



HAL
open science

Tubular Structure Analysis by Ranking the Orientation Responses of Path Operators

Odyssée Merveille, Hugues Talbot, Laurent Najman, Nicolas Passat

► **To cite this version:**

Odyssée Merveille, Hugues Talbot, Laurent Najman, Nicolas Passat. Tubular Structure Analysis by Ranking the Orientation Responses of Path Operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017, 10.1109/TPAMI.2017.2672972 . hal-01262728v1

HAL Id: hal-01262728

<https://hal.science/hal-01262728v1>

Submitted on 27 Jan 2016 (v1), last revised 21 Feb 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Tubular Structure Analysis by Ranking the Orientation Responses of Path Operators

Odyssee Merveille, Hugues Talbot, *Member, IEEE*, Laurent Najman, Nicolas Passat

Abstract—The analysis of thin tubular objects in 3D images is a complex and challenging task. In this article, we introduce a new, nonlinear operator, called RORPO (Ranking Orientation Responses of Path Operators). Inspired by the multidirectional paradigm currently used in linear filtering for thin structure analysis, RORPO is built upon the notion of path operator from mathematical morphology. This operator, unlike the gold-standard Hessian-based operators commonly used for 3D tubular structure analysis, is discrete, non linear and non-local. From this new operator, two main tubular structure characteristics can be estimated: an intensity feature, that can be assimilated to a quantitative measure of tubularity; and a directional feature, providing a quantitative measure of the tubular structure orientation. We provide a full description of the structural and algorithmic details for computing these two features from RORPO, and we discuss computational issues. We experimentally assess RORPO by comparing with gold standard Vesselness and we show that our method performs better on both features in realistic conditions.

Index Terms—Thin structures, nonlinear filtering, direction estimation, mathematical morphology, path opening, 3D grey-level imaging, blood vessel enhancement, vesselness.

1 INTRODUCTION

THIN structures in n D images are characterized by a significantly smaller size in at least one of their n dimensions. In most image-related applications, n is 2 or 3. We can then define two kinds of thin structures: 1D thin structures, typically line-like or tube-like objects respectively in 2D or 3D images; and 2D thin structures which are plane-like objects in 3D images. In this article, we mainly focus on 1D thin structures in 3D images, which we term *tubular structures*.

Images of 1D thin structures are among the hardest to handle in image processing. The difficulty lies both in the sparsity of the images and their complex geometry. In addition, such structures are naturally fragile, in that even a small amount of noise may be enough to disrupt their contours, leading to disconnections or misconnections. 1D thin structures can also be very tortuous, possess different orientations and scales and form a network, making geometric prior difficult to use.

Odyssee Merveille, Laurent Najman and Hugues Talbot are with Université Paris-Est ESIEE, Laboratoire d'informatique Gaspard-Monge, France (odyssee.merveille,laurent.najman,hugues.talbot@esiee.fr).

Odyssee Merveille and Nicolas Passat are with Université de Reims Champagne-Ardenne, CReSTIC, Reims, France (nicolas.passat@univ-reims.fr).

Many applications involve 1D thin structures e.g., road extraction in remote sensing images, vessels segmentation in medical images, or fibre detection in material science. In order to segment these 1D thin structures, it is frequent to use a general segmentation framework (e.g., active contours, region-growing, machine learning) enriched with some specific prior knowledge associated with the thin structures.

Various methods have been developed to extract specific thin structure features. These methods aim at enhancing the thin structure signal and decreasing the response of non-thin structures, which results in the extraction of features based on their intensity. Higher-level features can also be estimated, such as orientation, diameter or curvature of thin structures (see Sec. 2).

The current gold-standard for feature extraction of thin structures relies on linear, local, continuous, operators, based on a multidirectional and multiscale paradigm. More precisely, Hessian-based filters have contributed to the development of vesselness functions [1]. We recently introduced a new operator called RORPO (Ranking Orientation Responses of Path Operators) [2]. This operator still relies on a multidirectional and multiscale paradigm but it is a non-linear, global, discrete operator. In [3], it was experimentally observed that RORPO is robust, competitive and complementary to Hessian-based operators for the low-level processing of thin structures.

In this article, we present RORPO in more details, and we propose two fundamental thin structure features, namely (1) a low-level intensity feature that preserves the thin structures while removing the signal of other objects; and (2) a directional feature providing an estimation of the orientation of a putative thin structure.

In Sec. 3, we develop the general strategy behind RORPO and its underlying concepts. In Sec. 4 we present the computation of RORPO and of its associated features. Sec. 5 discusses algorithmic considerations including RORPO parameters, computational cost, and robustness to noise. An experimental validation of RORPO is carried out in Sec. 6. A discussion on extensions and future work concludes this article (Sec. 7).

2 RELATED WORKS: TUBULAR STRUCTURE FILTERING

Tubular structure filtering in 3D grey-level images is an active research area, that led to the proposal of many methods over the last two decades. A complete survey is beyond the scope of this article. The reader may refer to [4, Chs. 1, 2] for a general overview, and to [5], for a medical-oriented survey. We propose here a global, but non-exhaustive vision of the principal families of approaches, illustrated by representative methods.

Methods devoted to filtering tubular structures can be divided into two categories, namely those relying on differential (mostly linear scale-space) operators, and those relying on nonlinear (often mathematical morphology) operators.

2.1 Linear approaches

The majority of linear approaches are based on local, scale-space, differential analysis of images, viewed as continuous functions. In particular, the analysis of second-order derivative properties of the image were first proposed in [6], [7]. In these methods, the eigenvectors of multi-scale Hessian matrices and their associated eigenvalues are analyzed to characterize blobs, planar and tubular structures as well as their scale and orientation.

This strategy has led to the proposal of several “vesselness” measures, combining differential information into heuristic formulations. The vesselness proposed by Frangi et al. [1] is often considered as the current gold-standard. Many methods/variants have been proposed since then. Some of them also used the eigenvectors obtained from the Hessian matrix [8], [9], for instance for guidance of a diffusion framework [10]. In [11] the second derivatives were associated to first derivatives and a Canny filter, while in [12], a strain energy function used a stress tensor computed from the Hessian tensor.

To achieve multi-scale detection, derivative operators are combined with a convolution kernel. In order to avoid the induced blurring effects, these may be replaced by a gradient vector flow [13]. The use of a bi-Gaussian kernel was proposed to better take into account the anisotropy of tubular structures [14]. Optimally oriented flux have also been considered [15].

Alternatively to Hessian-based approaches, steerable filters are anisotropic filters that can be expressed in terms of a linear combination of basis filters [16]. In [17], the convolution between a bar profile and the second derivative is used to introduce a multi-scale approach. 3D steerable filters were first proposed in [18], using a n th Gaussian derivative basis. In [19] 3D steerable filters based on the second and fourth Gaussian derivatives were used to detect dendritic profiles.

Discrete gradients have also been used. In [20], tube-like orientation is first estimated using a set of discrete orientations. Then instead of a classical low pass filter, an anisotropic filter is used to enhance tube-like features.

The maximum curvature of tube-like structure is also computed by the second derivative operator along the 13 discrete lines of a $3 \times 3 \times 3$ kernel [21], [22].

2.2 Nonlinear approaches

2.2.1 SE-based approaches

Nonlinear approaches include those based on mathematical morphology [23]. In this framework, a common notion is the structuring element (SE), a geometric pattern from which basic operators (erosions, dilations, openings, closings, etc.) can be defined. Considering local straightness and extremal intensity of tubular structures, a basic idea is to propose filters using small structuring elements fitting these properties.

Two dual approaches have been considered. The first models a tube by a small straight line SE [24], [25], to carry out opening or closing operations by line segments of arbitrary orientation [26], [27], [28]. The second models the background in the orthogonal hyperplane of the segment, to carry out grey-level hit-or-miss transform [29], [30], [31].

Even if the shapes of the SEs are hard parameters, the orientation parameters need to be specified, to form filter banks. Several strategies were experimented: rotating structuring elements [32], knowledge-based parametrization [30], or spatially-variant mathematical morphology [33]. Blurred operators [31] were also designed to increase the robustness of the filtering.

Nevertheless, the hard, straight geometry of such SEs remains a limitation to the accuracy of these approaches, progressively leading to using more flexible SEs [34].

2.2.2 Connectivity-based approaches

A second notion is that of connectivity, generally handled on graphs. The key notion of connectivity is no longer the local notion of SE, but a more global notion of connected component. In this context, the concept of component-tree [35] and attribute-based methods was specifically investigated and developed for extensive, anti-extensive and self-dual filtering.

The attributes were mostly scalar [36], [37], allowing for threshold-based interaction. Recent effort were specifically conducted towards the development of geodesic attributes [38], designed for thin structures. Vectorial attributes have been less frequently used [39], due to more complex handling.

By construction, connected filters cannot split connected component, which may result in erroneous connections between tubes and artefacts, or between several branching tubular segments. Some attempts to minimize these drawbacks were proposed, either via thinning approaches [39], or with non-directed variants of component-trees [40].

It is worth mentioning that hybrid SE/connected strategies were proposed, e.g., in [41] for reconnection purposes. The links between connectivity-based and path-based approaches (described below) were also investigated in [42].

2.2.3 Path-based approaches

The SE- and connectivity-based approaches present dual intrinsic strengths and weaknesses. SE-based approaches can naturally handle anisotropy, which is highly desirable for linear structure filtering, but require explicitly defined families of SEs for orientation sampling. In comparison, connectivity-based approaches lead to more global descriptors; unfortunately the isotropic notion of adjacency cannot efficiently model the anisotropy of linear structures.

To address this problem, geodesic paths [43] were introduced to consider long-range, non local interactions while still coping with the constraints of thin objects, in particular noise. A thin object detector was proposed in [44] using geodesic voting, similar to path density. Polygonal path images [45] extended this idea allowing for better regularization and fewer artefacts. A recent formal discussion on optimal global paths was also proposed in [46]. All these solutions remained costly in 3D.

In [47], a notion of local optimal path was pioneered. The purpose was to restrict the research to a given distance, and in a given cone of orientations, in order to find the best paths starting from a given point. This paradigm led to the development of a notion of path operator [48]. These use a family of paths – i.e., connected sets – instead of a fixed shape as SE, thus enabling a higher flexibility in geometry and size, while preserving a 1D semantics. Algorithmic efforts were conducted to make such approach computationally efficient [49] and robust to noise [50], [51], leading to a notion of robust path opening. Sparse paradigms [52] were also proposed to avoid redundant computation.

3 PROPOSED FRAMEWORK

In this article, we aim to provide two new features characterizing tubular structures: an intensity feature which can be seen as a measure of presence of such a structure, and a directional feature providing at each point the orientation of these tubular structures. Here, we explain the distinguishing features of tubular structure, then we motivate the choice of the operator used in our method: the path opening, and we recall some basic notions.

3.1 General Strategy

Our strategy for distinguishing tubular structures from planar and blob structures is based on a simple geometric observation, illustrated in Fig. 1.

Let \mathcal{F} be any sort of oriented filter and \mathcal{O} be a set of chosen orientations such as $\mathcal{F}^o(I)$, $o \in \mathcal{O}$ is the response of this filter using orientation o on image I . Without loss of generality, we assume a bright structure on a dark background and a scale compatible with the size of the considered structure. As the name suggests, 1D thin (tubular) structures lie in fewer dimensions than 2D (planar) or 3D (blob) structures. Consequently, the




Fig. 1. In 3D, (a) a blob structure, (b) a planar structure, and (c) a tubular structure, in blue. Arrows show sampling along some orientation; a green (resp. red) arrow represents a high (resp. low) response of an oriented filter. The blob presents a high response in every orientations; the planar structure presents a high response in 4 out of the 8 orientations; the tubular structure only presents a high response in 1 orientation.

number of high responses among $\{\mathcal{F}^o(I), \forall o \in \mathcal{O}\}$ for an nD structure ($n > 1$) is higher than the number of high responses for a tubular structure. Therefore, counting the number of high responses of an oriented filter should allow us to distinguish each kind of nD structures.

Among the large choice of oriented filters in the literature, we chose the path operators. The main criterion we considered was their non-locality. Indeed, the majority of oriented filters compute the response of a structure in an isotropic neighborhood whose size depends on the scale, which itself depends on the size of the sought structure. This approach is not optimal for tubular structures which are highly anisotropic. In particular, it may lead to false detection and wrong orientation estimation, especially near structures borders. Path operators, by computing the response along an anisotropic neighborhood fitting in the tubular structure, avoid this pitfall.

Moreover, unlike classical openings, path openings can efficiently deal with local tortuosity. If a structure locally diverts from a chosen orientation, path openings can still detect this structure (see Fig. 2).

As the purpose of our operator is to provide features suitable for segmentation, an edge preserving filter is preferable. This excludes all the filters based on the Gaussian scale-space paradigm.

Finally, the path operators are nearly parameter-free. The only real parameter is the path length which is semantically meaningful as related to the length of the structure of interest.

3.2 Path operators

Path operators include the dual path openings and path closings. Without loss of generality, we will focus our explanations on path openings.

Intuitively, a path opening uses a set of oriented connected pixels of fixed length as a family of structuring elements. If no such SE fits entirely inside a structure at some local threshold level, the threshold value is decreased until at least one does. In this section we explain more formally this operator.




Fig. 2. Comparison on image (a) of a classical path opening (b) and a path opening (c) with the same SE length of 10 pixels. The bottom right line is globally vertical but presents local tortuosity. Only the path opening is able to detect such structure.




Fig. 3. (a) An example of adjacency relation \rightarrow . (b–d) The connectivity Γ which is a periodic repetition of (a) over X . (b) A path of length 3 (resp. 4) is depicted in red (resp. blue). (c) A binary image defined on X . (d) The result of a path opening of length 4 on the image of (c).

3.2.1 Paths

Let X be the support of an image and \rightarrow be an irreflexive, non-symmetric binary adjacency relation. For two points x and y of X , the relation $x \rightarrow y$ means that y is connected to x . We define a path π of length L as a set $\sigma(\pi)$ of L successively connected pixels (see Eq. (1)).

$$\sigma(\pi) = \{\{x_1, x_2, \dots, x_L\}, x_i \rightarrow x_{i+1}\} \quad \forall i \in \llbracket 1, L-1 \rrbracket \quad (1)$$

Practically, a path of length L is defined on a graph (X, Γ) , such that the connectivity Γ is the elementary adjacency relation \rightarrow periodically reproduced over X (see Fig. 3.(a,b)). We note $\Pi_L^\Gamma(X)$ the set of all paths of length L on X with connectivity Γ .

3.2.2 Binary and grey-level path opening

The binary path opening of length L with connectivity Γ is defined as the union of all paths of length L in X

(see Eq. (2)). This operator preserves each point of X belonging to at least one path of $\Pi_L^\Gamma(X)$ and removes the others; an example is shown in Fig. 3.(c,d).

$$\alpha_L^\Gamma(X) = \bigcup \{\sigma(\pi) \mid \pi \in \Pi_L^\Gamma(X)\} \quad (2)$$

Let I be a grey-level image with intensity $I(x) \in \mathbb{R}$ at each point $x \in X$. We note G the set of grey-levels of I . The extension of the path opening to grey-level is obtained by stacking the path openings on the thresholdings of I at every grey-level of G [48]. More formally, the thresholding of I at grey-level λ is noted $I_\lambda = \{x \in X \mid I(x) \geq \lambda\}$, and Eq. (3) defines the grey-level path opening. As openings, path openings are increasing, anti-extensive and idempotent.

$$A_L^\Gamma(I(x)) = \max\{\lambda \mid x \in \alpha_L^\Gamma(I_\lambda)\} \quad (3)$$

3.2.3 Orientation-space sampling

A path opening is defined over an adjacency relation, which provides a general orientation. A path opening preserves the structures compatible with this orientation. In order to preserve structures in all orientations, several path openings with n different orientations must be combined. Consequently, a space sampling is required.

In this article, we chose to consider a set of $n = 7$ orientations: the 3 main orientations corresponding to the vectors of the orthogonal basis $\{0, 0, 1\}$, $\{0, 1, 0\}$ and $\{1, 0, 0\}$ denoted $\{e_1, e_2, e_3\}$, and the 4 diagonals corresponding to $\{1, 1, 1\}$, $\{1, 1, -1\}$, $\{1, -1, 1\}$, $\{-1, 1, -1\}$ denoted $\{d_1, d_2, d_3, d_4\}$.

There is a degree of arbitrariness to this choice. Indeed it depends on some of the underlying properties of the path operator, like the angular aperture of the cones we use, themselves dependent on the adjacency relation we choose. However, this choice is motivated in Sec. 5. In the remainder we will focus our description using only these 7 orientations, for simplicity.

In order to cover the whole space using only these 7 vectors, our adjacency relation consists of a 3D discrete cone centered on them. We associate to e_* (resp. d_*) the cone \mathcal{C}_e (resp. \mathcal{C}_d) bounded by the vectors $\{d_1, d_2, d_3, d_4\}$ (resp. $\{e_1, e_2, e_3\}$). We note \mathcal{C} the set of the 7 cones corresponding to the 7 orientations (see Fig. 4). In the remainder of this article, $A_L^c(I)$, $c \in \mathcal{C}$ denotes the path opening of length L with the orientation c .

4 RORPO: RANKING ORIENTATION RESPONSES OF PATH OPERATORS

So far we have presented the purpose of our work: characterizing tubular structures; and the strategy we have developed: detecting these structures by counting the number of high responses of an oriented filter, the path operators. In this section, we first present our methodology (Sec. 4.1) and how we derive our intensity (Sec. 4.2) and directional (Sec. 4.3) features from it.




Fig. 4. The 7 orientations of \mathcal{C} . The center of the cube, the centers of faces, and corners represent points of the image. The blue arrows represent the adjacency relation.

4.1 Methodology

Path operators, as an oriented filter, preserve structures compatible with specific orientations. However, this is not sufficient to distinguish tubular from planar and blob structures. For instance, a path opening on a 2D image containing an isotropic object (a blob) and a line-like object with similar diameter/length will preserve both. The 3D case is even harder since we have to take into account planar structures, which are also detected by a path opening.

To tackle this problem, we propose to count the number of high responses of path openings. This is done by pointwise ranking the 7 path openings responses. Let $\gamma_L^i(I)$ be the image obtained by applying a pointwise rank filter, RF_i , of order i (see Eq. (4)). In particular, RF_1 , RF_4 and RF_7 are respectively the pointwise maximum, median and minimum operators.

$$\gamma_L^i(I(x)) = RF_i\{A_L^c(I(x)), c \in \mathcal{C}\} \quad (4)$$

From the 7 responses of the path opening, $\{A_L^c(I)\}_{c \in \mathcal{C}}$, we have built the 7 ranked filtered images $\{\gamma_L^i(I)\}_{i \in \llbracket 1, 7 \rrbracket}$. Let us consider a structure appearing in n ($n \leq 7$) of the 7 $A_L^c(I)$; then this structure now appears in $\gamma_L^i(I)_{i \in \llbracket 1, n \rrbracket}$. In particular, the image $\gamma_L^3(I)$ contains all the structures detected in at least 3 orientations. An illustration in the 2D case is shown on Fig. 5.

4.2 Filtering tubular structures

We have seen that tubular structures are detected in fewer path opening orientations than other structures. In the previous section, we have also proposed a method to count the number of orientations preserving a structure.

Let i_t be the maximal number of path opening orientations preserving a tubular structure. Then, we define the RORPO filter, Φ^L as follows:

$$\Phi_L(I) = \gamma_L^1(I) - \gamma_L^{i_t+1}(I) \quad (5)$$

γ_L^1 contains structures preserved in all orientations and $\gamma_L^{i_t+1}$ contains no tubular structure by definition of i_t .




Fig. 5. 2D illustration of filter response ranking. A structure present in one orientation (the vertical line) remains in RF_1 ; structures present in two orientations (the oblique lines) remain in RF_1 and RF_2 ; the disc is present in all the orientations and thus remains in all the RF s.

All other structures with a diameter/thickness greater than L are preserved in more orientations than tubular structures. Consequently, the residual operator $\Phi_L(I)$ reduces the signal of all structures except for those detected in i_t orientations, and these are the tubular structures.

The orientation threshold value i_t depends on the chosen orientation space sampling. With the one we selected (see Sec. 3.2.3), up to some limit cases we will deal with later, tubular structures can only be preserved by up to 3 orientations. We confirmed this geometric result experimentally by exhaustively analysing a dense sampling of the orientation space using randomly generated tubular structures. This experiment is fully developed in Sec. 5.2.3. Based on these results, we selected $i_t = 3$.

4.3 Direction estimation of tubular structures

Since RORPO is based on oriented path operators, it intrinsically carries information about the direction of the detected tubular structures. In order to retrieve this directional information, we have to identify which paths opening orientations fit best at each point. Typically there will be more than one, so by combining these orientations, we may obtain a reasonable evaluation of the tubular structure orientation.

Finding which orientations detect a tubular structure requires, for each orientation, a binary decision: either an orientation detects the tubular structure or not. However, the output of each path opening filter is a grey-level response. A solution with a threshold t such that if $A_L^\Gamma(I(x)) > t$, orientation Γ detects the tubular structure would be simple but flawed. Indeed the intensity of the responses depends on the initial grey-level of the image and the degree of noise of the tubular structure.




Fig. 6. At each point of an image (red square), 7 path opening responses $A_L^\Gamma(I(x))$, associated to each orientation C_* , are computed (top right values). After ranking, the orientations of interest are the first one, two or three orientations.

Consequently, we propose to determine which path opening orientations detect a given tubular structure by separating the 7 orientations in two classes: Those that detect the tubular structures (called orientations of interest) and the others. Our expectation is to achieve homogeneous responses $A_L^\Gamma(I(x))$ inside each class: a high response in the orientations of interest and a low one in the others. (see Fig. 6). For simplicity and efficiency, we chose the standard deviation at the homogeneity criterion.

The classification requires little extra computation. Indeed the 7 $A_L^\Gamma(I)$ have already been ranked into the $\gamma_L^i(I)_{i \in [1,7]}$ from which we can derive the ranked orientations $(O_L^i)_{i \in [1,7]}$ such that O_L^i is the orientation associated to γ_L^i . Moreover, we know that a tubular structure is detected in at most 3 orientations, so if a structure is detected in k orientations, $1 \leq k \leq 3$, the orientations of interest are the k $(O_L^i)_{i \in [1,k]}$. This results in 3 possible classes: $\{\{O_L^1\}, \{O_L^1, O_L^2\}, \{O_L^1, O_L^2, O_L^3\}\}$. Therefore, for each pixel, we only have to compute the standard deviation of these 3 possible classes and choose the class with the lowest (see Alg. 1). Finally, the direction of the structure is estimated by the mean of the orientations in the selected class.

Orientation vs. direction: An important point to note is that we encode a tubular structure orientation by a 3-vector. However, a vector has a *direction* which is more specific than an *orientation*. If a structure is horizontal, both vectors $[0, 0, 1]$ and $[0, 0, -1]$ encode its orientation. When averaging path opening orientations, as we use the vectors e_* and d_* , we must ensure that all the vectors are set to the same half-space, e.g. the mean of the vectors $[0, 1, 0]$, $[0, 0, 1]$ and $[1, -1, -1]$ is not $[\frac{1}{3}, 0, 0]$ but $[-\frac{1}{3}, \frac{2}{3}, \frac{2}{3}]$.

4.4 Multiscale length analysis

One of the difficulties in tubular structure analysis is dealing with multiple scales. Real applications need to cope with varying diameter, length and curvature, which are generally highly correlated. To tackle this issue,

Algorithm 1: FindSetOrientations

Data: $(\gamma_L^i)_{i \in [1,7]}$, and $(O_L^i)_{i \in [1,7]}$

Result: d : image of directions

begin

for each pixel x **of** I **do**

 # Compute the standard deviation of the 3 possible classes

$g_1 = \text{std}((\gamma_L^i(x))_{i \in [2,7]})$

$g_2 = \text{std}((\gamma_L^i(x))_{i \in [1,2]}) + \text{std}((\gamma_L^i(x))_{i \in [3,7]})$

$g_3 = \text{std}((\gamma_L^i(x))_{i \in [1,3]}) + \text{std}((\gamma_L^i(x))_{i \in [4,7]})$

 # Choose the lowest class

$g^* = \underset{i \in [1,3]}{\text{argmin}}(g_i)$

 # Compute the final direction

$d(x) = \text{mean}((O_L^i)_{i \in [1, g^*]})$

end

end

multiscale approaches have been developed. A common multiscale approach consists of applying the filter multiple times while changing scale-related parameters, and merging the results. In this section, we propose a multiscale version of our operator based on the length of the structure.

In the literature on multiscale tubular object analysis, the scale parameter is usually a diameter. However, in the case of RORPO, length is the only one available. Had we wanted to vary the diameter, we would have had to combine RORPO with another filter. Fortunately, in many applications, the diameter, curvature and length of tubular structures are highly correlated. For example small blood vessels, are generally more tortuous and shorter than large vessels like the aorta. The same argument can be made for insulation glass fibers or country roads vs. highway.

With RORPO, path lengths depend on the curvature of the structures. Indeed, to be retained, a structure of a given length must stay within a single cone, which puts a limit on the degree of its large-scale curvature. Consequently, a tubular structure with a high large-scale curvature can only be detected with a smaller path length than if it were straight. This consideration is only relevant for large-scale curvature, since small-scale curvature is already handled by path flexibility. The reader may refer to Sec. 5.2.1 for more explanations about path length.

More formally, let $S = \{L_1, L_2, \dots, L_n\}$ be a set of path lengths. The multiscale RORPO filter, $\Phi(I)$, is obtained by taking the maximum of each scale response $\Phi_L(I)$:

$$\Phi(I) = \bigcup_{L \in S} \Phi_L(I) \quad (6)$$

This multiscale paradigm also applies to the directional feature. Indeed, a direction can be computed for




Fig. 7. Illustration of our proposed simplified robust path opening (see Sec. 5.1).

each scale. The final direction is the one associated with the highest RORPO response with the lowest scale.

5 ALGORITHMIC CONSIDERATIONS

5.1 Simplified robust path opening

RORPO is based on the path opening. It then inherits the main weakness of path operators: a relative fragility to disconnections induced by noise, a fortiori for high values of L . Several works were conducted to propose various noise resistant versions of path operators [48], [50], in particular the robust path opening (RPO) of Cokelaer et al. [51]. RPO allows for K disconnections (noisy pixels) between two successive pixels path.

RPO suffers from 2 main drawbacks: 1) the noisy pixels are not taken into account in the overall path length calculation; so the parameter L may become almost meaningless. 2) The algorithmic layer handling the robustness is time and memory consuming.

We propose an alternative to RPO which yields similar results but is less time consuming and preserves the real path length. Our method relies on a mask-based second-generation connectivity strategy [53] in order to reconnect the noisy parts of tubular structures. A dilation by a cubical structuring element of size N is performed on the initial image I . This dilated image is used to compute the regular path opening. In order to preserve the anti-extensivity of the path opening, an infimum operator is applied (see Eq. (7)). An illustration of this strategy is shown in Fig. 7.

$$A_{L,N}^c = \bigwedge \{I, A_L^c(\delta_N(I))\} \quad (7)$$

5.2 Parameters

RORPO only requires a few parameters compared to other similar filters. In this section, we explain why the path length L is the only tunable parameter. Then we analyse why and how the other parameters are set.




Fig. 8. The path length also depends on the scale of the curvature. (a) A thin structure of length L_r presenting small-scale and large-scale curvature. (b) A path length of length $L_1 < L_r$ along one orientation detects a first part of the structure. (c) A path length of length $L_2 < L_r$ along another orientation detects a second part of the structure. (d) The two detected parts overlap.

5.2.1 Path length

Path length intuitively corresponds to the length of a path that lies in a given orientation. As evoked in Sec. 4.4, the path length carries both length and curvature information. A tubular structure with a high large-scale curvature will be detected in several orientations, each detecting a part of the complete structure. As the orientations overlap, the parts generally overlap as well (see Fig. 8).

5.2.2 Robustness parameter

In our simplified version of path opening robustness (see Sec. 5.1), the robustness parameter N is the size of the cubical structuring element used for the dilation of the filtered image. This parameter corresponds to the maximal number of noisy pixels in a row allowed in a path, which is exactly $N - 1$. We do not consider the robustness parameter tunable, given that in practice, we recommend $N = 3$ or $N = 5$. For higher values, false detections start appearing and the background noise reduction is less effective.

5.2.3 Orientations sampling

Working with a finite number of orientations implies to choose a sampling policy, in order to determine the number and the shape of the orientations for the computation of the path openings. Considering the structure of \mathbb{Z}^3 , isotropy requirements, and the algorithmic constraints of path opening, three main families of sampling policies can be considered: along the 3 principal orientations; along the 3+4 principal orientations and principal diagonals; and along the 3+4+6 principal orientations and principal/secondary diagonals.

As stated earlier, 3 orientations are not sufficient to accurately capture tubular structures due to the necessity for orientations to overlap. The 4 principal diagonals seem essential and offer a high degree of overlap with the 3 principal orientations. Our working hypothesis was that, with 3+4=7 orientations, such as described in




Fig. 9. Number of orientations, computed from 100 synthetic samples, detecting (a) tubes, (b) planes, and (c) curved planes, within the set of 7 orientations illustrated in Fig. 4.

Section 3.2.3, a quantitative analysis of the path openings responses should be sufficient to decide whether a point belongs to a tubular structure. In order to experimentally verify this hypothesis, we computed the number of high responses within the 7 orientations for 100 synthetic binary tubes, planes and curved planes (half ellipsoid). The results of this study (Fig. 9) validate this conjecture. Indeed, nearly all tubular structures are detected in at most 3 orientations whereas all planar structures are detected in at least 4 orientations.

Finally, two main reasons led us to reject the possibility of using all 13 orientations: the multiplication of angular cones induces the handling of considerably more limit cases; and the computation time is nearly doubled, for a questionable benefit.

5.2.4 Angular cones

The basic patterns for each of the 7 orientations need to fully cover the immediate neighbourhood of any point $\mathbf{x} = (x, y, z)$ of \mathbb{Z}^3 , namely, the 26 points forming a $3 \times 3 \times 3$ cube around \mathbf{x} . Two policies may be considered: choosing patterns that induce either a partition or a cover of these 26 points. A partition is not acceptable as it would eliminate all structures with an orientation lying between the bounds of two neighbors orientations. A cover is then the only choice.

In order to respect isotropy requirements and minimal overlapping, this cover was defined as illustrated in Fig. 4. The drawback of any covering policy is the existence of limit cases, corresponding to the paths that lie exactly at the frontier between three or more orientations. To solve this problem, we devised a virtually cost-free solution, described in details in [2]. The essentials of this solution are given in Appendix.

5.3 Computational cost

The computation cost of RORPO is dominated by the path openings. Indeed, the robustness step (Eq. (7)), ranking (Eq. (4)) and limit case handling only require infimum / supremum operations which have linear cost $\mathcal{O}(|\Omega|)$ with respect to the size of the image, Ω .

Two different algorithms were proposed to compute path openings: the Talbot algorithm [50] and the Luengo




Fig. 10. Computational cost of RORPO, with respect to $|\Omega|$ (left) and L (right, log scale).

algorithm [49]. The Talbot algorithm has a $\mathcal{O}(|\Omega|\log(L))$ complexity. We verified experimentally on synthetic images [54] (see Fig. 10) that the Luengo algorithm is equivalent. The Luengo algorithm is slightly slower and is not natively robust, however it is more generic for defining orientations, and with simplified robustification (see Sec. 5.1), native robustness does not matter. Consequently we used the Luengo algorithm.

6 EXPERIMENTS AND RESULTS

In this section we present results of both the RORPO features (intensity and direction). Then we quantitatively evaluate and compare to gold standard methods.

6.1 Evaluation and comparison of the intensity feature

6.1.1 Evaluation criteria

For the quantitative evaluation of our method, we used the standard Dice coefficient and the Matthews Correlation Coefficient (MCC) which is well adapted to sparse images, i.e. when the structures of interest only represent a small proportion of the total image size. These coefficients are defined as follows:

$$\text{Dice} = \frac{2 \text{TP}}{\text{TP} + \text{FN} + \text{TN} + \text{FP}}$$

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP}) \times (\text{TP} + \text{FN}) \times (\text{TN} + \text{FP}) \times (\text{TN} + \text{FN})}}$$

With TP/TN the true positives/negatives and FP/FN the false positives/negatives.

The classical manner for evaluating a grey-level filtering method is to compare all the possible thresholding results to the ground truth. Results are represented as a curve: the Receiver Operating Characteristic (ROC). We chose to compute ROCs of the true positive rate (TPR) vs. the false positive rate (FPR). Then, the closer to the point $(0, 1)$ the better the results.

As we deal with sparse features, the set of voxels belonging to the ground truth objects (TP_{GT}) is always much smaller than the set of voxels belonging to its background (TN_{GT}). Consequently, the number of false positives is much larger than the number of true positives. To present meaningful results, we defined the TPR and FPR with respect to the ground truth objects, i.e. $\text{TPR} = \frac{\text{TP}}{\text{TP}_{\text{GT}}}$ and $\text{FPR} = \frac{\text{FP}}{\text{TP}_{\text{GT}}}$. With this convention, the FPR can exceed 100%.




Fig. 11. (a–c) Synthetic image: (a) isosurface rendering; (b) ground truth; (c) maximum intensity projection. (d–f) Filtered synthetic image: maximum intensity projection. (d) RORPO; (e) Frangi's vesselness; (f) and RPO-top-hat.

6.1.2 Comparison methods

As indicated in Sec. 2, most filtering methods for tubular structures use a multiscale eigenvalue analysis of the Hessian tensor. In recent comparison [5], [55], the gold standard in tubular structure filtering is the Frangi Vesselness (FV). We also wanted to compare it to a reference method in mathematical morphology: a combination of RPO and a top-hat operator.

6.1.3 Synthetic Image

In a first series of experiments, we use a synthetic image, see Fig. 7.(a,b) containing isotropic structures synthesized with an additive Gaussian random field, a thin rectangle as a planar structure and a tubular structure: an helix with varying diameter and thickness. The ground truth is a binary image of the helix (see Fig. 11(b)).

We computed RORPO, the Frangi Vesselness and the RPO + top hat filters on this image. Results and scores are presented in Fig. 11 and the ROC curve in Fig. 13(a). Qualitatively, we observe that RORPO is the only filter that fully removes the rectangle. Frangi Vesselness removes the interior of the rectangle but the borders

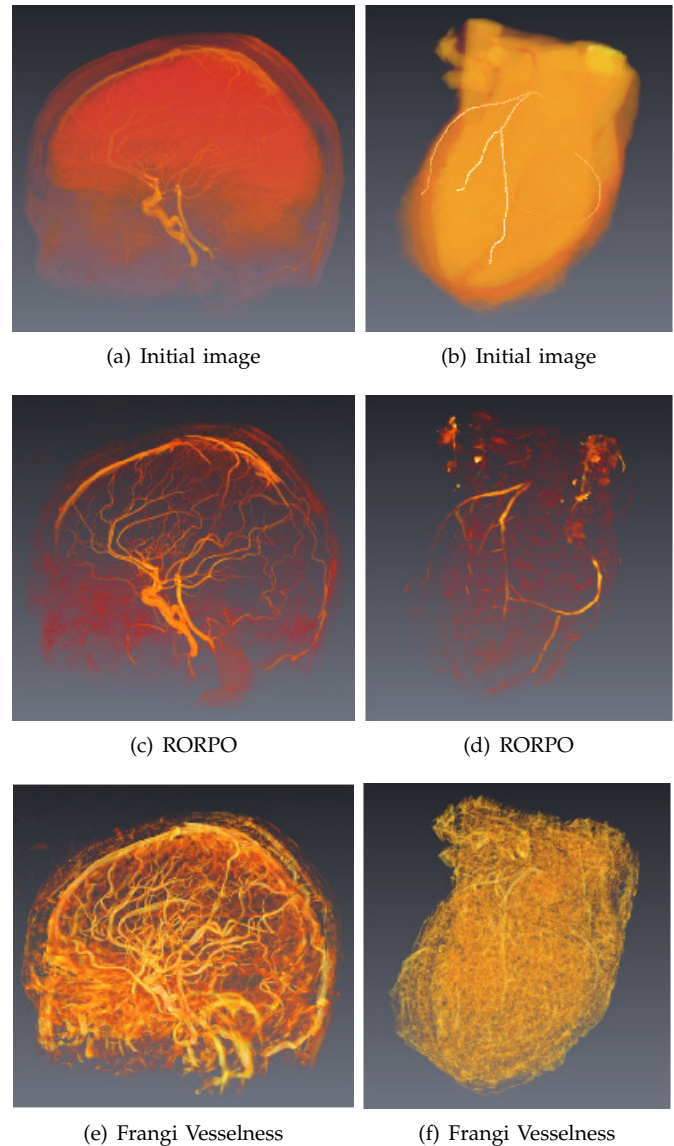


Fig. 12. (a,c,e) Brain arteries. (a) Time-of-flight magnetic resonance angiography volume rendering (from low intensities, in red, to high intensities, in yellow). (c) RORPO: 8 lengths from $L = 30\text{px}$ to 130px ; computation time: 32mn. (e) Frangi Vesselness: 4 scales from 1 to 3; computation time: 13mn. (b,d,f) Coronary arteries. (b) Computed tomography angiography, viewed in volume rendering (from low intensities, in red, to high intensities, in yellow); the coronary centerlines are outlined in white. (d) RORPO filter: 4 scales from $L = 20\text{px}$ to 50px ; computation time: 50mn. (f) Frangi Vesselness: 4 scales from 0.4 to 1.8; computation time: 6mn.

are retained. Hessian eigen-analysis cannot easily distinguish a tubular object from the border of a planar object, which is a known limitation. The RPO + top hat does not distinguish between tubular and planar objects, and so fully preserved the plane as expected. These limitations were among those we wanted to address with RORPO.

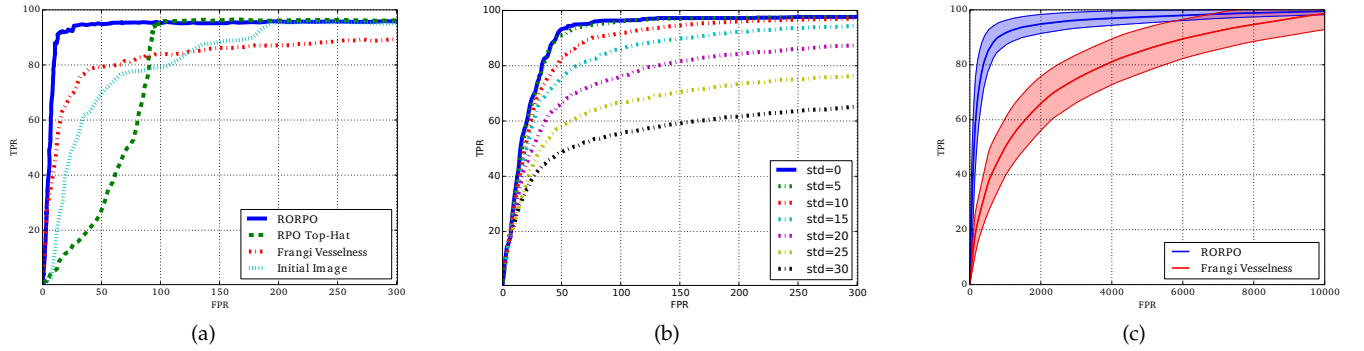


Fig. 13. ROC curves on (a,b) synthetic data, and (c) real data. (a) Comparison of the three filters, plus the native image. (b) Noise robustness of the RORPO filter. (c) Comparison between RORPO and Frangi Vesselness on the Rotterdam repository (average ROC curve \pm one standard deviation).

6.1.4 Real Images

Next, we compared the results of RORPO with Frangi vesselness on real 3D images. We considered two types of images :

- Time-of-Flight Magnetic Resonance Angiography (MRA) images of the brain vascular network acquired on a whole body scanner (Siemens Magnetom Verio 3.0T);
- Computed Tomography Angiography (CTA) images of the heart from the Rotterdam repository [56] used in a 2012 MICCAI Challenge.

The analysis presented here is only qualitative for the MRA data because, to the best of our knowledge, no ground truth is available for such complicated networks; for the CTA, both qualitative and quantitative analyses are presented as the Rotterdam data contains the centerlines of the coronaries as ground truth.

Visually, on the MRA image (Fig. 12), the RORPO results seems nicer as it detects less false positives than FV. Indeed, the Frangi Vesselness tends to detect more false positives and less false negatives. For example, FV, due to the Gaussian kernel convolution, overestimates the volume of the vessels leading to false positives, while RORPO preserves the initial geometry. These trends are confirmed by the quantitative evaluation on the CTA images. Indeed, CTA images are denser leading to even more false positives for the Frangi Vesselness.

The Rotterdam repository provides 17 different clinical CT angiography images. We eliminated 2 images due to poor ground truth and we computed both RORPO and FV on the remaining 15 images. One result is shown on Fig. 12, and the average \pm 1 standard deviation of the 15 ROC curves is presented in Fig. 13.(c).

The CTA Frangi Vesselness results are even less satisfactory than the MRA; the still detected coronaries are surrounded by a lot of false positives, which prevent a good visualization. In contrast, RORPO efficiently remove the heart tissues while preserving the coronaries. The ROC curves show that the TPR of Frangi Vesselness is significantly worse than that of RORPO for the entire

exploitable range of false positives.

6.2 Evaluation of the directional feature

Whereas some ground truth is publicly available for tubular structure segmentation, to the best of our knowledge, no such ground truth for directions exists. Consequently, we mostly evaluated the directional feature on synthetic images where we can generate the direction ground truth.

6.2.1 Evaluation criteria

We evaluate directions by estimating the angle between the ground truth and the RORPO direction at each voxel included in the segmentation ground truth. I.e, if a voxel belongs to a tubular structure, we compute the angle between the ground truth and the RORPO directions at this voxel, which provides an error measure. Then, we observe the distribution of the errors over the whole image.

One should note that the maximum error between two angle is 90° as we compute an orientation difference (i.e. an angle between lines).

We also need to take into account the false positive rate (FPR) and false negative rate (FNR) in this evaluation. While they indicate the quality of the intensity feature, they also matter for the evaluation of the orientation. Indeed, the direction is only relevant where the intensity feature is high (i.e. where the structure is tubular). Consequently, if the FNR is high, we incorrectly omit many directions; and if the FPR is high, orientations are incorrectly computed on isotropic structures.

6.2.2 Evaluation of the intrinsic directional error

The directional feature is computed only from the 7 orientation vectors e_* and d_* . Consequently, it can never be very precise.

Our first experiment is designed to evaluate this precision. We sampled a half sphere homogeneously with 1300 orientations. In each of these, we generated a tube of length 50 pixels and diameter around 11 pixels. Then

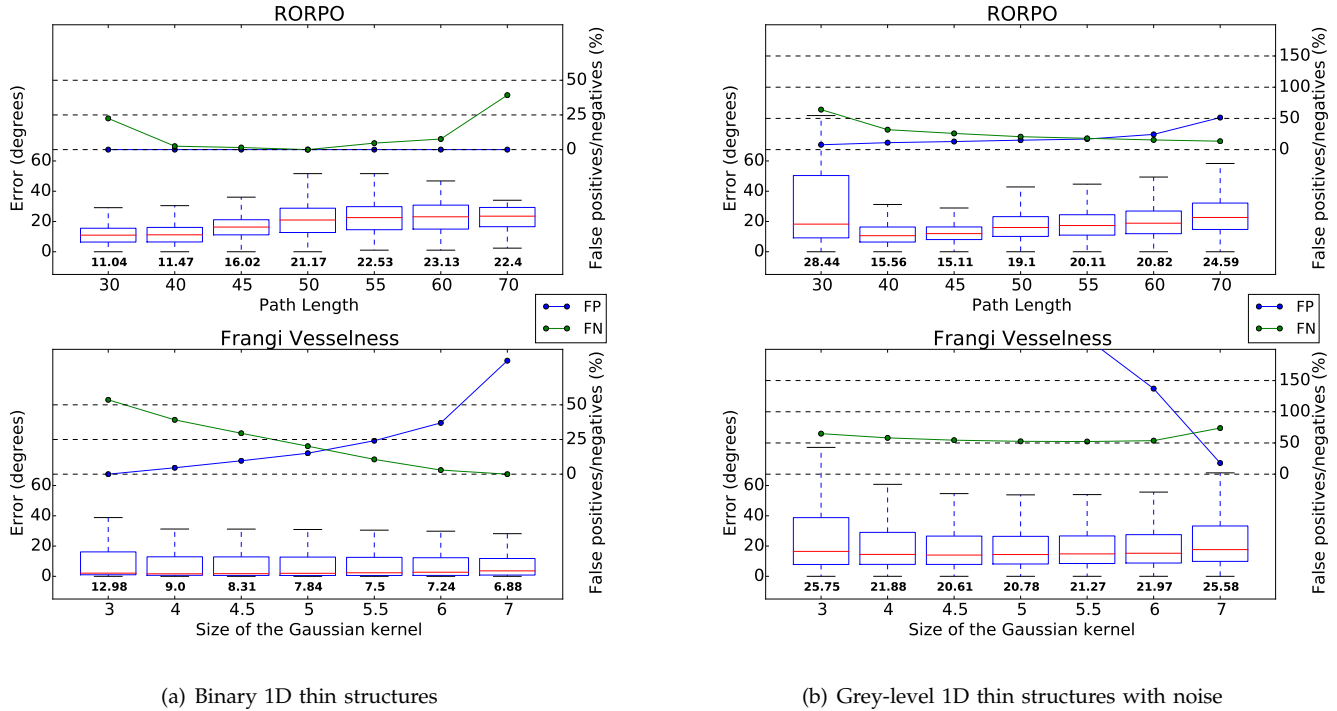


Fig. 14. Comparison of the direction feature by RORPO and Frangi Vesselness on binary tubular structures (a) without and (b) with noise (b). The box plot shows the distribution of errors (in degrees) over the 1300 samples for each parameter of RORPO (path length) or Frangi Vesselness (size of the Gaussian kernel). The green (resp. blue) curves are the rate of false negatives (resp. positives) compared with the size of the 1D thin structure (e.g., FP = 100% means that as many true positives as false positives have been detected).

we applied RORPO on each binary tube and we computed the orientation error with respect to the generated ground truth.

Using a path length of 40, we obtained a mean error of 12 degrees with a standard deviation of 7. These results are encouraging given the small number of path opening orientations used.

One characteristic of our method is that the computed directions tend to be homogeneous forming large zones with the same direction. This tends to limit outliers due to noise, but causes sharp direction transitions even where a tubular structure is smoothly curved. To enhance the results, we propose to apply a local averaging filter as a post processing on the directional feature using a $7 \times 7 \times 7$ window. In the following experiments we always use this post processing.

6.2.3 Comparison with Frangi Vesselness

We have established that the intrinsic directional error of our method is low, despite the small number of orientations we used. We now compare our directional feature with that of the Frangi Vesselness, first on binary tubes, then on grey level tubes plus noise and finally, we compare both features on a real application.

In our first experiment, we compare Frangi Vesselness vs. RORPO on the same binary tubes as in the previous

section. The directions were estimated for several scale parameters (path length for RORPO and size of the Gaussian kernel for FV) and the angle error, the FPR and the FNR were measured in each case. Then, we performed the same experiment but with Gaussian noise and small isotropic structures in the background (see Fig. 15).

Results shown in Fig. 14 indicate that the Frangi Vesselness generally estimates a more accurate direction than our method in the perfect case (i.e. a noise-free binary tube). However, when noise and other structures are present, the Frangi Vesselness results are significantly worse than ours. Indeed, as RORPO uses paths, which are semi global structures. The RORPO directions integrate the orientation information over the whole path. It appears that paths, by their non-locality and anisotropy, are a more meaningful neighborhood for the analysis of tubular structures than the isotropic neighborhood used by Frangi. We can observe that the RORPO length parameter is less sensitive than the Frangi Vesselness diameter parameter. Indeed, a large range of RORPO parameters give similar, good results whereas only one scale of Frangi Vesselness yields the best results. This stability with respect to the parameters is important in practice because one generally does not know the best set of parameter for a given application.




Fig. 15. Projection (MIP) of the noisy image (a) and the directional feature from RORPO (b) and Frangi Vesselness (c). Median and standard deviation of the noise are shown below each results.




Fig. 16. Directional feature of RORPO (a) and Frangi Vesselness (b) on the segmentation ground truth (in pink).

To the best of our knowledge, no freely available ground truth for directions exists for real 3D images. The company HeartFlowtm [57] kindly provided us with a CTA exam plus their manually corrected central line and segmentation of the coronaries. From these data we were able to generate good enough directional ground truth to validate our results.

With optimal parameters, we computed both RORPO

	mean error	std error	FP	FN
RORPO	18.15	11.21	37.36	71.32
FV	14.85	13.31	566	39.48

Fig. 17. Evaluation of the directional features of RORPO and FV on a CTA image.

and FV directional features on this exam and evaluated the angle error and FPR/FNR. A selected example of results is shown on Fig. 16 and quantitative results of the comparison are shown on Fig. 17. We see that FV performs slightly better on the accuracy of directions. However it is important to note that this error is computed only for the voxels inside the segmentation ground truth. That means that all the false positive directions detected are not taken into account in the error. Indeed, the Frangi Vesselness provides fourteen times more false directions on isotropic structures. In contrast, RORPO provides a slightly worse error (18° vs. 14°) for the blood vessel but also compute far fewer false positives directions.

In real image processing applications, a directional feature can usually be used at two different stage: either to guide a segmentation method directly on the grey-level image or to guide a post processing pipeline on a segmentation image. In the first case, RORPO provides better results than FV as it computes directions with a similar error but much less false positives, and so resulting in generally more accurate directions. In the second case, when the segmentation is already available, the Frangi Vesselness should be used, but it was not the problem we sought to address.

7 CONCLUSION

In this article, we have proposed RORPO, a new framework for the characterization of tubular structures. It consists of an intensity feature which measures the presence of a tubular structure and a directional feature providing at each point the orientation of this tubular structure. We have shown that RORPO is a convincing alternative to Hessian-based filters which were, until now, the reference framework proposing both intensity and directional features for grey level images. Indeed, the Hessian-based filters are based on local and isotropic eigen-analysis whereas tubular structures are intrinsically anisotropic. RORPO, by using path operators, considers non local and anisotropic neighborhoods which are better adapted to tubular structures.

Experimental quantitative and qualitative evaluations show that both of our features give better results for the analysis of tubular structures. The intensity feature is more reliable as it detects much less false positives and the directional feature is much more robust as it relies on non local SE.

As a low level operator for the characterization of tubular structures, RORPO can be used as prior information in segmentation frameworks in the same way as the Hessian. As future work, we plan on embedding RORPO features in a learning-based framework for segmentation, and we are working on accelerating the computations.

RESOURCES

Free, open-source C++ implementation, test data and documentation is available at <http://path-openings.github.io/RORPO>.

ACKNOWLEDGMENTS

This work was partially funded with French *Agence Nationale de la Recherche* grant agreement ANR-12-MONU-0010 (VIVABRAIN Project: <http://vivabrain.fr>). MRA data was provided by the In Vivo Imaging Platform at Strasbourg University (<http://piiv.u-strasbg.fr>). CTA data for the intensity feature experiments was provided by the Rotterdam Coronary Artery Algorithm Evaluation Framework (<http://coronary.bigr.nl>). The CTA and ground truth for the directional feature experiments was provided by HeartFlow (<http://hearflow.com>).

REFERENCES

- [1] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," in *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, ser. Lect. Notes Comput. Sci., vol. 1496. Springer, 1998, pp. 130–137.
- [2] O. Merveille, H. Talbot, L. Najman, and N. Passat, "Ranking orientation responses of path operators: Motivations, choices and algorithmics," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 9082. Springer, 2015, pp. 633–644.
- [3] —, "Tubular structure filtering by ranking orientation responses of path operators," in *Proc. Eur. Conf. Comput. Vis.*, ser. Lect. Notes Comput. Sci., vol. 8690. Springer, 2014, pp. 203–218.
- [4] H. Talbot, "Étude des directions en analyse d'image," Habilitation Thesis, Université Paris-Est, 2013, (In English).
- [5] D. Lesage, E. D. Angelini, I. Bloch, and G. Funka-Lea, "A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes," *Med. Image Anal.*, vol. 13, no. 6, pp. 819–845, 2009.
- [6] Y. Sato *et al.*, "3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images," in *Proc. joint CVRMed-MRCAS*, ser. Lect. Notes Comput. Sci., vol. 1205. Springer, 1997, pp. 213–222.
- [7] C. Lorenz, I.-C. Carlsen, T. M. Buzug, C. Fassnacht, and J. Weese, "Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2D and 3D medical images," in *Proc. joint CVRMed-MRCAS*, ser. Lect. Notes Comput. Sci., vol. 1205. Springer, 1997, pp. 233–242.
- [8] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Troussset, "Model-based detection of tubular structures in 3D images," *Comput. Vis. Image Understand.*, vol. 80, no. 2, pp. 130–171, 2000.
- [9] P.-E. Danielsson, Q. Lin, and Q.-Z. Ye, "Efficient detection of second-degree variations in 2D and 3D images," *J. Vis. Comm. Image Repr.*, vol. 12, no. 3, pp. 255–305, 2001.
- [10] R. Manniesing, M. A. Viergever, and W. J. Niessen, "Vessel enhancing diffusion: A scale space representation of vessel structures," *Med. Image Anal.*, vol. 10, no. 6, pp. 815–825, 2006.
- [11] H. E. Bennink *et al.*, "A novel 3D multi-scale lineness filter for vessel detection," in *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, ser. Lect. Notes Comput. Sci., vol. 4792. Springer, 2007, pp. 436–443.
- [12] C. Xiao *et al.*, "A strain energy filter for 3D vessel enhancement with application to pulmonary CT images," *Med. Image Anal.*, vol. 15, no. 1, pp. 112–124, 2011.
- [13] C. Bauer and H. Bischof, "A novel approach for detection of tubular objects and its application to medical image analysis," in *Pattern Recogn.*, ser. Lect. Notes Comput. Sci., vol. 5096. Springer, 2008, pp. 163–172.
- [14] C. Xiao, M. Staring, Y. Wang, D. P. Shamonin, and B. C. Stoel, "Multiscale bi-Gaussian filter for adjacent curvilinear structures detection with application to vasculature images," *IEEE Trans. Image Process.*, vol. 22, no. 1, pp. 174–88, 2013.
- [15] M. W. K. Law and A. C. S. Chung, "Three dimensional curvilinear structure detection using optimally oriented flux," in *Proc. Eur. Conf. Comput. Vis.*, ser. Lect. Notes Comput. Sci., vol. 5305. Springer, 2008, pp. 368–382.
- [16] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906, 1991.
- [17] T. M. Koller, G. Gerig, G. Székely, and D. Dettwiler, "Multiscale detection of curvilinear structures in 2D and 3D image data," in *Proc. IEEE Int. Conf. Comput. Vis.* IEEE, 1995, pp. 864–869.
- [18] K. G. Derpanis and J. M. Gryn, "Three-dimensional nth derivative of Gaussian separable steerable filters," in *Proc. Int. Conf. Image Process.*, vol. 3. IEEE, 2005, pp. 553–556.
- [19] G. González, F. Aguet, F. Fleuret, M. Unser, and P. Fua, "Steerable features for statistical 3D dendrite detection," in *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, ser. Lect. Notes Comput. Sci., vol. 5762. Springer, 2009, pp. 625–632.
- [20] M. Őrkisz, M. Hernández-Hoyos, P. Douek, and I. Magnin, "Advances of blood vessel morphology analysis in 3D Magnetic Resonance Images," *Mach. Graphics Vision*, vol. 9, no. 1–2, pp. 463–472, 2000.
- [21] Y. P. Du, D. L. Parker, and W. L. Davis, "Vessel enhancement filtering in three-dimensional MR angiography," *J. Magn. Reson. Imaging*, vol. 5, no. 3, pp. 353–359, 1995.
- [22] B. É. Chapman and D. L. Parker, "3D multi-scale vessel enhancement-filtering based on curvature measurements: Application to time-of-flight MRA," *Med. Image Anal.*, vol. 9, no. 3, pp. 191–208, 2005.
- [23] L. Najman and H. Talbot, Eds., *Mathematical morphology: From theory to applications*. ISTE/John Wiley & Sons, 2010.
- [24] F. Zana and J.-C. Klein, "Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation," *IEEE Trans. Image Process.*, vol. 10, no. 7, pp. 1010–1019, 2001.
- [25] O. Tankyevych, H. Talbot, and P. Dokládál, "Curvilinear morpho-Hessian filter," in *Proc. IEEE Int. Symp. Biomed. Imaging*. IEEE, 2008, pp. 1011–1014.
- [26] P. Soille, E. Breen, and R. Jones, "Recursive implementation of erosions and dilations along discrete lines at arbitrary angles," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 5, pp. 562–567, 1996.
- [27] P. Soille and H. Talbot, "Directional morphological filtering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1313–1329, 2001.
- [28] P. Dokládál and D. Jeulin, "3D extraction of fibres from microtomographic images of fibre-reinforced composite materials," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 5720. Springer, 2009, pp. 126–136.
- [29] N. Passat, C. Ronse, J. Baruthio, J.-P. Armspach, and C. Maillot, "Magnetic resonance angiography: From anatomical knowledge modeling to vessel segmentation," *Med. Image Anal.*, vol. 10, no. 2, pp. 259–274, 2006.
- [30] B. Naegel, N. Passat, and C. Ronse, "Grey-level hit-or-miss transforms—Part II: Application to angiographic image processing," *Pattern Recogn.*, vol. 40, no. 2, pp. 648–658, 2007.
- [31] B. Bouraoui, C. Ronse, J. Baruthio, N. Passat, and P. Germain, "3D segmentation of coronary arteries based on advanced mathematical morphology techniques," *Comput. Med. Imag. Graph.*, vol. 34, no. 5, pp. 377–387, 2010.
- [32] B. D. Thackray and A. C. Nelson, "Semi-automatic segmentation of vascular network images using a rotating structuring element (ROSE) with mathematical morphology and dual feature thresholding," *IEEE Trans. Med. Imag.*, vol. 12, no. 3, pp. 385–392, 1993.
- [33] O. Tankyevych, H. Talbot, P. Dokládál, and N. Passat, "Spatially variant morpho-Hessian filter: Efficient implementation and application," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 5720. Springer, 2009, pp. 137–148.

- [34] M. Buckley and H. Talbot, "Flexible linear openings and closings," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Computational Imaging and Vision, vol. 18. Springer, 2000, pp. 109–118.
- [35] P. Salembier, A. Oliveras, and L. Garrido, "Antiextensive connected operators for image and sequence processing," *IEEE Trans. Image Process.*, vol. 7, no. 4, pp. 555–570, 1998.
- [36] M. H. F. Wilkinson and M. A. Westenberg, "Shape preserving filament enhancement filtering," in *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, ser. Lect. Notes Comput. Sci., vol. 2208. Springer, 2001, pp. 770–777.
- [37] Y. Xu, T. Géraud, and L. Najman, "Two applications of shape-based morphology: Blood vessels segmentation and a generalization of constrained connectivity," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 7883. Springer, 2013, pp. 390–401.
- [38] V. Morard, E. Decencièrre, and P. Dokládal, "Efficient geodesic attribute thinnings based on the barycentric diameter," *J. Math. Imaging Vision*, vol. 46, no. 1, pp. 128–142, 2013.
- [39] B. Caldaïrou, B. Naegel, and N. Passat, "Segmentation of complex images based on component-trees: Methodological tools," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 5720. Springer, 2009, pp. 171–180.
- [40] B. Perret, J. Cousty, O. Tankyevych, H. Talbot, and N. Passat, "Directed connected operators: Asymmetric hierarchies for image filtering and segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 6, pp. 1162–1176, 2015.
- [41] A. Dufour, O. Tankyevych, B. Naegel, H. Talbot, C. Ronse, J. Baruthio, P. Dokládal, and N. Passat, "Filtering and segmentation of 3D angiographic data: Advances based on mathematical morphology," *Med. Image Anal.*, vol. 17, no. 2, pp. 147–164, 2013.
- [42] M. H. F. Wilkinson, "Hyperconnectivity, attribute-space connectivity and path openings: Theoretical relationships," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 5720. Springer, 2009, pp. 47–58.
- [43] L. Cohen and T. Deschamps, "Segmentation of 3D tubular objects with adaptive front propagation and minimal tree extraction for 3D medical imaging," *Comput. Meth. Biomech. Biomed. Eng.*, vol. 10, no. 4, pp. 289–305, 2007.
- [44] Y. Rouchdy and L. D. Cohen, "Geodesic voting for the automatic extraction of tree structures. Methods and applications," *Comput. Vis. Image Understand.*, vol. 117, no. 10, pp. 1453–1467, 2013.
- [45] V. Bismuth, R. Vaillant, H. Talbot, and L. Najman, "Curvilinear structure enhancement with the polygonal path image – Application to guide-wire segmentation in X-ray fluoroscopy," in *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, ser. Lect. Notes Comput. Sci., vol. 7511. Springer, 2012, pp. 9–16.
- [46] J. Stawiaski, "Optimal path: Theory and models for vessel segmentation," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 6671. Springer, 2011, pp. 417–428.
- [47] L. Vincent, "Minimal path algorithms for the robust detection of linear features in gray images," in *Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Computational Imaging and Vision, vol. 12. Springer, 1998, pp. 331–338.
- [48] H. J. A. M. Heijmans, M. Buckley, and H. Talbot, "Path openings and closings," *J. Math. Imaging Vision*, vol. 22, no. 2–3, pp. 107–119, 2005.
- [49] C. L. Luengo Hendriks, "Constrained and dimensionality-independent path openings," *IEEE Trans. Image Process.*, vol. 19, no. 6, pp. 1587–1595, 2010.
- [50] H. Talbot and B. Appleton, "Efficient complete and incomplete path openings and closings," *Image Vis. Comput.*, vol. 25, no. 4, pp. 416–425, 2007.
- [51] F. Cokelaer, H. Talbot, and J. Chanussot, "Efficient robust d-dimensional path operators," *IEEE J. Sel. Topics Signal Process.*, vol. 6, no. 7, pp. 830–839, 2012.
- [52] V. Morard, P. Dokládal, and E. Decencièrre, "Parsimonious path openings and closings," *IEEE Trans. Image Process.*, vol. 23, no. 4, pp. 1543–1555, 2014.
- [53] G. K. Ouzounis and M. H. Wilkinson, "Mask-based second-generation connectivity and attribute filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 6, pp. 990–1004, 2007.
- [54] G. Hamarneh and P. Jassi, "VascuSynth: Simulating vascular trees for generating volumetric image data with ground-truth segmentation and tree analysis," *Comput. Med. Imag. Graph.*, vol. 34, no. 8, pp. 605–616, 2010.
- [55] O. Tankyevych, H. Talbot, N. Passat, M. Musacchio, and M. Lagneau, "Angiographic image analysis," in *Medical Image Processing: Techniques and Applications*. Springer, 2011, ch. 6, pp. 115–144.
- [56] M. Schaap *et al.*, "Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms," *Med. Image Anal.*, vol. 13, no. 5, pp. 701–714, 2009.
- [57] C. A. Taylor, T. A. Fonte, and J. K. Min, "Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: scientific basis," *J. Am. Coll. Cardiol.*, vol. 61, no. 22, pp. 2233–2241, 2013.



Odysée Merville received the engineering degree from the Institut Supérieur de Bio-Sciences de Paris (Université Paris-Est Créteil / ESIEE Engineering) in 2013. She is currently a PhD student at Université Paris-Est. Her scientific interests include mathematical morphology and medical imaging.



Hugues Talbot received the Habilitation from Université Paris-Est in 2013, the PhD from École des Mines de Paris in 1993 and the engineering degree from École Centrale de Paris in 1989. He was principal research scientist at CSIRO, Sydney, from 1994 to 2004. He is now a professor at Université Paris-Est. He is the co-author of 6 books and over 140 articles in the area of mathematical morphology, discrete geometry, combinatorial and continuous optimization.



Laurent Najman received the Habilitation from the University of Marne-la-Vallée, in 2006, the PhD in applied mathematics from Paris-Dauphine University in 1994, and the Engineering degree from the École des Mines de Paris, in 1991. After 10 years of research work on image processing and computer graphics problems in several companies, he joined ESIEE Paris in 2002, where he is currently a professor. His research interests include discrete mathematical morphology and discrete optimization.



Nicolas Passat obtained the MSc and PhD from Université Strasbourg 1 in 2002 and 2005, and Habilitation from Université de Strasbourg in 2011. He was an assistant professor at Université de Strasbourg, between 2006 and 2012. He is now a full professor at Université de Reims Champagne-Ardenne. His scientific interests include mathematical morphology, discrete topology, medical imaging and remote sensing.

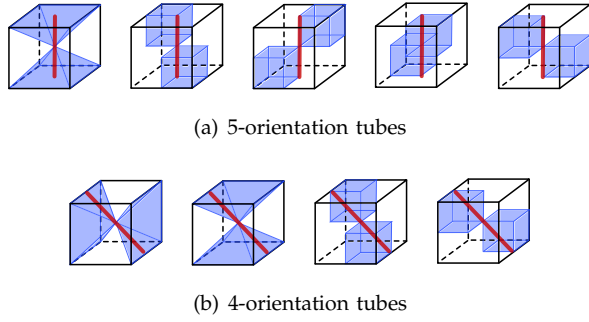


Fig. 18. The red tube lies within the border between 5 (a) or 4 (b) cones resulting in a limit case.

APPENDIX A LIMIT CASES

In Sec. 4, we propose a method to compute both intensity and directional features for regular tubular structures, i.e. the one detected in 1, 2 or 3 path opening orientations. Here we propose a simple post processing procedure to deal with the 4 and 5-orientations tubes.

Intensity feature

Cones need to overlap in order to detect some tortuous structures. Tubular structures lying entirely within the border between more than two cones may be detected either 4 times or 5 times. We propose a three-step post processing approach to deal with them.

Detecting the 4- and 5-orientation tubes

The 5-orientation tubes lie in one cone C_{e_x} plus the 4 cones C_{d_x} (an example is shown in Fig. 18.(a)) so the intersection between the 4 cones C_{d_x} , denoted Γ_5 , detects them:

$$\Gamma_5(I) = \min_{i \in \llbracket 1,4 \rrbracket} \{A_L^{C_{d_i}}(I)\} \quad (8)$$

The 4-orientation tubes lie in 2 cones C_{e_x} and 2 cones of C_{d_x} (an example is shown in Fig. 18.(b)) so the union of every combinations \mathcal{O}_i of intersections between 2 pairs of such cones, denoted Γ_4 , detects them:

$$\Gamma_4(I) = \max_{i \in \llbracket 1,6 \rrbracket} \min_{c \in \mathcal{O}_i} \{A_L^c(I)\} \quad (9)$$

Removal of remaining non-1D thin structures

Γ_4 and Γ_5 are computed by intersection of 4 path opening responses, which means that both may contain non tubular structures. We define Δ_i as the first ranked path opening orientation that does not contain i -orientation tubes. In particular, we have $\Delta_4(I) = \gamma_L^5(I)$ and $\Delta_5(I) = \gamma_L^6(I)$. $\Delta_i(I)$ contains more non tubular structures than those present in $\Gamma_i(I)$. Consequently, the geodilation $\rho(\Delta_i(I), \gamma_L^i(I))$ of Δ_i in $\gamma_L^i(I)$ is used instead of Δ_i . Finally, the removal of remaining non tubular structures is performed as follows:

$$LC_i(I) = \Gamma_i(I) - \min \{\Gamma_i(I), \rho(\Delta_i(I), \gamma_L^i(I))\} \quad (10)$$




Fig. 19. 2D orientations. The blue arrows represent the adjacency relation.

Addition of limit tubes to RORPO result

The post-processed RORPO result is then obtained from the standard RORPO (Eq. (5)), by adding the extracted limit cases (Eq. (10)):

$$\Phi^{post}(I) = \max \{\Phi(I), LC_4(I), LC_5(I)\} \quad (11)$$

More details and justifications about these steps can be found in [2].

Directional feature

To deal with 4 and 5-orientations tubes, we need to find the 4 or 5 path opening orientations involved in their detection. Then, we can apply the standard procedure, i.e. compute the mean of the orientations of interest. Fortunately, the 4 and 5 path opening orientations involved in each limit case have already been computed while handling the limit cases as shown above. Each set of intersection gives us one group of limit cases that we can then label with the correct direction, i.e. the mean of all the path openings orientations involved in this intersection.

APPENDIX B THE 2D CASE

In this article, we chose to present our operator in 3D. Nevertheless, a 2D version is simple. The few changes required are presented here.

2D Orientations

In 2D, space is discretized in 4 orientations. Two main orientations corresponding to the vectors of the orthogonal basis $\{0,1\}$, $\{1,0\}$ denoted e'_1 and e'_2 , and two diagonals corresponding to $\{1,1\}$ and $\{-1,1\}$ denoted d'_1 and d'_2 . These orientations are illustrated in Fig. 19.

2D RORPO operator

The only necessary changes is the value of i_t . In 2D, we only have to distinguish 1D thin structures (line-line) from isotropic (blob-like) structures. Blob structures are detected in all the 2D orientations. Line-like structures cannot be detected in more than 3 orientations. It is then trivial to set $i_t = 3$ in the 2D case. The reader should note that this is only a coincidence that the threshold value i_t is the same for the 2D and 3D case.