A Fully Automatic Method For Segmenting Retinal Artery Walls in Adaptive Optics Images

Nicolas Lermé, Florence Rossant, Isabelle Bloch, Michel Paques, Edouard Koch, Jonathan Benesty

HAL Id: hal-02559191
https://hal.science/hal-02559191

Submitted on 18 Jun 2020

A Fully Automatic Method For Segmenting Retinal Artery Walls in Adaptive Optics Images

Nicolas Lermé[a,b,**], Florence Rossant[a], Isabelle Bloch[b], Michel Paques[c], Edouard Koch[d], Jonathan Benesty[c]

[a]Institut Supérieur d'Électronique de Paris, Paris, France
[b]Institut Mines-Télécom, Télécom ParisTech, CNRS LTCI, Paris, France
[c]CIC 503, Centre Hospitalier National d'Ophtalmologie des Quinze-Vingts, Paris, France
[d]Centre Hospitalier de Versailles, Versailles, France

ABSTRACT

Adaptive optics imaging of the retina has recently proven its capability to image micrometric structures such as blood vessels, involved in common ocular diseases. In this paper, we propose an approach for automatically segmenting the walls of retinal arteries in the images acquired with this technology. The walls are modeled as four curves approximately parallel to a previously detected reference line located near the vessel center (axial reflection). These curves are first initialized using a tracking procedure and then more accurately positioned using an active contour model embedding a parallelism constraint. We consider both healthy and pathological subjects in the same framework and show that the proposed method applies in all cases. Extensive experiments are also proposed, by analyzing the robustness of the axial reflections detection, the influence of the tracking parameters as well as the performance of the tracking and the active contour model. Noticeably, the results show a good robustness for detecting axial reflections and a moderate influence of the tracking parameters. Compared to a naive initialization, the active contour model coupled with the tracking also offers faster convergence and better accuracy while keeping an overall error smaller or very near the inter-physicians error.

1. Introduction

This paper deals with the segmentation and the quantification of retinal blood vessels in Adaptive Optics (AO) images. This recent and non-invasive technique provides a new insight on retinal vessels and their diseases. In comparison to classical eye fundus images, AO images have a better lateral resolution (Rossant et al., 2011) and allow us to visualize microstructures such as photoreceptors (Liang et al., 1997), capillaries (Martin and Roorda, 2005; Paques et al., 2014b) and vascular walls (Chui et al., 2012). This technique offers a new diagnosis and prognosis investigation tool to study the diseases affecting the retinal blood vessels of small diameter ($\leq 150\mu$m), which are major causes of morbidity and mortality, such as Hypertensive Retinopathy (HR) and Diabetic Retinopathy (DR). Early treatment of these diseases is crucial to avoid visual loss.

This requires objective and accurate quantification of vessel features, such as wall morphometry, which can be derived from an automated segmentation of AO images. The measure of wall thickness appears to be of great importance for physicians. The automatic segmentation of AO images is the topic of this paper.

Data and challenges. The images used in this study were acquired with a rtx1 camera (Viard et al., 2011) with flood illumination at 10Hz using a 850nm LED light source with a pixel resolution of 1.33μm. Flood-illumination systems usually produce noisy images making walls hardly visible. A common solution is to geometrically align a stack of images acquired in a short time (4.2s, 9.5 frames/s) and average them to increase the signal-to-noise ratio (Kulcsár et al., 2013). In these images, blood vessels appear as dark elongated structures with a bright linear axial reflection, over a textured background. These characteristics will be exploited in the proposed method. Parietal structures (arterial walls) appear as gray lines along both sides of the lumen (blood column), with a typical thickness of about 15% of the latter (Koch et al., 2014) (see Figure 1).

**Corresponding author: Tel.: +33(0)1.64.69.48.00.
e-mail: nicolas.lerme@mines-paristech.fr (Nicolas Lermé)

(a) 924×947 pixels   (b) 218×218 pixels   (c) 218×218 pixels   (d) 900×900 pixels

Figure 1: Examples of images acquired by the rtx1 camera and a detailed view of them for an healthy subject (a,b) and a pathological one (c,d) (Lermé et al., 2014a).


Figure 2: Flow diagram of the proposed approach.

Segmenting arterial walls in these images is however very challenging for multiple reasons: (i) the background is highly textured, (ii) the lumens are globally dark but with significant intensity variation along them, (iii) the axial reflections may locally show discontinuities or poor contrast, (iv) the outer borders of walls are poorly contrasted, (v) the vessel segments can be locally blurred due to the geometry of the retina, and (vi) morphological deformations can occur in case of pathologies.

Method overview. In this paper, we propose a fully automatic method for segmenting arterial walls in a selected region of interest[1] in the averaged images produced by (Kulcsár et al., 2013). Figure 2 outlines the different steps of the proposed approach. To the best of our knowledge, this is the first method addressing this problem in such images. To overcome the difficulties previously outlined, we propose a strategy exploiting both geometric, radiometric and topological prior information regarding vessels. More precisely, we model arterial walls as four curves approximately parallel to a common reference line located near the vessel center (axial reflection). Once this line is detected, the curves are simultaneously initialized as close as possible to the borders of walls using a tracking procedure to cope with morphological deformations along vessels (pre-segmentation). Then, these curves are more accurately positioned using a parallel active contour model where each curve evolves independently of the others towards large image gradients under a parallelism constraint (Ghorbel et al., 2011). This approach allows us to control the distance of the curves to their reference line, without knowing it accurately as prerequisite. This work has also permitted the physicians to establish relationships between morphometric measurements and clinical parameters (Koch et al., 2014).

This paper is an extension of our previous publication (Lermé et al., 2014b), considering still healthy subjects but also pathological ones in the same framework. Additionally, we provide more details on the method where some steps (such as the axial reflection detection) have been improved. The experiments and the evaluation have also been substantially expanded with the analysis of the robustness for detecting axial reflections, the influence of the tracking parameters as well as the performance of the tracking and the active contour model. This work is complementary to our previous publication (Lermé et al., 2014a) where curves are linked to each others to improve the robustness of the model.

The rest of this paper is organized as follows. In Section 2, we detail the steps for detecting axial reflections inside vessels. Next, we introduce in Section 3 the pre-segmentation and the active contour model for segmenting arterial walls. Finally, we evaluate the performance of the method against manual segmentations in Section 4 and discuss perspectives in Section 5.

2. Axial reflection detection

All along this section, we consider 2D images as functions mapping pixels from $\Omega \subset Z^2$ into the interval $[0, 1]$.

2.1. Pre-processing

The original image (see Figure 3(a)) is first pre-processed by applying a median filter with a square structuring element of size 5 followed by a morphological closing with a circular structuring element of radius 3, in order to enhance the axial reflection. We denote the resulting image by I_{P_1} (see Figure 3(b)).

The source image is also denoised by a non-linear diffusion filter (Weickert et al., 1998) with the contrast parameter and the space regularization parameter respectively set to 0.2 and 2.0. We denote by I_{P_2} the resulting image (see Figure 3(c)). This filter allows us smoothing the vessel lumen while preserving the contrast along its edges.

2.2. Detection of bright elongated structures

Two filters are sequentially applied on the pre-processed image I_{P_1} in order to further enhance the bright elongated structures. The first one is a white top-hat with a binary mask whose radius has a fixed size of 1/3 of the axial reflection diameter [2].

---

[1]In particular, we do not aim at segmenting the whole vascular tree.

[2]Despite the variety of images, this parameter appears to be stable in our experiments (including those presented in Section 4).

(a) Source image   (b) I_{P_1} (c) I_{P_2} (d) I_{T_1} (e) I_{T_2}

(f) $I_{ES}^{(0)}$ (g) $I_{ES}^{(1)}$ (h) I_{ES} (i) k-mean   (j) I_{DA}

(k) $I_{ES'}$ (l) $I_{ES''}$ (m) I_{LM} (n) Labeled axial segments   (o) Labeled lumens
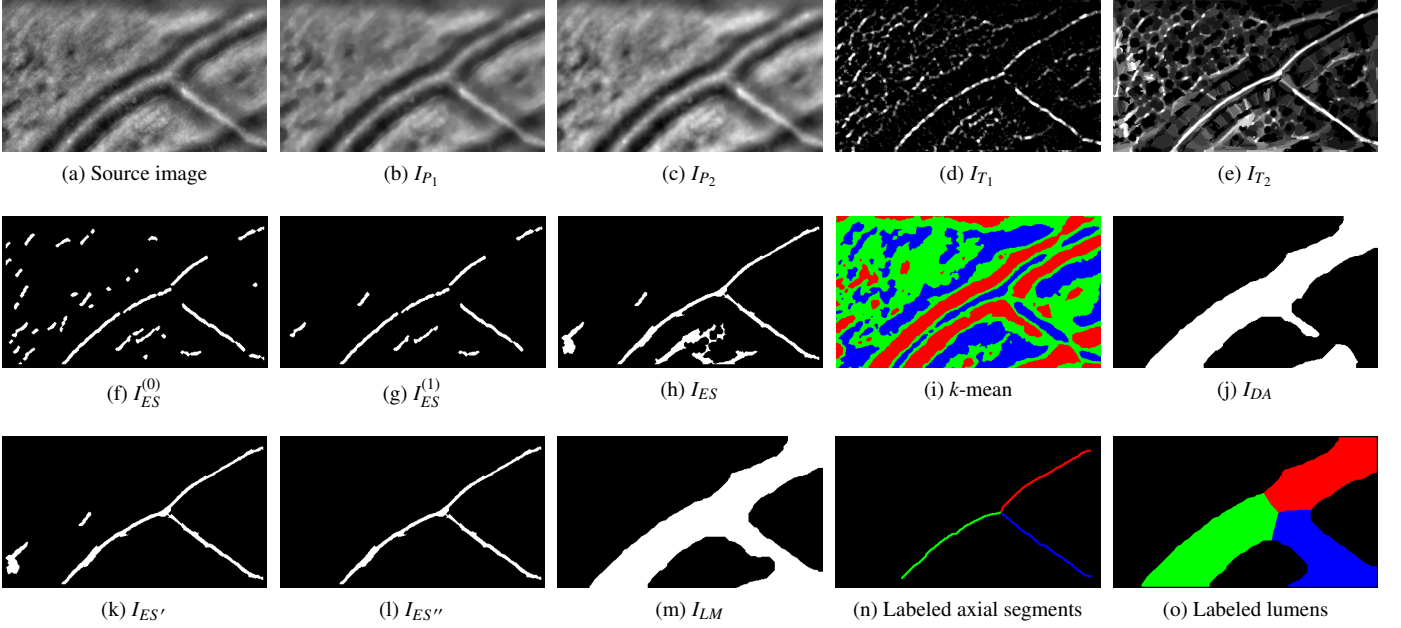
Figure 3: Steps for axial reflection detection: The source image (a) is pre-processed in two different ways (b,c). Bright elongated structures are detected from (b) by top-hat (d) and adapted linear filters (e). From (d), these structures are extracted by hysteresis thresholding (f), filtered (g) and finally combined with those detected in (e) to give (h). Darkest areas are detected from (c) by k-means (i) followed by post-processing (j). From (h) and (j), the most relevant axial segments are kept (k,l) and combined with (j) to give (m). Final axial segments (i) and lumens (j) are labeled from (d), (l) and the pruned skeleton of (m).

We denote by I_{T_1} the top-hat image (see Figure 3(d)). The second one is a series of adapted linear filters designed to estimate the local direction of white linear structures. The mean grey-level is calculated at every pixel $(i, j) \in \Omega$ along segments of fixed length but with different orientations, centered on it. The segment length is about 55 pixels ($\simeq 73\mu m$) and the orientation step is equal to $\Delta\theta = 5°$ ($N = 36$ filters). Let us denote by $I_{LF}^{(k)}$ the image output by the filter with orientation $k\Delta\theta$, $k \in \{0, \ldots, N-1\}$, and by I_D the image storing the estimated direction:

$$k_{opt}(i, j) = \operatorname*{argmax}_{k} I_{LF}^{(k)}(i, j),$$

$$I_D(i, j) = \begin{cases} k_{opt}(i, j) & \text{if } 0.75 I_{LF}^{(k_{opt}(i,j))(i,j)} > I_{LF}^{((k_{opt}(i,j)+\frac{N}{2}) \bmod N)}(i, j) \\ -1 & \text{otherwise,} \end{cases}$$

where $x \bmod y$ is the remainder of x divided by y. In the latter equation, the threshold 0.75 enables us to distinguish the pixels within linear bright features with well-defined local direction from all the others. From the top-hat image I_{T_1}, we compute a second image I_{T_2}, defined for any pixel $(i, j) \in \Omega$ as follows:

$$I_{T_2}(i, j) = \begin{cases} (I_{T_1} \bullet S l_{I_D(i,j)})(i, j) & \text{if } I_D(i, j) \in \{0, \ldots, N-1\} \\ I_{T_1}(i, j) & \text{otherwise,} \end{cases}$$

where $\bullet$ denotes the morphological closing operator and $S l_k$ is a binary linear structuring element whose length and orientation are respectively set to 21 pixels ($\simeq 28\mu m$) and $k\Delta\theta$ degrees. Compared to the image I_{T_1}, axial reflections show better properties in terms of continuity in the image I_{T_2} (see Figure 3(e)) but the lumens are more noisy.

Afterwards, we binarize the image I_{T_1} using hysteresis thresholding (lower and upper thresholds are set to 0.2 and 0.6,

respectively) and denote the resulting image by $I_{ES}^{(0)}$ (see Figure 3(f)). Parts of the axial reflection of vessels are thus extracted, but also other bright areas of the textured background. Further processing steps are therefore needed to discard the undesired areas. For doing so, we propose to keep the 10 largest connected components and denote the resulting binary image by $I_{ES}^{(1)}$ (see Figure 3(g)). To retrieve the connectivity, a geodesic reconstruction by dilation of the resulting components contained in $I_{ES}^{(1)}$ is performed in a binary mask I_M that is dynamically built from the image I_{T_2} as follows:

$$I_M(i, j) = \mathbf{1}_{\{I_{T_2}(i,j)>S_M\}} \vee I_{ES}^{(1)}(i, j), \quad \forall (i, j) \in \Omega,$$

with

$$S_M = \left(\frac{1}{\sharp I_{ES}^{(1)}} \sum_{\substack{(i,j)\in\Omega \\ I_{ES}^{(1)}=1}} I_{T_2}(i, j) \right) - 0.4,$$

where $\mathbf{1}$, $\sharp$ and $\vee$ respectively denote the indicator function, the cardinality of a set and the logical *OR* operator. The threshold S_M is computed from the mean grey-level of the selected bright features in I_{T_2}, which is then lowered (-0.4) in order to surely get the whole components. We denote by I_{ES} the final binary image of the axial reflection components (see Figure 3(h)).

2.3. Detection of the darkest areas

k-means classification ($k = 3$) is performed on the pre-processed images I_{P_2} (see Figure 3(i)) and the cluster of lowest mean intensity value provides a first binary image of the darkest regions. It is then post-processed with morphological operations to get the main connected components corresponding to the dark areas of lumens. In particular, we retain the

components whose area is more than 1/5 of the largest one, do a morphological closing with a circular structuring element of radius 15 and fill holes that are smaller than 10% of the area of the component having the largest surface. We denote by I_{DA} the resulting image (see Figure 3(j)).

2.4. Extraction of vascular segments by information fusion

A first selection of vascular segments is performed based on a simple measure of the tortuosity. Let us denote by I^L_{ES} the binary image of a tested connected component of the image I_{ES}. This component is retained if

$$\frac{\sharp I^L_{ES}}{\sharp(I^L_{ES} \bullet S)} > 0.8, \tag{1}$$

where S is a binary disk whose radius is 15 pixels. Moreover, a segment of axial reflection must lie inside a dark area, and conversely, a dark region of the lumen must contain at least one axial reflection segment. We denote by $I_{ES'}$ the binary image made of the components satisfying (1) (see Figure 3(k)) and $I^L_{ES'}$ a tested component of it. The component $I^L_{ES'}$ is kept as part of an axial reflection segment if

$$\sharp(I_{DA} \cap (I^L_{ES'} \oplus S')) > \frac{\sharp I^L_{ES'}}{5}, \tag{2}$$

where $\oplus$ denotes the morphological dilation operator and S' is a binary disk whose radius is 15 pixels. Notice that the radius of S and S' are determined according to the minimum size of the vessels that are studied. We denote by $I_{ES''}$ the binary image made of the components satisfying (1) and (2) (see Figure 3(l)). Morphological operations are applied to I_{DA}, including reconstruction by dilation with the marker $I_{ES''}$, in order to get the final lumen mask I_{LM} (see Figure 3(m)).

2.5. Segment labeling and reconnection of branches

We first compute the skeleton of the image $I_{ES''}$ to get the end-points of the retained segments (see Section 2.4). These end-points are then reconnected using minimal path techniques (Cohen and Kimmel, 1997; Sethian, 1999). These techniques aim at extracting curves of minimal length, according to a Riemannian metric computed from the image and depending on the targeted application. A minimal path C connecting two end-points p and q is obtained by minimizing the following functional:

$$L[C] = \int_p^q \mathcal{P}(C(s))\, ds,$$

where s denotes the curvilinear abscissa and $\mathcal{P}$ is a potential inducing the metric, defined as

$$\mathcal{P}(x) = w_1(1 - I_{T_M}(x))^2 + w_2(1 - I_{SM}(x))^2 + w_3, \quad \forall x \in \Omega,$$

where I_{T_M} is the averages of images I_{T_1} and I_{T_2}, I_{SM} is the pruned skeleton of the lumen mask I_{LM} filtered by a Gaussian of standard deviation σ and $w_1, w_2, w_3 \in \mathbb{R}^+_*$ are free parameters. These parameters are empirically set to $\sigma = 10$, $w_1 = 0.5$, $w_2 = 0.45$ and $w_3 = 0.05$. In the latter expression, the first term should ideally be close to zero along the axial reflection.

The second term encourages the path C to pass near the middle of the lumen mask I_{LM}. The last term is a regularization constant. The combination of the above criteria allows for a good robustness against the variety of the encountered images.

Two end-points form a candidate pair for reconnection if they belong to the same connected component in the lumen mask I_{LM} and if they do not belong to the same connected component in the image $I_{ES''}$. The candidate pairs are then processed by decreasing order of the Euclidean distance (to start with points that are close to each other) and reconnected using the above procedure. A new skeleton is then calculated, providing the axial reflection of the vessels, and the vessel branches are then labeled (see Figure 3(n)). The vessel branches are individually regularized using a classical parametric active contour (Kass et al., 1988) with Gradient Vector Flow (Xu and Prince, 1998). The lumen mask I_{LM} is also labeled such that every non-null pixel receives the label of the closest branch (see Figure 3(o)).

Although the above described steps rely on a number of parameters, we empirically found that they are all stable for the images studied (including those presented in Section 4).

3. Segmentation of arterial walls

In what follows, we detail the procedure for segmenting arterial walls of a single regularized vessel branch obtained at the end of the axial reflection detection step (see Section 2). We denote this regularized branch as the reference line $V(s) = (x(s), y(s))^T \in \mathbb{R}^2$ of the vessel, parameterized by s. Once obtained, this line is considered to be fixed and will therefore no longer evolve in the subsequent steps. Additionally, we choose to model the artery wall by four curves $\{V_k\}_{k=1}^4$ approximately parallel to the reference line V where V_1, V_2 and V_3, V_4 represent the inner and the outer borders of this wall, respectively. These curves are defined as follows:

$$\begin{cases} V_1(s) &=& V(s) &+& b_1(s)\vec{n}(s) \\ V_2(s) &=& V(s) &-& b_2(s)\vec{n}(s) \\ V_3(s) &=& V(s) &+& b_3(s)\vec{n}(s) \\ V_4(s) &=& V(s) &-& b_4(s)\vec{n}(s), \end{cases}$$

where $\vec{n}(s) \in \mathbb{R}^2$ is the normal vector to the curve V and $b_k(s) \in \mathbb{R}^+$ is the local distance (or half-diameter) of any curve V_k to the reference line V (see Figure 4). This model establishes a direct correspondence between the points of any curve V_k and those of the reference line V. The segmentation aims at computing the half-diameters $\{b_k\}_{k=1}^4$ that are the most suitable for the delineation of the artery wall. By making strong (but realistic) assumptions, we first detail how these curves can be simultaneously initialized using a tracking procedure to cope with morphological deformations. Next, we relax these assumptions and present the parallel active contour model used to position these curves closer to artery walls.

3.1. Pre-segmentation

For convenience, we first sample the reference line V into equally spaced points and denote by $V(i) = (x(s = ih), y(s = ih))^T \in \mathbb{R}^2$ the coordinates of the i^{th} point along the reference

Figure 4: Parametric representation of the model (Lermé et al., 2014b).

line V, and by $\vec{n}(i) \in \mathbb{R}^2$ the associated normal vector ($h > 0$ denotes the sampling step, set to 1 here). For every $i \in \{1, \ldots, \sharp V\}$, we sample the curves representing the artery walls by

$$\begin{cases} V_1(i) &= V(i) + b^{int}(i)\vec{n}(i) \\ V_2(i) &= V(i) - b^{int}(i)\vec{n}(i) \\ V_3(i) &= V(i) + b^{ext}(i)\vec{n}(i) \\ V_4(i) &= V(i) - b^{ext}(i)\vec{n}(i) \end{cases}, \quad \text{s.t.} \quad b^{int}(i) < b^{ext}(i), \quad (3)$$

where $b^{int}, b^{ext} \in \mathbb{R}^+$ respectively denote the half-diameters of the inner and outer curves. Notice that the model (3) assumes that the inner (respectively outer) curves lie at the same distance from the reference line V (i.e. $b_1 = b_2 = b^{int}$ and $b_3 = b_4 = b^{ext}$). Additionally, we assume on both sides of V that the wall thickness is constant along the vessels (i.e. $b^{ext} - b^{int} = cst$). Although these assumptions could appear to be somewhat strong, they are verified for a large number of images.

More generally, the goal is to obtain a robust initialization of the active contour algorithm, which in turn, will refine the curves positioning so as to reach a better accuracy. The pre-segmentation aims at simultaneously computing the half-diameters b^{int} and b^{ext}. In this way, the robustness of the pre-segmentation against noise and poor contrast is substantially improved. For instance, this is typically useful when some of the borders of the artery walls are low contrasted since we can rely on those which are well contrasted.

Before presenting the pre-segmentation, we need to introduce some notations. First, we denote by $I : \Omega \subset \mathbb{Z}^2 \to [0, 1]$ a grayscale image and $D_{\vec{u}}I(p)$ the derivative of I in the direction $\vec{u}$ at a pixel $p \in \Omega$. For a half-diameter $b \in \mathbb{R}^+$, a point $V(i)$ on the reference line V and a window of size $(2r + 1)$, we also define the mean local gradient along the curves V_1 and V_2 by

$$\bar{D}^{int}(b, i, r) = \frac{1}{2(2r+1)} \sum_{j=-r}^{r} \Big(D_{\vec{n}(i+j)}I(V_1(i+j)) +$$
$$D_{-\vec{n}(i+j)}I(V_2(i+j)) \Big), \quad (4)$$

and the mean local gradient along the curves V_3 and V_4 by

$$\bar{D}^{ext}(b, i, r) = \frac{1}{2(2r+1)} \sum_{j=-r}^{r} \Big(|D_{\vec{n}(i+j)}I(V_3(i+j))| +$$
$$|D_{-\vec{n}(i+j)}I(V_4(i+j))| \Big). \quad (5)$$

Using (4) and (5), the pre-segmentation is based on a criterion to be maximized, defined for every $i \in \{1, \ldots, \sharp V\}$ by:

$$G(b^i, b^e, i, r) = \bar{D}^{int}(b^i, i, r) + \bar{D}^{ext}(b^e, i, r). \quad (6)$$

This criterion encourages inner and outer curves to be located near large gradients in the image. Notice that (4) does not use absolute values on directional derivatives unlike (5). Increasing the window radius r makes gradient measures more robust to noise but less reliable when strong deformations occur along vessels. This parameter therefore requires a trade-off. In what follows, we detail the steps necessary to estimate the half-diameters b^{int} and b^{ext} in (3).

Step 1. Instead of jointly estimating constant half-diameters b^{int_0} and b^{ext_0} as in (Lermé et al., 2014b), we propose to estimate them as piecewise constant. In this way, the estimated half-diameters are less prone to morphological deformations while keeping a good robustness against intensity changes along the artery walls. The robustness of the tracking (see Step 3 below) is therefore improved, especially for pathological cases. Let us denote a vessel segment k whose starting and ending indices in the reference line are k_{start} and k_{end}, respectively. For each vessel segment k, the half-diameters are estimated using the criterion (6) as follows:

$$(b^{int_0^k}, b^{ext_0^k}) = \operatorname*{argmax}_{\substack{b^i, b^e \in \mathbb{R}^+ \\ b^i < b^e}} \frac{1}{2(k_{end} - k_{start} + 1)} \sum_{i=k_{start}}^{k_{end}} G(b^i, b^e, i, r).$$

We then set $b^{int_0}(i) = b^{int_0^k}$ and $b^{ext_0}(i) = b^{ext_0^k}$, $\forall i \in \{k_{start}, \ldots, k_{end}\}$. To speed up this step, the intervals for b^i and b^e are restricted to typical values (Koch et al., 2014). In accordance to morphometric features, the piecewise constant estimation of the half-diameters is performed on vessels segments whose length is 50 pixels ($\simeq 66.5\mu m$). Such an approach can however fail to accurately segment the walls when deformations occur inside the same vessel segment. Estimating varying half-diameters is therefore essential to take into account these deformations.

Step 2. We determine the position along the curves having the largest contrast by maximizing the criterion (6):

$$i^* = \operatorname*{argmax}_{i \in \{1, \ldots, \sharp V\}} G(b^{int_0}, b^{ext_0}, i, r).$$

Step 3. Finally, we jointly estimate varying half-diameters b^{int} and b^{ext} whose difference (the wall thickness) is constant along vessels. Again, this is achieved by using the criterion (6). These half-diameters are constructed iteratively under a regularity constraint from each side of the position i^* found at Step 2 and using the estimates b^{int_0} and b^{ext_0} found at Step 1. Let us denote by $\bar{e} = (b^{ext_0}(i^*) - b^{int_0}(i^*))$ the wall thickness at the position i^*. For a fixed error e, the half-diameters are constructed as follows:

$$\begin{aligned} b^{int}(i) &= \begin{cases} b^{int_0}(i^*) & \text{if } i = i^* \\ \operatorname{argmax}_b E(b, b^{int}(i-1), i, e, r) & \text{if } i > i^* \\ \operatorname{argmax}_b E(b, b^{int}(i+1), i, e, r) & \text{if } i < i^*, \end{cases} \\ b^{ext}(i) &= b^{int}(i) + \bar{e} + e \end{aligned}$$
$$(7)$$

with

$$E(b, b', i, e, r) = \alpha G(b, b + \bar{e} + e, i, r) - (1 - \alpha)(b - b')^2, \quad (8)$$

and where $\alpha \in [0, 1]$ is a regularization parameter. The closer the parameter α is to zero, the more the right term of (8) penalizes large deviations of b^{int}, and conversely. To get a more robust estimate of the wall thickness $\bar{e} + e$, we propose to slightly vary e around $\bar{e}$ and select the error e^* achieving the maximum energy (8) summed along the vessel:

$$e^* = \underset{e}{\operatorname{argmax}} \sum_{i>0} E(b^{int}(i), b^{int}(i - 1), i, e, r).$$

In the latter equation, b^{int} and b^{ext} are constructed using (7) and (8). The resulting half-diameters b^{int^*} and b^{ext^*} are those found for e^*. It is not difficult to see that the assumptions made at the beginning of this section hold for these half-diameters.

3.2. Refined segmentation

The model proposed by (Ghorbel et al., 2011) simultaneously evolves two curves under a parallelism constraint. In what follows, we describe an extension of this model for extracting four curves V_1, V_2, V_3 and V_4 almost parallel to a reference line V. Since this line is fixed, the energy becomes

$$E(V_1, \ldots, V_4, b_1, \ldots, b_4) = \sum_{k=1}^{4} \left(E_{Image}(V_k) + R(V_k, b_k) \right), \quad (9)$$

where the term

$$E_{Image}(V_k) = \int_0^1 P(V_k(s)) \, ds,$$

is designed to attract the curve V_k towards large intensity gradients (see (Kass et al., 1988)). In this context, the term E_{Image} is based on the Gradient Vector Flow (Xu and Prince, 1998). The role of the term R in (9) is to control the variation of the distance b_k, thus imposing a local parallelism. (Ghorbel et al., 2011) proposed a function of the derivative of b_k with

$$R(V_k, b_k) = \int_0^1 Q(s, b_k') \, ds = \int_0^1 \varphi_k(s)(b_k'(s))^2 \, ds,$$

where $\{\varphi_k\}_{k=1}^4$ are application-dependent parameters that locally control the strength of the parallelism of the curve V_k with respect to the reference line V. More precisely, the larger these parameters are, the stricter the parallelism to the reference line V is. Unlike previous active contour methods embedding a parallelism constraint, it is important to note that the distance between any curve V_k and the reference line V has not to be provided in the model. It is adjusted during the evolution process and can vary along boundaries. It is worth noting that the assumptions made for the pre-segmentation (see Section 3.1) are relaxed, i.e. the curves can now evolve independently of each other (instead of (Lermé et al., 2014a) where curves are linked to others). Also, we want to mention that the energy (9) does not ensure that $b_1(s) < b_3(s)$ and $b_2(s) < b_4(s)$, $\forall s$. However, we never encountered such a behavior in our experiments (including those presented in Section 4).

Since the energy (9) does not have crossing terms involving different curves, the minimization can be independently done for each curve V_k. For any $k \in \{1, \ldots, 4\}$, the Euler-Lagrange equation expresses the minimization of (9) with respect to $b_k(s)$

$$\frac{\partial P(V_k(s))}{\partial b_k} - \frac{d}{ds} \frac{\partial Q(s, b_k')}{\partial b_k'} = 0,$$

and the evolution of the distance b_k to the reference line V is driven by

$$\langle \vec{n}, -\nabla P(V_k(s)) \rangle - 2 \left(\varphi_k(s) b_k''(s) + \varphi_k'(s) b_k'(s) \right) = 0.$$

The latter equation is solved by discretizing it and introducing the time variable using standard numerical approximations of derivatives (central difference in space, backward difference in time). The resolution of the above equations stops when

$$\max_s |b_k^n(s) - b_k^{n-1}(s)| \le \varepsilon, \quad \forall k \in \{1, \ldots, 4\},$$

where $b_k^n(s)$ is the distance of the curve V_k to the reference line V at iteration n and $\varepsilon \simeq 0$ is an accuracy parameter.

4. Experimental results and discussion

4.1. Datasets

A first set of 16 images from healthy and pathological subjects (denoted by D_1) is used for evaluating the axial reflections detection in complex situations (e.g. with several branches). Since well contrasted vessels are generally limited to small localized segments, these images are not adapted to the evaluation of walls. To this aim, a second set of images is used (denoted by D_2), distinct from D_1, composed of 14 and 17 images from healthy and pathological subjects, respectively. D_2 is also used for evaluating the axial reflection detection. All these images were selected by physicians having several years of experience in the field of AO image interpretation to ensure the representativeness of the situations encountered during clinical routine, in terms of noise levels, contrast and morphological deformations.

4.2. Axial reflections

From each image of D_1 and D_2, the axial reflections of the vessels whose lumen is visible were identified by a physician by first interactively entering a few points along them and then applying a classical parametric active contour algorithm (Kass et al., 1988) with Gradient Vector Flow (Xu and Prince, 1998).

When no bifurcations occur, the automatic detection of axial reflections (see Section 2) can result in fragmented reference lines due to intensity changes along them. To not overpenalize the automatic procedure, we propose to put in correspondence these lines with those obtained semi-interactively by the physician. Let us denote by V^A and V^M a sampled reference line obtained from the automatic procedure and from the physician, respectively. Then, we consider that these reference lines match with each other if V^A overlaps V^M by at least 50% of its length, or more formally, if

$$\frac{L(Proj(V^A, V^M))}{L(V^A)} > 0.5, \quad (10)$$

Table 1: Description of the measures used for comparing the number of reference lines in manual and automatic segmentations.

True Positives (TP)	Number of correctly matched reference lines between automatic and manual segmentations.
False Positives (FP)	Number of reference lines in automatic segmentation that do not match a manual one.
False Negatives (FN)	Number of reference lines in manual segmentation that do not match an automatic one.

where $Proj(V^A, V^M)$ is a subset of points of V^A such that each of them has its nearest point in V^M that lies at a distance less than or equal to 5 pixels and $L(x)$ denotes the piecewise linear approximation of the length of the line x, defined by

$$L(x) = \sum_{i=2}^{\sharp x} \|x_i - x_{i-1}\|_2.$$

Notice that the points of $Proj(V^A, V^M)$ in (10) can be unequally spaced due to the distance threshold. The latter permits us to tolerate an imperfect positioning between the reference lines V^A and V^M. Once the matching is done between these reference lines and to give an idea of the performance of the axial reflection detection, we propose to compare an automatic segmentation and a manual one by relying on measurements based on (i) the number of reference lines detected (see Table 1) and (ii) their relative overlap. Given a set of automatically detected reference lines $\{V^{A_k}\}_{k=1}^{n}$ that match with a manual reference line V^M using (10), this overlap (in percentage) is defined by

$$RO(\{V^{A_k}\}_{k=1}^{n}, V^M) = \left(1 - \frac{|L(V^M) - \sum_{k=1}^{n} L(V^{A_k})|}{L(V^M)}\right) \times 100. \quad (11)$$

The performance of the automatic axial reflection detection against manual segmentations using the above measures is detailed in Table 2 and illustrated in Figure 5. To address the generalization capability, the automatic procedure is optimized on a training set (D_1). The test set corresponds to D_2. In Table 2, we can see that the automatic procedure has a good robustness with very few false positives / false negatives and a large mean RO, meaning that most of the axial reflections are detected and overlap well with the ideal ones. In particular, the automatic procedure demonstrates its capacity to properly manage single or multiple artery branches (see Subjects 8 and 10). These results appear to be in the same range for the test set D_2, meaning a good generalization of the automatic procedure. However, we can notice that some images contain falsely detected axial reflections (see Subject 15). While undesired, these false detections remain minor in practice and can be easily removed by medical experts in our software. Also, we can notice that one axial reflection is missing (see Subject 11). All these particular cases are depicted in Figure 5. In particular, we observe that the detection of axial reflections is very challenging for the Subjects 11 and 15 due to the large amount of blur on some vessel segments. However, such vessels are not relevant for physicians since the arterial walls are not visible.

Table 2: Performance of the axial reflection detection on the training set D_1 and the test set D_2. False positives (FP), false negatives (FN), true positives (TP) and the relative overlap (RO, in %) are provided (see Table 1 and (11)) for each set of images.

	Number of branches	FP	FN	TP	RO
Training set	39	3	1	38	92.38% ± 8.28%
Test set	37	4	1	36	85.26% ± 20.01%

4.3. Artery walls

From each image of D_2, a single artery branch was manually segmented two times by three physicians, with an interval of several weeks between the inputs to attenuate the memory effect. To measure the accuracy of a segmentation A obtained with our approach against a segmentation M realized by a physician, we use the absolute relative error on the inner diameter, the outer diameter and the wall thickness, respectively denoted by δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$. For each image, these measures are taken where no vessel bifurcation occur and on the intersection between automatic and manual segmentations. Considering the wall thickness on each side separately (δ_{wt_1} and δ_{wt_2}) instead as a whole (as previously done in (Lermé et al., 2014b)) is more sensitive but also more realistic. Notice that the measure $\delta_{wt_{1,2}}$ is of great importance for physicians.

First, we estimate the intra-physician variability by computing the mean and the standard deviation of the measures δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ for all images. To put in perspective the results and since arterial walls are very difficult to delineate up to a pixel, we provide the mean and the standard deviation of δ_{id} and δ_{od} when one of the two delineated curves is shifted by one pixel all along it (see Section 3.2). For $\delta_{wt_{1,2}}$, the two concerned curves depend on which side of the reference line the measure is performed. The resulting statistics are summarized in Table 3. As expected, the mean error on the walls thickness appears to be much larger than for the inner and outer diameter due to its high sensitivity. Also, we can also notice that the physician $Phys_1$ has the smallest variability on both healthy and pathological subjects for two measures out of three. Because this physician produced the most stable segmentations, we decided to choose him as a reference for the remainder of this evaluation and denote him by $Phys_{Ref}$.

In what follows, the influence of the parameter α involved in the pre-segmentation step (see Section 3.1) is studied. In particular, the window radius r is not considered here since our experiments reveal that this parameter has a limited influence on the results. We therefore choose to set $r = 1$. We believe that this behavior is due to the dominant role of (8) in the tracking.

First, we study the sensitivity to the parameter α. Figure 6 shows how the parameter α affects the accuracy of the pre-segmentations compared to those performed by the physician $Phys_{Ref}$ for both healthy and pathological subjects. Clearly, the segmentation error is stable for $\alpha \in [0, 0.65]$ and relatively small around 0.95 on the interval $[0.9, 1]$.

Secondly, we study the impact of the parameter α on pre-segmentations and segmentations. For doing so, we constitute a training set and a set set using all the images from healthy and pathological subjects (D_2). The training set is composed of one

| Subject 8 | Subject 10 | Subject 11 | Subject 15 |

Figure 5: Good (half-left) and poor (half-right) results from Table 2 for the automatic detection of axial reflections for four distinct subjects. Top and bottom rows depict respectively automatic and manual segmentations, superimposed on the original image. White and cyan dashed lines correspond to missed and falsely detected reference lines, respectively. Dashed lines that are neither white or cyan correspond to correctly matched reference lines.

Table 3: Intra-physician variability with the absolute relative error on the inner diameter (ID), the outer diameter (OD) and the wall thickness (WT) (see δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ in the text). All measures are expressed in percentages and averaged over vessel and images. For each measure, the numbers reported between parentheses are the errors made by shifting one of the four curves by one pixel, all along it.

	$Phys_1$ / $Phys_1$	$Phys_2$ / $Phys_2$	$Phys_3$ / $Phys_3$
ID	3.74 ± 3.38 (1.03 ± 0.32)	4.09 ± 3.49 (1.02 ± 0.31)	3.44 ± 2.84 (1.01 ± 0.31)
OD	2.56 ± 2.04 (0.75 ± 0.21)	3.52 ± 2.99 (0.76 ± 0.21)	3.03 ± 2.58 (0.74 ± 0.21)
WT	20.33 ± 20.43 (6.01 ± 2.09)	27.52 ± 31.95 (6.76 ± 3.08)	22.01 ± 22.33 (6.03 ± 2.37)

Table 4: Robustness of the pre-segmentation on healthy and pathological subjects against manual segmentations from the physician $Phys_{Ref}$. The parameter α^* achieving the best accuracy on the training set is computed and used for pre-segmenting the images from the test set. Parallel snakes (PS) are then applied. The accuracy is then estimated on the inner diameter (ID), the outer diameter (OD) and the wall thickness (WT) (see δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ in the text). The inter-physicians error is also provided. For each measure, the numbers reported between parentheses are the errors made by shifting one of the four curves by one pixel, all along it.

		Pre-segmentation (α^*)	Pre-segmentation (α^*) + PS	Inter-physicians error
	ID	5.62 ± 7.31 (1.13 ± 0.49)	5.33 ± 7.11 (1.13 ± 0.49)	4.09 ± 3.71 (1.13 ± 0.49)
Training set	OD	3.46 ± 3.45 (0.83 ± 0.31)	3.35 ± 3.51 (0.83 ± 0.31)	3.24 ± 2.79 (0.83 ± 0.31)
	WT	16.60 ± 12.93 (6.74 ± 2.65)	16.49 ± 13.17 (6.74 ± 2.65)	21.84 ± 17.44 (6.75 ± 2.68)
	ID	5.97 ± 4.74 (0.98 ± 0.17)	5.15 ± 4.28 (0.98 ± 0.17)	4.07 ± 3.37 (0.98 ± 0.17)
Test set	OD	3.51 ± 2.77 (0.71 ± 0.11)	3.19 ± 2.46 (0.71 ± 0.11)	3.34 ± 2.81 (0.71 ± 0.11)
	WT	19.14 ± 15.55 (5.60 ± 1.65)	18.11 ± 16.81 (5.60 ± 1.65)	23.42 ± 23.36 (5.65 ± 1.63)

Table 5: Comparison of the proposed pre-segmentation (computed using α^* that gives the best accuracy against $Phys_{Ref}$) compared to a naive one (computed using $\alpha = 0$) in terms of convergence speed and accuracy when coupled to parallel snakes (PS). The accuracy is estimated using the absolute relative error on the inner diameter (ID), outer diameter (OD) and the walls thickness (WT) (see δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ in the text) against manual segmentations from the physician $Phys_{Ref}$. For these measures, the inter-physicians error is also provided. The convergence speed is estimated by the number of iterations needed by the PS model to converge.

	Pre-segmentation ($\alpha = 0$)	Pre-segmentation ($\alpha = 0$) + PS	Pre-segmentation (α^*)	Pre-segmentation (α^*) + PS	Inter-physicians error
ID	9.48 ± 9.01	7.78 ± 9.94	5.86 ± 5.70	5.21 ± 5.36	4.08 ± 3.48
OD	6.34 ± 5.02	5.40 ± 5.68	3.49 ± 3.00	3.24 ± 2.84	3.31 ± 2.80
WT	27.16 ± 14.90	30.15 ± 28.27	18.32 ± 14.80	17.59 ± 15.75	22.91 ± 21.64
Iterations	/	90.74 ± 50.23	/	59.03 ± 31.68	/

third of images from healthy subjects and one third of images from pathological subjects, all randomly selected. The test set is composed of the remaining images.

Then, we search for the value of the parameter α giving the best accuracy against manual segmentations from the physician $Phys_{Ref}$ on the training set. More precisely, we choose the optimized value α^* as the one that minimizes the mean of δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ along arteries and over images. We pre-segment the images from the test set using α^* and apply the parallel active contour model (PS) (see Section 3.2) on the training and test sets using the following parameter values: $\varepsilon = 0.1$ and $\varphi_k = 100$, $\forall k \in \{1, \ldots, 4\}$. To estimate the accuracy of our approach on the two sets of images (D_2), we again compute the mean and the standard deviation of the measures δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ between the resulting segmentations and those performed by the physician $Phys_{Ref}$. The inter-physicians variability is also computed for the same images by comparing the segmentations performed by all the physicians, except $Phys_{Ref}$, with respect to those realized by $Phys_{Ref}$. All these statistics are summarized in Table 4. In this table, we can see that the resulting measures appear to be approximately in the same range, meaning good generalization properties of the parameter α. The pre-segmentation offers a good accuracy when compared to the inter-physicians error. When parallel snakes are applied from these pre-segmentations, the accuracy is improved for all measures on both sets of images. For two measures out of three, the accuracy on the test set is very good given the inter-physicians error. A subset of the results is given in Figure 7 for two healthy subjects (first and third rows) and two pathological subjects (second and fourth rows). Mislocations in manual and automatic segmentations are pointed by orange and green arrows in images, respectively. On the first two rows, we see that the pre-segmentation is robust to morphological deformations and is able to delineate thin walls. However, the last row depicts an extreme situation where the pre-segmentation fails due to the poor contrast on the inner borders of walls and a fake line in the textured background that mislead the tracking procedure (green arrows). Also, we believe that the least performance of our approach is due to a lack of accuracy in the manual segmentations because artery walls are difficult to delineate. This issue is shown in the two last rows of Figure 7 (orange arrows). It is worth noting that the walls are not strictly parallel to the reference line for pathological cases. However, we do believe that it still makes sense to consider an approximate parallelism since its deviation remains moderate. While the model imposes a local parallelism, it allows diameter variations at a larger scale.

Finally, Table 5 demonstrates, on both healthy and pathological subjects, the benefit of using a pre-segmentation that follows morphological deformations along vessels compared to a naive pre-segmentation that does not. For the first one, we use the optimized value α^* (see above). For the second one, we use $\alpha = 0$ that leads to constant half-diameters along arteries. We then apply the PS model on all these pre-segmentations and compare the required number of iterations as well as the accuracy with respect to the manual segmentations from the physicians $Phys_{Ref}$. In words, the pre-segmentation using α^* both leads in average to a better accuracy for all measures and diminishes the


Figure 6: Sensitivity of the parameter α on both healthy and pathological subjects. The y-axis is the mean error of the measures δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ with respect to the segmentations performed by the reference physician $Phys_{Ref}$.

number of iterations by about 35%.

4.4. Practical use

We discuss below the potential usage of the method for a non-specialist in image processing. To ensure a good accuracy of the resulting segmentations, the method should be applied on a clean image where the walls of arteries are well contrasted, i.e. visually distinguishable from surrounding structures. Whenever it is possible, it is preferable to focus the analysis on a region of medical interest (Paques et al., 2014a; Koch et al., 2014) to limit the computation time and to avoid wrongly detected vessels due to linear structures in the background. In case of failure, manual corrections can be performed on axial reflections and wall detection, by manually initializing the snakes. Finally, note that the user interface developed to address these tasks is easy to handle, stable and reactive as acknowledged by the positive feedback we had from the collaborating hospitals (Hôpital des Quinze-Vingts and Hôpital la Pitié Salpétrière). The overall time for automatically delineating walls of a single artery from a region of interest of 900×900 pixels is $\simeq 4$ minutes on a PC with a Intel Core i7-3612QM @ 2.10GHz. However, our software has not been optimized and we do believe that this time would be inferior by using a C implementation.

5. Conclusion

In this paper, we have proposed an automatic method for delineating the walls of retina arteries that apply on healthy and pathological cases. Noticeably, the results showed a good robustness for detecting axial reflections and an overall error on arterial walls measurements smaller or very near the inter-physicians error. For future work, we plan to evaluate the benefit of the coupled parallel snakes model introduced in (Lermé et al., 2014a) for different pathologies. Finally, notice that this work is not limited to ophthalmology applications and could be for instance applied to arterial hypertension in cardiology.

6. Acknowledgments

Pre-segmentations   Refined segmentations

2.85 ± 2.14 (ID), 1.91 ± 1.39 (OD), 11.13 ± 7.70 (WT) 2.76 ± 1.56 (ID), 1.49 ± 1.04 (OD), 11.18 ± 7.64 (WT)

4.91 ± 5.18 (ID), 3.04 ± 2.79 (OD), 14.32 ± 14.52 (WT) 4.33 ± 4.94 (ID), 2.26 ± 1.74 (OD), 15.36 ± 14.67 (WT)

3.48 ± 2.42 (ID), 3.25 ± 1.89 (OD), 21.60 ± 20.30 (WT) 2.17 ± 1.65 (ID), 4.77 ± 1.13 (OD), 20.77 ± 19.18 (WT)

13.71 ± 17.00 (ID), 7.54 ± 6.11 (OD), 20.18 ± 13.21 (WT) 14.03 ± 18.13 (ID), 7.55 ± 6.70 (OD), 20.04 ± 14.47 (WT)

Figure 7: From top to bottom: good (upper-half) and poor (lower-half) segmentation results against the physician $Phys_{Ref}$. First and third rows correspond to healthy subjects while second and fourth rows correspond to pathological ones. The left column shows pre-segmentations (α^*) while the right column shows segmentations obtained by applying parallel snakes from them. Manual and automatic segmentations are drawn in purple and cyan, respectively, superimposed on the original images. The white dashed line is the reference line V while the white circle denotes the position along it from which the tracking operates in the pre-segmentation. The absolute relative error on inner diameter (ID), outer diameter (OD) and walls thickness (WT) (see δ_{id}, δ_{od} and $\delta_{wt_{1,2}}$ in the text) is indicated below each image. Mislocations in manual and automatic segmentations are pointed by orange and green arrows in images, respectively.

References

Chui, T., Vannasdale, D., Burns, S., 2012. The use of forward scatter to improve retinal vascular imaging with an adaptive optics scanning laser ophthalmoscope. Biomedical Optics Express 3, 2537–2549.

Cohen, L., Kimmel, R., 1997. Global minimum for active contour models: A minimal path approach. International Journal of Computer Vision 24, 57–78.

Ghorbel, I., Rossant, F., Bloch, I., Paques, M., 2011. Modeling a parallelism constraint in active contours. Application to the segmentation of eye vessels and retinal layers, in: International Conference on Image Processing (ICIP), pp. 445–448.

Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: Active contour models. International Journal of Computer Vision 1, 321–331.

Koch, E., Rosenbaum, D., Brolly, A., Sahel, J.A., Chaumet-Riffaud, P., Girerd, X., Rossant, F., Paques, M., 2014. Morphometric analysis of small arteries in the human retina using adaptive optics imaging: Relationship with blood pressure and focal vascular changes. Journal of Hypertension 32, 890–898.

Kulcsár, C., Le Besnerais, G., Ödlund, E., Levecq, X., 2013. Robust processing of images sequences produced by an adaptive optics retinal camera, in: Optical Society of America, Adaptive Optics: Methods, Analysis and Applications, p. OW3A.3.

Lermé, N., Rossant, F., Bloch, I., Paques, M., Koch, E., 2014a. Coupled parallel snakes for segmenting healthy and pathological retinal arteries in adaptive optics images, in: International Conference on Image Analysis and Recognition (ICIAR), pp. 311–320.

Lermé, N., Rossant, F., Bloch, I., Paques, M., Koch, E., 2014b. Segmentation of retinal arteries in adaptive optics images, in: International Conference on Pattern Recognition (ICPR), pp. 574–578.

Liang, J., Williams, D., Miller, D., 1997. Supernormal vision and high-resolution retinal imaging through adaptive optics. Journal of Optical Society America A 14, 2884–2892.

Martin, J., Roorda, A., 2005. Direct and noninvasive assessment of parafoveal capillary leukocyte velocity. Ophthalmology 112, 2219–2224.

Paques, M., Rossant, F., Brolly, A., Benesty, J., Lermé, N., Koch, E., Bloch, I., Girmens, J.F., 2014a. Venous nicking without arteriovenous contact: Evidence for the role of the arteriolar microenvironment in arteriovenous nickings. Published online on JAMA Ophtalmology .

Paques, M., Rossant, F., Lermé, N., Miloudi, C., Kulcsár, C., Sahel, J.A., Mugnier, L., Bloch, I., Loquin, K., Koch, E., 2014b. Adaptive optics imaging of retinal microstructures: Image processing for medical applications, in: International Workshop on Computational Intelligence for Multimedia Understanding (IWCIMU).

Rossant, F., Badellino, M., Chavillon, A., Bloch, I., Paques, M., 2011. A morphological approach for vessel segmentation in eye fundus images, with quantitative evaluation. Journal of Medical Imaging and Health Informatics 1, 42–49.

Sethian, J., 1999. Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science. Cambridge University Press.

Viard, C., Nakashima, K., Lamory, B., Paques, M., Levecq, X., Château, N., 2011. Imaging microscopic structures in pathological retinas using a flood-illumination adaptive optics retinal camera, in: Photonics West: Biomedical Optics (BiOS), pp. 788 509+.

Weickert, J., Romeny, B., Viergever, M., 1998. Efficient and reliable schemes for nonlinear diffusion filtering. IEEE Transactions on Image Processing 7, 398–410.

Xu, C., Prince, J., 1998. Snakes, shapes and gradient vector flow. IEEE Transactions on Image Processing 7, 359–369.