

Efficient automatic segmentation of vessels

Adriana Romero¹
aromero@cvc.uab.es

Simeon Petkov²
spetkov@cvc.uab.es

Carlo Gatta¹
carlo.gatta@ub.edu

Dr. Manel Sabaté³
masabate@clinic.uab.es

Petia Radeva¹
petia.ivanova@ub.edu

¹ Universitat de Barcelona
Barcelona
Spain

² Centre de Visió per Computador
Bellaterra
Spain

³ Institut Clinic del Torax
Hospital Clinic Barcelona
Spain

Abstract

The segmentation of tubular structures is still an open field of investigation, particularly in medical imaging, where the quality of the image is poor with respect to natural images. Despite the quality of state-of-the-art segmentation methods, little effort has been devoted to the computational efficiency of the algorithms. Efficiency is an important topic, since intra-operative computer assisted interventions require near real-time performance. In this paper, we present a simple, yet effective, algorithm that efficiently segments vessels in 2D images. The algorithm requires no initialization and has a computational cost of $\mathcal{O}(SN \log N)$, where S is the number of scales and N is the number of image pixels. Results on the DRIVE dataset show that the proposed method has near state-of-the-art performance with very little computational burden.

1 INTRODUCTION

There is a vast literature in the area of vessel analysis and segmentation. A detailed review of existing vessel extraction techniques is given in [5]. Among the existing techniques, machine learning-based techniques may use knowledge about appearance of blood vessels in the segmentation process [11]. Hessian-based techniques analyze the eigenvalues of the Hessian matrix to perform the vessel segmentation [3]. Other techniques use medialness functions [6] as vessel detection filters. The great majority of these methods have high computational costs and/or require manual initialization of one or more seed points. We propose an efficient algorithm which is based on the following ideas: (1) detect blobs, tubular structures and edges through linear filtering, without having to compute the Hessian matrix for all pixels/voxels; (2) select a number of seeds based on the novel EdgeLoG measure and reject blob-like points by means of eigenvalues analysis, so that we limit the Hessian and eigenvalues computation only to potential seed points; (3) use a region growing approach in which edge information is used to stop the propagation. We validate our method on the




Figure 1: Edgeness, Laplacian of Gaussian and EdgeLoG responses of (a) a 1-D edge signal and (b) a 1-D tubular signal.

DRIVE dataset¹, obtaining a kappa value of 0.6991, which compares well with much more sophisticated state-of-the-art methods.

2 METHOD

The method combines an edge detection operator with a blob detection operator and exploits their properties to obtain a negative signal in the presence of vessels. The approach computes the Edgeness and the Laplacian of Gaussian of an image. Edgeness obtains a positive response in the presence of edges and the Laplacian of Gaussian obtains a positive/negative response where there is contrast and 0 in constant regions. The combination of both signals is able to distinguish between tubular structures and edges and provides a negative response with a local minimum in the middle of the vessels.

Fig. 1 outlines the basic idea of the approach on a 1D toy problem. Fig. 1(a) shows the response of the operators to an edge signal. The response of the combined operator is always positive. Fig. 1(b) shows the response of the operators to a tubular-like signal. In this case, the response of the combined operators has a negative local minimum in the middle of the tubular structure. Then, we select reliable local minima in the middle of tubular structures and use them as seeds in a region growing process. Since the Laplacian of Gaussian signal can detect blobs, we reject all seeds that represent blob-like structures.

2.1 EdgeLoG computation

Let $I(x,y)$ be a grayscale image of any size. The scale space representation of $I(x,y)$ is defined as the image convolved with a family of Gaussian kernels [7]: $L^{(1)}(x,y;\sigma_i) = G(x,y;\sigma_i) * I(x,y)$ where $G(x,y;\sigma_i)$ is a Gaussian kernel of standard deviation σ_i . We also define a k-modified scale space: $L^{(k)}(x,y;k\sigma_i) = G(x,y;k\sigma_i) * I(x,y)$ where k is a constant and $G(x,y;k\sigma_i)$ is a Gaussian kernel of standard deviation $k\sigma_i$. We define the Edgeness signal as $E(x,y) = \max_{\sigma_i} \|\nabla L^{(\cdot)}(x,y;\sigma_i)\|$ where $L_x^{(\cdot)}(x,y;\sigma_i)$ and $L_y^{(\cdot)}(x,y;\sigma_i)$ are the normalized first derivatives of $L^{(\cdot)}(x,y;\sigma_i)$ as in [7]. The Laplacian of Gaussian is $\nabla^2 L^{(\cdot)}(x,y;\sigma_i) = L_{xx}^{(\cdot)}(x,y;\sigma_i) + L_{yy}^{(\cdot)}(x,y;\sigma_i)$ where $L_{xx}^{(\cdot)}(x,y;\sigma_i)$ and $L_{yy}^{(\cdot)}(x,y;\sigma_i)$ are the normalized scale-space second derivatives [7].

¹<http://www.isi.uu.nl/Research/Databases/DRIVE/>

We propose a new measure, named EdgeLoG which conveniently combines the Laplacian of Gaussian and the Edgessness signals.

$$EL(x,y) = \min_{\sigma_i} (\nabla^2 L^{(1)}(x,y; \sigma_i) + \|\nabla L^{(k)}(x,y; \sigma_i)\|) \quad (1)$$

The rationale behind this new operator is the following: (1) the Laplacian of Gaussian has a strong negative response in the presence of bright blobs, tubular structures and edges; (2) Edgessness obtains a positive response in the presence of edges and is 0 in the middle of vessels; (3) we remove the negative response of the Laplacian of Gaussian at edges while maintaining negative values in the middle of the vessels, using the Edgessness response (see e.g. Fig. 1). However, using the same standard deviation for both detectors has proven not to work properly. Therefore we introduce a constant k which modifies the Edgessness response and compensates responses to edges and tubular structures. The constant k can be computed as the minimum value that provides a negative response in ridges and a strictly positive response at edges.

To allow the detection of different vessel sizes, we compute the response of the EdgeLoG operator at multiple scales and we select the minimum EdgeLoG response over scales.




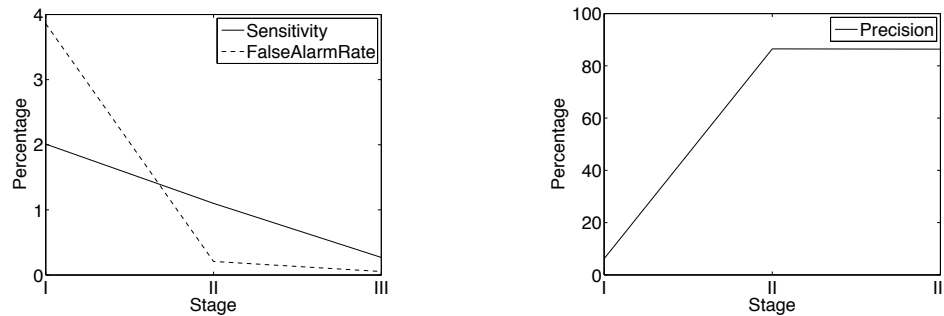
Figure 2: 2D Segmentation of retinal vessels.

2.2 Seed selection/pruning

STAGE I: Once the EdgeLoG signal is obtained, we search for local minima to define the initial set of seeds.

STAGE II: Not all local minima are reliable seeds. To reduce the number of seeds, we compute an image dependent threshold $\varepsilon < 0$ using the Otsu method [10] on the EdgeLoG intensities of the previously computed seeds, pruning seeds with EdgeLoG $> \varepsilon$. This process already rejects a consistent quantity of seeds (96.07 % on average).

STAGE III: Some of these minima may belong to blob-like structures. To identify blob-like seeds, we compute the Hessian matrix of the seed locations at the respective dominant scale, and compute their eigenvalues $|\lambda_1| \leq |\lambda_2|$. The ratio $r = |\lambda_1|/|\lambda_2| \in [0, 1]$ is a nice indicator of the blob-ness and saddle-ness of the structure. We want to reject as many non-vessel structures as possible. Therefore, we remove seed points having $r > 0.5$. In this way, we limit the possibility of using some false positives as seeds, which could induce severe errors in the region growing process. Nonetheless, removing a large amount of true positives is not a problem to our algorithm since the proposed region growing approach can produce satisfactory results with just one seed per vessel branch (see Fig. 2(d)). Fig. 3(a) and Fig. 3(b) present respectively the sensitivity and false alarm rate and the precision of the seed points in three different stages. Note that the sensitivity and the false alarm rate decrease as we remove potentially erroneous seeds while the precision tends to increase.



(a) Variations in sensitivity & false alarm rate (b) Variations in precision
 Figure 3: Seed pruning.

2.3 Region growing

To segment the vessel regions using the selected seeds, we employ a classical region growing method. The region growing process is driven by two hypotheses: (1) the region should grow in the direction of vessels to include areas where no seeds are present and (2) the region should grow in the orthogonal direction until we reach the edges of vessels to provide an accurate segmentation. To achieve these aims, we use the EdgeLoG and the Edgeness operators. Being (x_p, y_p) a seed or a pixel that belongs to the region growing frontier, we analyze its 8-neighborhood to check if to propagate the frontier or not. Being (x_j, y_j) one neighbor of (x_p, y_p) that has not been analyzed previously, we set the pixel as the new frontier if the following constraint is satisfied:

$$EL(x_j, y_j) < t_{EL} \cup \frac{E(x_j, y_j) - E(x_p, y_p)}{\|(x_p, y_p) - (x_j, y_j)\|} \geq t_E \quad (2)$$

where $t_{EL} \leq 0$ and $t_E \geq 0$ are two algorithm parameters. Equation (2) makes the algorithm grow through the negative valleys of the EdgeLoG signal (see Fig.2(b)) and allows the algorithm to “climb” the Edgeness map, until the ridges are reached (see Fig. 2(c)). A nice property of the proposed method is that the theory behind the EdgeLoG and Edgeness maps suggests a theoretical value for the parameters: $t_{EL} = 0$ and $t_E = 0$. Nonetheless, due to the discrete domain of the image and noise, we set $t_{EL} = -2.5 \cdot 10^{-3}$ and $t_E = 2.5 \cdot 10^{-3}$.

3 VALIDATION

The method has been tested on the DRIVE database. The database contains 40 retinal images divided into a training set and a test set, containing 20 images each. Since our algorithm does not require a training step, we limit our evaluation to the test set. In all the results shown in this paper, the set of scales is $\Sigma = \{0.8, 1.75, 3.5\}$. For the given scales, a value of $k = 1.256$ has been obtained. Table 1 shows the results of our method compared to state-of-the-art algorithms, as reported in the DRIVE website². Our method performs well compared to the state-of-the-art while being less computational expensive and much less complicated than other algorithms.

²<http://www.isi.uu.nl/Research/Databases/DRIVE/results.php>

Table 1: Performance of vessel segmentation methods on the DRIVE database.

Method	Accuracy	Kappa
Human observer	0.9473 (0.0048)	0.7589
Staal [12]	0.9442 (0.0065)	0.7345
Niemeijer [9]	0.9416 (0.0065)	0.7145
<i>Our method</i>	<i>0.9345 (0.0060)</i>	<i>0.6991</i>
Zana [13]	0.9377 (0.0077)	0.6971
Al-Diri [1]	0.9258 (0.0126)	0.6716
Jiang [4]	0.9212 (0.0076)	0.6399
Martínez-Pérez [8]	0.9181 (0.0240)	0.6389
Chaudhuri [2]	0.8773 (0.0232)	0.3357
All background	0.8727 (0.0123)	0

4 DISCUSSION

The proposed algorithm aims at performing a reasonably good segmentation while limiting the computational cost to the very minimum. The use of linear filtering at different scales allows a quick detection of a certain number of candidate seeds, without the need of computing the Hessian and its eigenvalues for all the pixels in the image. While this could be a minor gain in the algorithm speed for 2D images, it is a huge advantage for 3D data, where the computation of the Hessian and its eigenvalues is extremely costly. To properly deal with possible false positives in the first part of the detection, we employ a Hessian analysis *only on previously selected seeds*; in the proposed experiments the selected seeds are 0.16% of the total image. Moreover, the dominant scale can be easily computed from the EdgeLoG operator, then reducing the computation of eigenvalues to only one scale per seed pixel. The region growing is very efficient and uses a very limited local information, namely the EdgeLoG map and the local gradient of the Edginess map. The computational cost is mainly dominated by the Gaussian filtering, so that the computational cost is $\mathcal{O}(SN \log N)$, where S is the number of scales and N is the number of image pixels. It is worth mentioning that this part of the algorithm can be easily parallelized.

The method has several interesting properties: (1) the use of the LoG allows having negative EdgeLoG values in presence of bifurcations, so that the region growing can actually fill bifurcation areas; (2) differently than methods that use the eigenvalues of the Hessian matrix, the method does not produce less accurate results in vessels presenting high curvature; (3) finally, thanks to the Edginess signal, strong edges are suppressed while other methods tend to identify straight edges in 2D images as vessels.

Nonetheless, the method has some limitations to be approached in future research: (1) the use of the Edginess signal, especially at coarser scales, could potentially mask important information at finer scales, causing e.g. poor segmentation of orthogonal vessels branches whose contrast is much lower than the main vessel; (2) the implemented region growing has no sub-pixel accuracy and very little vessels can be lost in the region growing process due to EdgeLoG and Edginess discrete lattice.

Finally, a discussion should be devoted to the similarity and differences with the method presented in [8]. Our method does not need the computation of eigenvalues of Hessian matrix. The proposal of the EdgeLoG signal, with appropriate k parameter is totally novel. Despite both methods use a region growing approach, we have theoretical basis on parameter selection, and our method requires much less computational cost: we do not perform region growing of the “background” class.

5 CONCLUSION AND FUTURE WORKS

In this paper we presented an efficient method for segmentation of vessels in medical images. Despite its low computational cost, it demonstrated to perform as good as state-of-the-art algorithms on the DRIVE dataset. Future works encompass: (1) the use of a classifier to learn the constraints for the region growing propagation as a function of local image features; (2) implementation of a 3D version of the method and subsequent testing on CT data.

References

- [1] B. Al-Dir, A. Hunter, and D. Steel. An active contour model for segmenting and measuring retinal vessels. *IEEE TMI*, 28(9):1488–97, 2009.
- [2] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE TMI*, 8(3):263–269, 1989.
- [3] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *Proc. of MICCAI*, volume 1496, pages 130–137, 1998.
- [4] X. Jiang and D. Mojon. Adaptive local thresholding by verification-based multithreshold probing with application to vessel detection in retinal images. *IEEE TPAMI*, 25(1):131–137, 2003.
- [5] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM CSUR*, 36(2):81–121, 2004.
- [6] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Troussel. Model-based detection of tubular structures in 3d images. *CVIU*, 80(2):130 – 171, 2000.
- [7] T. Lindeberg. Principle for automatic scale selection. Technical report, Royal Institute of Technology, 1998.
- [8] M. Martínez-Pérez, A. Hughes, A. Stanton, S. Thom, A. Bharath, and K. Parker. Retinal blood vessel segmentation by means of scale-space analysis and region growing. In *MICCAI*, volume 1679, pages 90–97, 1999.
- [9] M. Niemeijer, J. J. Staal, B. van Ginneken, M. Loog, and M. D. Abramoff. Comparative study of retinal vessel segmentation methods on a new publicly available database. In *SPIE MI*, volume 5370, pages 648–656, 2004.
- [10] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE TSMC*, 9(1):62–66, 1979.
- [11] C. Smets, G. Verbeek, P. Suetens, and A. Oosterlinck. A knowledge-based system for the delineation of blood vessels on subtraction angiograms. *PRL*, 8(2):113 – 121, 1988.
- [12] J. J. Staal, M. D. Abramoff, M. Niemeijer, M. A. Viergever, and B. van Ginneken. Ridge based vessel segmentation in color images of the retina. *IEEE TMI*, 23:501–509, 2004.
- [13] F. Zana and J. Klein. Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE TIP*, 10(7):1010–19, 2001.