-sets are detected, and it is set in 1. Since the extraction of this feature is computationally expensive, we only compute it using a single scale $\sigma = 1.5$. We name this feature **Staal *et al.* ridge-line elements (F2.3)**.

2.3 Line detectors

Line detectors (Figure 4) are based on the prior that blood vessels appear as elongated structures in the image. This type of features is obtained by evaluating the average intensity along lines passing through the target pixel at different orientations. The original idea was introduced in (Zwiggelaar et al., 2004) for the detection of linear structures in mammographic images, although it is modified by Ricci and Perfetti (2007) in order to fit with the retinal vessel segmentation problem. By taking the inverted green channel of the original image, the average intensity is evaluated along lines of fixed length $l = 15$ passing through each target pixel (i, j) at different orientations (angles from 0º to 165º with a step of 15º). The line with the largest average intensity, $L(i, j)$, is found, and the difference $S(i, j) = L(i, j) - N(i, j)$ is taken as the line

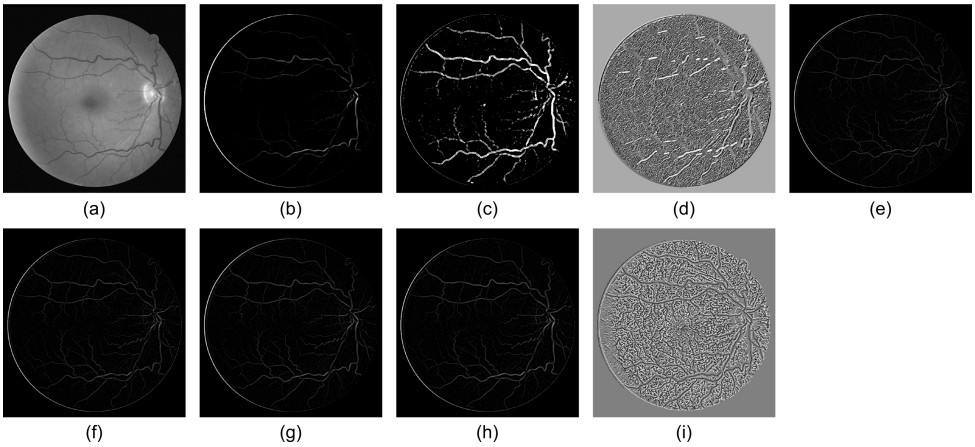

Figure 3: Ridge detectors based on second order derivatives. (a) Green band of a retinal image. (b-d) Frangi *et al.* vesselness measures (F2.1), (e-h) Maximum principal curvature (F1.2) over scales $\sigma^2 = \{1, 2, 3, 4\}$, (i) Staal *et al.* ridge-line elements (F2.3).

strength of the pixel, where $N(i,j)$ represents the average intensity in the square window, centered on the pixel, with edge length equal to l. Additionally, a line of three pixels is considered, centered on the midpoint of the main line and orthogonal to it. Its average intensity, $L_o(i,j)$, is utilized to obtain another strength feature $S_o(i,j) = L_o(i,j) - N(i,j)$ (Figure X (c)). The angle of the orthogonal line is approximated with the closest principal angle $(0^o, 45^o, 90^o, 135^o)$. Both features, jointly with the inverted green band of the image, are proposed as features by the authors. In this work, we only consider $S(i,j)$ and $S_o(i,j)$, and we call it **Ricci and Perfetti line detector (F3.1)**.

Ricci and Perfetti line detector gives false responses under certain circumstances such as at a background pixel between two close vessels, at a background pixel at the corner of a crossover point, or at a background pixel near a strong vessel (Nguyen et al., 2012). To reduce these wrong detections, Nguyen et al. (2012) propose to compute line responses at varying scales, combining them linearly with the inverted green channel into a single feature. Instead of taking $S(i,j) = L(i,j) - N(i,j)$, the new line strength measure is computed by taking $S(i,j) = L^l(i,j) - N^s(i,j)$, with $1 \leq l \leq s$. Parameter l is the length of the line passing through each target pixel, and s is the size of the neighbourhood from where the average intensity is obtained. Since our segmentation method is able to learn by itself the weights to combine each scale of the feature, we consider each scale as a feature, named **Nguyen *et al.* line detector (F3.2)**. The scales considered in our experiments are from 1 to 15 with a step of 2.

2.4 Wavelets

Wavelets have been extensively explored in the literature for signal and image processing (Soares et al., 2006). Regarding retinal blood vessels segmentation, wavelets have been applied to enhance vessels by Soares et al. (2006). We include **Soares *et al.* Gabor wavelet (F4.1)** in our experiments due to their capability of distinguish both narrow and wide vessels (Figure 5). By taking the inverted green channel of the retinal image, authors creates feature vectors including measurements at different scales taken from the two-dimensional (2-D) Gabor wavelet transform, also known as Morlet wavelet. Each response is utilized in this work as an individual feature.

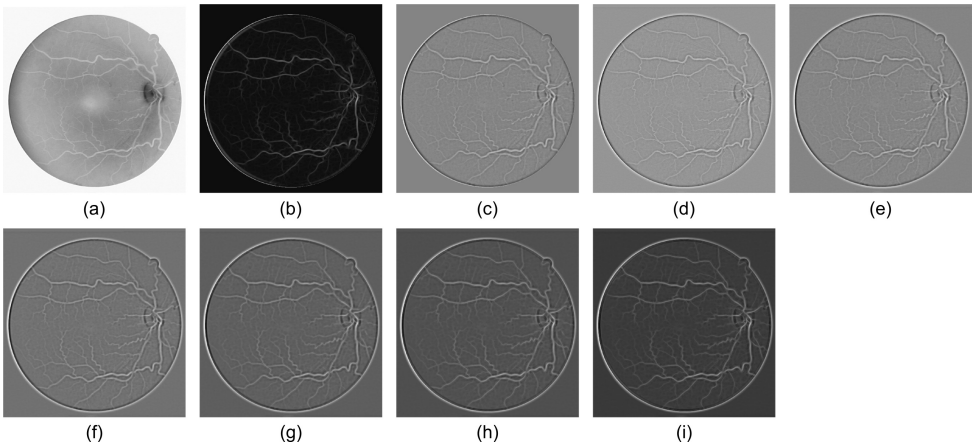

Figure 4: Line detectors. (a) Inverted green band of a retinal image. (b-c) Ricci and Perfetti line detectors (F3.1), (d-i) Nguyen *et al.* line detectors (F3.2).

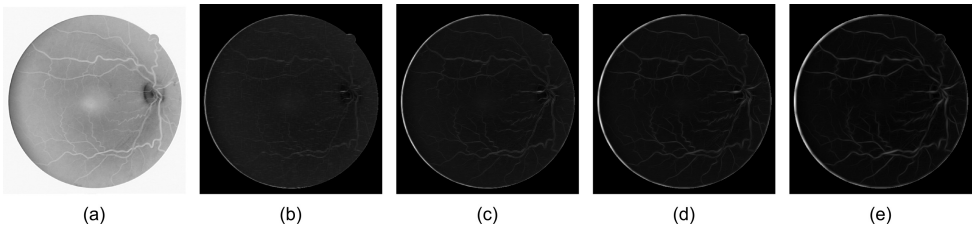

Figure 5: Wavelets. (a) Inverted green band of a retinal image. (b-e) Soares *et al.* Gabor wavelet (F4.1).

2.5 Mathematical morphology

Mathematical morphology is also well adapted for vessel-like patterns detection. By finding the correct combination of operations, it is possible to improve the contrast between vessels and other structures (Zana and Klein, 2001).

Zana and Klein (2001) introduce an unsupervised algorithm that combines morphological filters and cross-curvature evaluation to segment vessel-like patterns. We include the image resulted of those operations as a feature, named **Zana and Klein processed image (F5.1)** (Figure 6). First, a top-hat morphological operation is applied over the green band of the image, using lines of length 9 in a given angle θ. Those angles ranges from 0^o to 165^o, with a step of 15^o. The sum of all the top-hat responses is taken since it reduces the small bright noise and improve the contrast of all linear parts with respect to the original image. Further treatment to remove noisy responses is performed by applying a Laplacian of Gaussian filter using square windows of side 7 with $\sigma = 7/4$. A morphological opening is then applied on the resulting image using lines of length 9 in the angles mentioned above, and the maximum over angles is taking. Afterwards, a morphological closing is performed using the same structuring element, but taking the minimum over angles. Finally, the maximum over angles of an opening operation is applied again using line elements of length 29.

3 PREPROCESSING TECHNIQUES FOR ARTIFACT REDUCTION

Main problems of features extracted from the green band of fundus images are related to the lack of contrast between vessels and background, and artifacts introduced by the border of the FOV (Fraz et al., 2012a). In this section, we study two preprocessing techniques to deal with those problems.

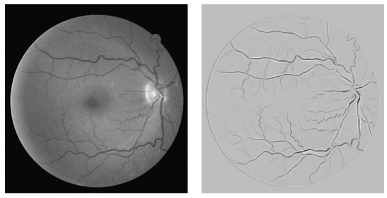

Figure 6: Mathematical morphology. (a) Green band of a retinal image. (b) Zana and Klein processed image (F5.1).

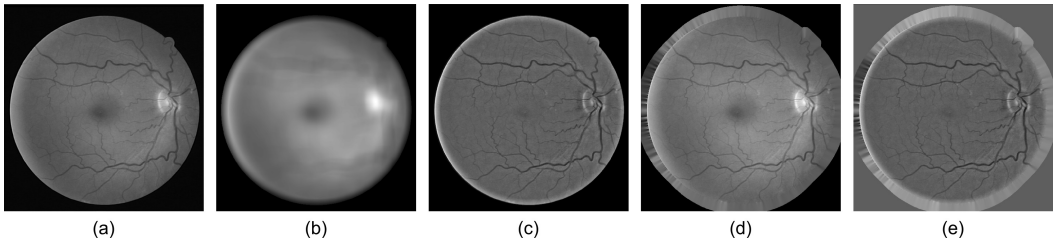

Figure 7: Preprocessing techniques. (a) Original green band of a retinal image. (b) Estimated background. (c) Background subtracted image. (d) Border extended image. (e) Background subtracted and border extended image.

Illumination flaws while capturing the image result in images with biased background. This bias reduce the contrast of vessels in different regions and decrease feature capability of detecting vessels properly (Marín et al., 2011). Several works propose to preprocess images before extracting features to overcome this difficulties. In general, bias reduction methods are based on subtracting to the green band of the original image an estimation of the background, obtained by applying noise filtering with large neighbourhoods. After an extensive, qualitative analysis of bias reduction methods, in this work we propose to estimate the background using mean filtering with 35×35 windows, and considering only internal pixels of the FOV to reduce the effect of possible bias in borders. Figura 7 (b) and (c) depict examples of the estimated background and the background subtracted image. As can be observed, this preprocessing technique reduce the bias and increase the contrast, although introducing high intensities next to the FOV.

Some features analysed in this work require filtering the image. However, filters are influenced by the black points outside the FOV and reduce the contrast of blood vessels, as in the case of Figura 7 (c). Moreover, other features has a high response in the borders of the FOV, resulting in false detection of the border of the camera's aperture. Soares et al. (2006) propose an iterative method to overcome this difficulties which is widely utilized in the state-of-the-art (un par de casos). The preprocessing algorithm starts with a region of interest (ROI) determined by the camera's aperture and iteratively grows this ROI. Each step of the algorithm consists in the following. First, the set of pixels of the exterior border of the ROI is determined, by taking the neighbours to pixels inside the FOV, considering four-connectivity. Then, each pixel value of this set is replaced with the mean value of its neighbours inside the ROI (this time considering eight-connectivity). Finally, the ROI is expanded by including this altered set of pixels. This process is repeated until a certain number of iterations. In this work, we repeat this process 30 times, and we consider square neighbourhoods of size 3 to perform the expansion. Figura 7 (d) shows the result obtained after applying this approach over the original green band of a retinal image. It can be observed that the border is expanded successfully in the right side of the image, although it has dark artifacts in the left side due to the biased background.

In this paper we analyse the effect of this two preprocessing techniques over the previously

mentioned features. We also incorporates a third preprocessing technique, based on the combination of the bias correction method and the algorithm for border expansion (Figura 7 (e)).

4 SEGMENTATION METHOD

As it was explained in Section 1, our supervised method for features evaluation is a linear SVM (Hastie et al., 2009). Let S be a set of N training samples $(\mathbf{x}_i, y_i)$, where $\mathbf{x}_i \in \mathbb{R}^d$ is the feature vector of a given pixel i and $y_i \in \{-1, +1\}$ its corresponding label (in our case, $+1$ represents the blood vessel class, and -1 any other non-vessel class). An SVM identifies the optimal separating hyperplane for which the margin of separation between classes in the d-dimensional feature space is maximized. This is achieved by optimizing the following expression:

$$\min_{\mathbf{w}\in\mathbb{R}^d, \xi_i \in \mathbb{R}^+} \frac{\lambda}{2}\|\mathbf{w}\|^2 + \sum_{i=1}^{N} \xi_i \tag{7}$$

subject to the linear constraints

$$y_i(\mathbf{w}^T\mathbf{x}_i + b) \geq 1 - \xi_i, i = 1, ..., N \tag{8}$$

where $\mathbf{w}$ is the weight vector, b is a bias term which is also learned from training data, each ξ_i is a slack variable, and the parameter $\lambda > 0$ controls the trade-off between the minimization of classification errors on the training set ($0 < \lambda \leq 1$) and the maximization of the margin ($\lambda \gg 1$).

Once the model is learned, the class of a new point $\mathbf{x}$ is given by the sign of the decision function:

$$f(\mathbf{x}) = \mathbf{w}^T\mathbf{x} + b \tag{9}$$

For further details on SVMs, we refer the interested reader to (Hastie et al., 2009).

We base our experiments in VLFeat implementation of the the linear SVM (Vedaldi and Fulkerson, 2010). Due to the large size of the training set (almost 3 million pixels), traditional optimization methods such as gradient descent are not feasible in a reasonable computational time to minimize the objective function. We utilize stochastic gradient descent instead, which converges faster that traditional gradient descent methods, although finding an approximation to the minimum (Spall, 2005).

5 VALIDATION & RESULTS

We performed all our experiments using the publicly available dataset DRIVE (Niemeijer et al., 2004). It comprises a total of 40 manually segmented color fundus images, all obtained from a diabetic retinopathy screening program in the Netherlands. Seven images in the set contains pathology, namely exudates, hemorrhages and pigment epithelium changes. The images were acquired using a Canon CR5 non-mydriatic 3-CCD camera with a 45º FOV, stored in JPEG format, 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. The set is divided into a training set and a test set, both containing 20 images. The test set includes two different segmentations per image, the first one assumed as the gold-standard segmentation and the second one utilized to study the performance of a second human observer.

Several works in the literature report their results in terms of global accuracy measures such as accuracy or area under the receiver operating characteristic (ROC) curve. Since vessel segmentation in retinal images corresponds to a typical skewed-class problem–only around 13% of

image pixels corresponds to blood vessels (Fraz et al., 2012a)–, using these measures to evaluate segmentation is not recommended: due to the small proportion of vessel pixels with respect to the entire image, big changes in the amount of well-classified pixels will not significantly affect global metrics; in the contrary, small changes in the amount of well-classified non-vessel pixels will be reflected in big changes in the performance. Following this prior, we evaluate our results in terms of specificity (Sp), precision (Pr) and recall (Re, also known as sensitivity) (Eq. 10), and also in terms of the F1-score (Eq. 11):

$$Sp = \frac{TN}{TN + FP}, \quad Pr = \frac{TP}{TP + FP}, \quad Re = \frac{TP}{TP + FN} \tag{10}$$

$$F1 = 2 \cdot \frac{Pr \cdot Re}{Pr + Re} \tag{11}$$

where TP is the number of true positives, TN is the amount of true negatives, FN is the number of false negatives and FP is the amount of false positives. TP and TN represent the number of pixels identified as vessel and non-vessel, respectively, in both the ground truth and the segmented images. FN is the amount of pixels classified as non-vessel in the segmentation but as vessel in the ground truth, while FP is the number of pixels classified as vessel in the segmentation but as non-vessel in the ground truth. Specificity measures the ability of the algorithm to detect non-vessel pixels, precision indicates the proportion of identified pixels which are true vessel pixels, and recall the ability to detect vessel pixels. Finally, F1-score is a standard measure for skewed-class classification problems, defined as the harmonic mean of precision and recall.

All the parameters of our method, including λ and the combination of features, were selected following the methodology proposed by Orlando and Blaschko (2014). First, we randomly partitioned DRIVE training set into two new sets, training* and validation, with 15 and 5 images each one, respectively. We utilized training* to train the SVM, and the parameters were selected by optimizing the F1-score over the validation set. Only the final configuration was evaluated using the test set. This methodology allows to avoid overfitting on the test set.

We first evaluate the performance of each preprocessing technique combined with each individual feature. For this purpose, we train the SVM on training*, using a feature extracted from images preprocessed with a given algorithm, and we evaluate the performance using the validation set. Values of λ are selected from the set 10^i, with $i = \{-5, -4, ..., 0, ..., 4, 5\}$, according to the validation set, and looking for the value which optimizes the F1-score measure. Results are summarized in Table 1. As can be seen, only 2 features (F1.2 and F3.1) are not improved by preprocessing images first with the proposed approaches. Applying bias reduction improves the discrimination of blood vessels when using features F2.1, F3.2 and F4.1; extending borders of the field of view upgrade the performance of features F1.4, F1.5 and F2.2; finally, applying both techniques together over the input images improve the F1-scores of features F1.1, F1.3 and F5.1.

In the following experiments, we fix the preprocessing algorithm for each feature according to previous results, and we apply a forward selection approach to select the best combination of features according to the validation set. The best configuration we found includes almost all the features we analysed, except F1.4 y F3.1. Once the model is learned, we evaluate it on the test set. We also evaluate our method using the same set of features but extracted from images without any preprocessing. Results obtained are compared to other state-of-the-art methods in Table 2. After an extensive review of the literature on blood vessel segmentation, we only

Feature	Without preprocessing	Bias reduction	Extended border	Both techniques
F1.1	0.5620	0.5614	0.5609	**0.5621**
F1.2	**0.5483**	0.5214	0.5462	0.5262
F1.3	0.5864	0.5736	0.5758	**0.5917**
F1.4	0.5060	0.4999	**0.5083**	0.5042
F1.5	0.6232	0.6233	**0.6293**	**0.6293**
F2.1	0.5874	**0.5926**	0.5922	0.5838
F2.2	0.4637	0.4608	**0.4744**	0.4508
F2.3	0.2547	0.2547	0.2547	0.2547
F3.1	**0.6497**	0.6488	0.6440	**0.6497**
F3.2	0.6468	**0.6640**	0.6518	0.6470
F4.1	0.5652	**0.5664**	0.5659	0.5646
F5.1	0.6423	0.6422	0.6401	**0.6429**

Table 1: F1-score values obtained on the validation set using features extracted over images preprocessed with different techniques. Bold numbers indicates the maximum F1-score value.

include in the comparison those works that report their results in terms of Pr, Re and Sp. As can be seen, preprocessing images reduces the average values of precision and specificity, but increasing recall and F1-score values.

Method	Pr	Re	Sp	F1-score
Espona et al. (2007)	–	0.6634	0.9682	–
Espona et al. (2008)	–	0.7436	0.9615	–
Fathi and Naghsh-Nilchi (2013)	0.7559	0.7768	0.9759	0.7662
Fraz et al. (2012b)	0.815	0.7152	0.9723	0.7618
Marín et al. (2011)	0.8433	0.7067	0.9801	0.7690
Soares et al. (2006)	–	0.7230	0.9762	–
Staal et al. (2004)	–	0.7190	0.9770	–
Mendonca and Campilho (2006)	–	0.7315	0.9781	–
Wang et al. (2007)	–	0.7800	0.9780	–
You et al. (2011)	–	0.7410	0.9751	–
Human observer	–	0.7761	0.9757	–
Our method without preprocessing	**0.8809**	**0.6321**	**0.9875**	**0.7338**
Our method with preprocessing	**0.8116**	**0.7189**	**0.9757**	**0.7596**

Table 2: Comparison of Pr, Re, Sp and F1-score of our method with and without preprocessing, with respect to the existing blood vessel segmentation algorithm, based on DRIVE data set.

Qualitative results are depicted in Figure 8. In general, the SVM trained with features extracted from non-preprocessed images result in good segmentations of wide vessels. However, narrow vessels are almost ignored by the method. Preprocessing the images improves segmentation of wide vessels, filling holes between their borders (Figure 9 (I)). Additionally, this approach incorporates several pixels of narrow vessels (Figure 9 (II)). Noisy pixels in the border of the FOV are also included into the segmentation (Figure 9 (III)), although there are not significant with respect of the large amount of well classified pixels incorporated.

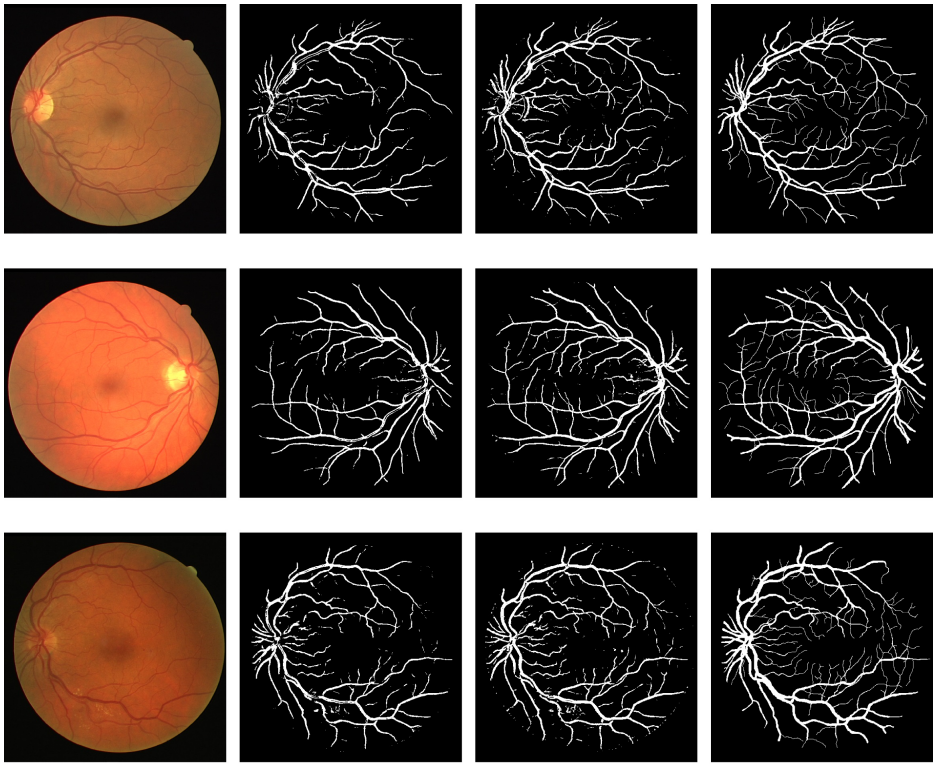

Figure 8: Examples of results obtained on the test set. (a) Original image. (b) Our method without preprocessing the images. (c) Our method preprocessing the images. (d) Ground truth segmentation.

6 CONCLUSIONS

In this paper we have presented an evaluation of different preprocessing and feature extraction techniques in the context of supervised segmentation of retinal blood vessels in fundus images. By taking a linear SVM as the segmentation method, we have reviewed the behaviour of several features when images are preprocessed using two state-of-the-art algorithms. Parameters of the method were fixed by considering a validation set of retinal images extracted from the original training set of DRIVE, a publicly available data set for the evaluation of blood vessel segmentation algorithms. Best combination of features was also selected according to the validation set, following a forward selection approach. We have showed that performance is considerably increased in terms of the F1-score when the best preprocessing algorithm is fixed for each individual feature. We have also demonstrated that applying a preprocessing algorithm before extracting each feature allows to achieve better results when segmentations are analysed qualitatively. In particular, this change incorporates a large amount of pixels corresponding to narrow vessels and filling holes in wide vessels. Finally, results were compared with respect to other state-of-the-art approaches, showing that this simple method achieves similar quality values than other available algorithms. Thus, combining our approach with more sophisticated segmentation techniques could reach better results. The entire feature evaluation methodology applied in this work could be implemented to fix features for more sophisticated supervised approaches.

REFERENCES

Abe S. *Support vector machines for pattern classification*. Springer, 2010.

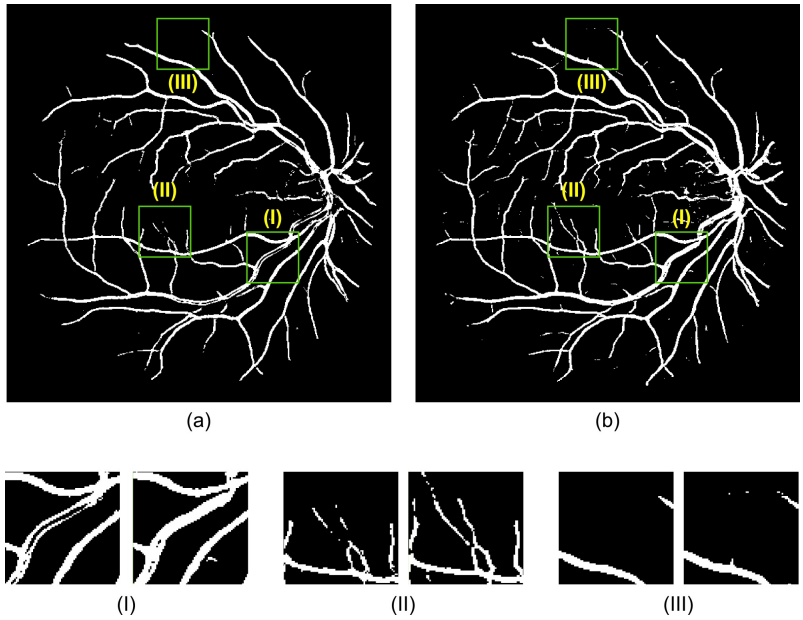

Figure 9: Results obtained using features extracted over images (a) without preprocessing and (b) with preprocessing. (I) Improvements on the detection of wide vessels. (II) Improvements on the detection of narrow vessels. (III) Noisy artifacts in the border of the FOV.

Abràmoff M.D., Garvin M.K., and Sonka M. Retinal imaging and image analysis. *Biomedical Engineering, IEEE Reviews in*, 3:169–208, 2010.

Becker C., Rigamonti R., Lepetit V., and Fua P. Supervised feature learning for curvilinear structure segmentation. In K. Mori, I. Sakuma, Y. Sato, C. Barillot, and N. Navab, editors, *MICCAI 2013, LNCS*, volume 8149, pages 526–533. Springer, 2013.

Chaudhuri S., Chatterjee S., Katz N., Nelson M., and Goldbaum M. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE T-MI*, 8(3):263–269, 1989.

Espona L., Carreira M.J., Ortega M., and Penedo M.G. A snake for retinal vessel segmentation. In J. Martí, J.M. Benedí, A.M. Mendonca, and J. Serrat, editors, *Pattern Recognition and Image Analysis, LNCS*, volume 4478, pages 178–185. Springer, 2007.

Espona L., Carreira M.J., Penedo M., and Ortega M. Retinal vessel tree segmentation using a deformable contour model. In *ICPR*. 2008.

Fathi A. and Naghsh-Nilchi A.R. Automatic wavelet-based retinal blood vessels segmentation and vessel diameter estimation. *Biomedical Signal Processing and Control*, 8(1):71–80, 2013.

Frangi A.F., Niessen W.J., Vincken K.L., and Viergever M.A. Multiscale vessel enhancement filtering. In W.M. Wells, A. Colchester, and S. Delp, editors, *MICCAI 1998, LNCS*, volume 1496, pages 130–137. 1998.

Fraz M.M., Remagnino P., Hoppe A., Uyyanonvara B., Rudnicka A.R., Owen C.G., and Barman S.A. Blood vessel segmentation methodologies in retinal images–a survey. *Computer methods and programs in biomedicine*, 108(1):407–433, 2012a.

Fraz M.M., Remagnino P., Hoppe A., Uyyanonvara B., Rudnicka A.R., Owen C.G., and Barman S.A. Ensemble classification system applied for retinal vessel segmentation on child images containing various vessel profiles. In *Image Analysis and Recognition*. 2012b.

Hastie T., Tibshirani R., Friedman J., Hastie T., Friedman J., and Tibshirani R. *The elements of statistical learning*, volume 2. Springer, 2009.

Lupascu C.A., Tegolo D., and Trucco E. Fabc: retinal vessel segmentation using adaboost.

IEEE T-ITB, 14(5):1267–1274, 2010.

Marín D., Aquino A., Gegúndez-Arias M.E., and Bravo J.M. A new supervised method for blood vessel segmentation in retinal images by using gray-level and moment invariants-based features. *IEEE T-MI*, 30(1):146–158, 2011.

Martinez-Perez M.E., Hughes A.D., Thom S.A., Bharath A.A., and Parker K.H. Segmentation of blood vessels from red-free and fluorescein retinal images. *Medical Image Analysis*, 11(1):47–61, 2007.

Mendonca A.M. and Campilho A. Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction. *IEEE T-MI*, 25(9), 2006.

Nguyen U.T., Bhuiyan A., Park L.A., and Ramamohanarao K. An effective retinal blood vessel segmentation method using multi-scale line detection. *Pattern Recognition*, 2012.

Niemeijer M., Staal J., van Ginneken B., Loog M., and Abramoff M.D. Comparative study of retinal vessel segmentation methods on a new publicly available database. In *Medical Imaging 2004*, pages 648–656. International Society for Optics and Photonics, 2004.

Orlando J.I. and Blaschko M. Learning fully-connected CRFs for blood vessel segmentation in retinal images. In *MICCAI 2014, LNCS*. Springer, 2014.

Osareh A. and Shadgar B. Automatic blood vessel segmentation in color images of retina. *Iran. J. Sci. Technol. Trans. B: Engineering*, 33(B2):191–206, 2009.

Ricci E. and Perfetti R. Retinal blood vessel segmentation using line operators and support vector classification. *IEEE T-MI*, 26(10):1357–1365, 2007.

Salem S.A., Salem N.M., and Nandi A.K. Segmentation of retinal blood vessels using a novel clustering algorithm (racal) with a partial supervision strategy. *Medical & biological engineering & computing*, 45(3):261–273, 2007.

Sinthanayothin C., Boyce J.F., Cook H.L., and Williamson T.H. Automated localisation of the optic disc, fovea, and retinal blood vessels from digital colour fundus images. *British Journal of Ophthalmology*, 83(8):902–910, 1999.

Soares J.V., Leandro J.J., Cesar R.M., Jelinek H.F., and Cree M.J. Retinal vessel segmentation using the 2-d Gabor wavelet and supervised classification. *IEEE T-MI*, 25(9), 2006.

Spall J.C. *Introduction to stochastic search and optimization: estimation, simulation, and control*, volume 65. John Wiley & Sons, 2005.

Staal J., Abràmoff M.D., Niemeijer M., Viergever M.A., and van Ginneken B. Ridge based vessel segmentation in color images of the retina. *IEEE T-MI*, 23(4):501–509, 2004.

Vedaldi A. and Fulkerson B. Vlfeat: An open and portable library of computer vision algorithms. In *Proceedings of the international conference on Multimedia*, pages 1469–1472. ACM, 2010.

Wang L., Bhalerao A., and Wilson R. Analysis of retinal vasculature using a multiresolution hermite model. *Medical Imaging, IEEE Transactions on*, 26(2):137–152, 2007.

You X., Peng Q., Yuan Y., Cheung Y.m., and Lei J. Segmentation of retinal blood vessels using the radial projection and semi-supervised approach. *Pattern Recognition*, 44(10), 2011.

Zana F. and Klein J.C. Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE TIP*, 10(7):1010–1019, 2001.

Zwiggelaar R., Astley S.M., Boggis C.R., and Taylor C.J. Linear structures in mammographic images: detection and classification. *Medical Imaging, IEEE Transactions on*, 23(9):1077–1086, 2004.