



Curvilinear 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. Curvilinear structure analysis by ranking the orientation responses of path operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018, 40 (2), pp.304-317. 10.1109/TPAMI.2017.2672972 . hal-01262728v2

HAL Id: hal-01262728

<https://hal.science/hal-01262728v2>

Submitted on 21 Feb 2017

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.

Curvilinear Structure Analysis by Ranking the Orientation Responses of Path Operators

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

Abstract—The analysis of thin curvilinear objects in 3D images is a complex and challenging task. In this article, we introduce a new, non-linear 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 most operators commonly used for 3D curvilinear structure analysis, is discrete, non-linear and non-local. From this new operator, two main curvilinear structure characteristics can be estimated: an intensity feature, that can be assimilated to a quantitative measure of curvilinearity; and a directional feature, providing a quantitative measure of the structure's 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 comparison with three of the most popular curvilinear structure analysis filters, namely Frangi Vesselness, Optimally Oriented Flux, and Hybrid Diffusion with Continuous Switch. In particular, we show that our method provides up to 8% more true positive and 50% less false positives than the next best method, on synthetic and real 3D images.

Index Terms—Thin structures, non-linear filtering, direction estimation, mathematical morphology, path opening, 3D grey-level imaging, curvilinear structure.

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 *curvilinear structures*.

Images of curvilinear structures are among the hardest to handle in image analysis. 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. Curvilinear structures can also be very




Fig. 1. An illustration of the RORPO intensity feature (b) applied on a 3D synthetic image (a) with a plane and a curvilinear structure. Both images are seen with a Maximum Intensity Projection (MIP). RORPO successfully removes the plane structure, decreases the background intensity while preserving the curvilinear structure.

tortuous, possess different orientations and scales and form a network, making geometric prior difficult to use. Nonetheless many applications involve curvilinear 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, it is frequent to use a general segmentation framework (e.g., active contours, region-growing, machine learning) enriched with some specific curvilinear prior knowledge.

Various methods have been developed to extract these specific features. Such methods aim at enhancing the curvilinear structure signal and decreasing the response of non-curvilinear 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 (see Sec. 2).

We recently introduced a new operator called RORPO (Ranking Orientation Responses of Path Operators) [1][2]. Similarly to operators already proposed in the literature, RORPO relies on a multidirectional and multi-scale paradigm. However, in contrast to most approaches from the state of the art, it is a non-linear, global discrete operator; this makes RORPO a powerful tool for the low-level processing of curvilinear structures (see Fig. 1).

In this article, we present RORPO in more details, and we propose two fundamental curvilinear structure features, namely (1) a low-level intensity feature that preserves the curvilinear structures while removing the

Odyssée Merveille, Laurent Najman and Hugues Talbot are with Université Paris-Est, LIGM (UMR 8049), CNRS, ENPC, ESIEE Paris, UPEM, F-93162, Noisy-le-Grand, France (*odyssée.merveille, laurent.najman, hugues.talbot@esiee.fr*).

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

signal of other objects; and (2) a directional feature providing an estimation of the orientation of a putative curvilinear 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 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, by comparison with three operators representative of the state of the art, namely Frangi Vesselness (FV) [3], Optimally Oriented Flux (OOF) [4], and Hybrid Diffusion with Continuous Switch (HDCS) [5]. A discussion on extensions and future work concludes this article (Sec. 7).

2 RELATED WORK: CURVILINEAR STRUCTURE FILTERING

Curvilinear structure filtering in 3D grey-level images is an active research area, that has 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 [6, Chs. 1, 2] for a general overview, and to [7], 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 curvilinear structures can be divided into two categories, namely those relying on differential (mostly linear scale-space) operators, and those relying on non-linear (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 [8], [9]. In these methods, the eigenvalues of multiscale Hessian matrices and their associated eigenvectors are analyzed to characterize blob-like, planar and curvilinear 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. [3] is the most commonly used vesselness. Many methods/variants have been proposed since then. Some of them also used the eigenvectors obtained from the Hessian matrix [10], [11], for instance for guiding a diffusion framework [12]. In [13], the second derivatives were associated to the first derivatives and a Canny filter, while in [14], a strain energy function used a stress tensor computed from the Hessian tensor. The so-called *structure tensor* can be used instead of the Hessian matrix. It is produced from the tensor product of the gradient vector convolved with a Gaussian kernel [15].

Eigenvalue analysis can also be performed to make use of anisotropy information [5], [16], [17]. It is usually less sensitive to noise than the Hessian, but cannot distinguish dark structures from light structures. A more recent alternative is to consider the projection of the image gradient onto the surface of a spherical region. This links to the notion of an oriented flux, which, when minimized, yields a method for detecting curvilinear structures [4].

To achieve multiscale 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 [18]. The use of a bi-Gaussian kernel was proposed to better take into account the anisotropy of curvilinear structures [19].

Alternatively to Hessian-based approaches, steerable filters are anisotropic filters that can be expressed in terms of a linear combination of basis filters [20]. In [21], 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 [22], using a n th Gaussian derivative basis. In [23], 3D steerable filters based on the second and fourth Gaussian derivatives were used to detect dendrite profiles.

Discrete gradients have also been used. In [24], line-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 curvilinear features. The maximum curvature of curvilinear structures is also computed by the second derivative operator along the 13 discrete lines of a $3 \times 3 \times 3$ kernel [25], [26].

2.2 Non-linear approaches

2.2.1 SE-based approaches

Non-linear approaches include those based on mathematical morphology [27]. 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 curvilinear 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 curvilinear structure by a small straight line SE [28], [29], to carry out opening or closing operations by line segments of arbitrary orientation [30], [31], [32]. The second models the background in the orthogonal hyperplane of the segment, to carry out grey-level hit-or-miss transform [33], [34], [35].

Even if the shapes of the SEs are fixed, orientation parameters can be specified to form filter banks. Several strategies were experimented: rotating structuring elements [36], knowledge-based parametrization [34], or spatially-variant mathematical morphology [37]. Blurred operators [35] were also designed to increase robustness.

The hard, straight geometry of such SEs remains a limitation to the accuracy of these approaches, progressively leading to the use of more flexible SEs [38].

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 local with SEs, but more global with connected components. In this context, the concept of component-tree [39] and attribute-based methods was specifically investigated and developed for extensive, anti-extensive and self-dual filtering.

Most considered attributes are scalar [40], [41], allowing for threshold-based interaction. Geodesic attribute were recently designed for thin structures [42]. Vectorial attributes remain infrequently used [43], due to a more complex handling.

By construction, connected filters cannot split connected component, which may result in erroneous connections between curvilinear structures and artifacts, or between branching curvilinear segments. Some attempts to minimize these drawbacks were proposed, either via tilling approaches [43], or with asymmetric notions of connectivity [44].

It is worth mentioning that hybrid SE/connected strategies were proposed, e.g. in [45] for reconnection purposes. The links between connectivity-based and path-based approaches (described below) were also investigated in [46].

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 curvilinear 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 curvilinear structures.

To address this problem, geodesic paths [47] were introduced to consider long-range, non-local interactions while still coping with the constraints of curvilinear objects, in particular noise. A curvilinear object detector was proposed in [48] using geodesic voting, similar to path density. Polygonal path images [49] extended this idea allowing for better regularization and fewer artefacts. A recent formal discussion on optimal global paths was also proposed in [50]. All these solutions remained costly in 3D.

In [51], a notion of local optimal path was pioneered. The purpose was to restrict the search 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 [52]. 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 [53] and robust to noise [54], [55], leading to a notion of robust path




Fig. 2. In 3D, (a) a blob structure, (b) a planar structure, and (c) a curvilinear 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 orientation; the planar structure presents a high response in 4 out of the 8 orientations; the curvilinear structure only presents a high response in 1 orientation.

opening. Sparse representations [56] were also proposed to avoid redundant computation.

3 PROPOSED FRAMEWORK

In this article, we aim to provide two new features characterizing curvilinear 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 structures. Here, we explain the distinguishing features of curvilinear structures; then we motivate the choice of operator used in our method: the path opening, and we recall some basic notions.

3.1 General strategy

Our strategy for distinguishing curvilinear structures from planar and blob structures is based on a simple geometric observation, illustrated in Fig. 2.

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. Curvilinear structures lie in fewer dimensions than planar or blob structures. Consequently, the number of high responses among $\{\mathcal{F}^o(I), \forall o \in \mathcal{O}\}$ for a blob or a planar structure is higher than the number of high responses for a curvilinear structure. Therefore, counting the number of high responses of an oriented filter should allow us to distinguish each kind of structure.

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 curvilinear structures which are highly anisotropic. In particular,




Fig. 3. Comparison on image (a) of a classical opening with line structuring elements (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.

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 curvilinear structure, avoid this pitfall.

Moreover, unlike classical openings, path openings can efficiently deal with local tortuosity. If a structure locally diverges from a chosen orientation, path openings can still detect this structure (see Fig. 3). 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, as they tend to move the edges. Finally, the path operators are nearly parameter-free. The only real parameter is the path length which is semantically 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 this operator more formally.

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 of L successively connected pixels (see Eq. (1)).

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

Practically, a path of length L is defined on a graph (X, Γ) , such that the connectedness Γ is the elementary adjacency relation \rightarrow periodically reproduced over X (see Fig. 4.(a,b)). We note $\Pi_L^\Gamma(X)$ the set of all paths of length L on X with connectedness Γ and $\sigma(\pi)$ the set $\{x_1, x_2, \dots, x_L\}$ of the pixels belonging to the path π .




Fig. 4. (a) An example of adjacency relation \rightarrow . (b–d) The connectedness 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).

3.2.2 Binary and grey-level path opening

The binary path opening of length L with connectedness 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. 4.(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 [52]. More formally, the thresholding of I at grey-level λ is noted $I_{\geq \lambda} = \{x \in X \mid I(x) \geq \lambda\}$, and Eq. (3) defines the grey-level path opening, $A_L^\Gamma(I(x))$. As openings, path openings are increasing, anti-extensive and idempotent.

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

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, each with a different orientation, 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 $\{d_1, d_2, d_3, d_4\}$.




Fig. 5. The 7 orientations of \mathcal{C} . The center of the cube, the centers of faces, and corners represent points of the image. The arrows represent the adjacency relation. The red arrows are the vectors e_* and d_* representing the global orientation of each cone.

There is a degree of arbitrariness for 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. 5). 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 THE ORIENTATION RESPONSES OF PATH OPERATORS

So far, we have presented the purpose of our work: characterizing curvilinear 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.

4.1 Methodology

Path operators, as an oriented filter, preserve structures compatible with specific orientations. However, this is not sufficient to distinguish curvilinear from planar and blob-like structures. For instance, a path opening on a 2D image containing an isotropic object (a blob) and a curvilinear 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.




Fig. 6. 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.

To tackle this problem, we propose to count the number of high responses of path openings. This is done by point-wise ranking of the 7 path openings responses. Let $\gamma_L^i(I)$ be the image obtained by applying a point-wise rank filter, RF_i , of order i (see Eq. (4)). In particular, RF_1 , RF_4 and RF_7 are respectively the point-wise 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 $\{\gamma_L^i(I)\}_{i \in [1,7]}$. Let us consider a structure appearing in n ($n \leq 7$) of the 7 $A_L^c(I)$; then, this structure now appears in $\{\gamma_L^i(I)\}_{i \in [1,n]}$. 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 in Fig. 6.

4.2 Filtering curvilinear structures

We have seen that curvilinear 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 curvilinear 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 curvilinear structure by definition of i_t . All other structures with a diameter/thickness greater than L are preserved in more orientations than curvilinear structures. Consequently, the residual operator $\Phi_L(I)$ reduces the signal of all structures except for those

detected in at most i_t orientations, and these are the curvilinear 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, curvilinear 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 curvilinear 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 curvilinear structures

Since RORPO is based on oriented path operators, it intrinsically carries information about the direction of the detected curvilinear structures. In order to retrieve this directional information, we have to identify which path 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 curvilinear structure orientation.

Finding which orientations detect a curvilinear structure requires, for each orientation, a binary decision: either an orientation detects the curvilinear 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 curvilinear structures 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 curvilinear structure.

Consequently, we propose to determine which path opening orientations detect a given curvilinear structure by separating the 7 orientations in two classes: those that detect the curvilinear structures (called orientations of interest) and the others. Our expectation is to achieve homogeneous responses $A_L^\Gamma(I(x))$ inside each class: high responses in the orientations of interest and low ones in the others (see Fig. 7). For simplicity and efficiency, we choose the standard deviation as homogeneity criterion and we reformulate the problem as a minimization of the intra-class standard deviation.

More formally, we recall that \mathcal{C} is the set of the 7 path opening orientations. Let S be the set of all the combinations of 1, 2 or 3 path opening orientations. One of the combinations of S is the set of the orientations of interest we seek. Let $P = \{\Gamma_1, \Gamma_2, \dots, \Gamma_k\}$, $1 \leq k \leq 3$, be one of these combinations ($P \in S$); we note $\sigma_P(x)$, the standard deviation of the path opening responses with the orientations of P at pixel x :

$$\sigma_P(x) = \sqrt{\frac{1}{k} \sum_{j=1}^k \left(A_L^{\Gamma_j}(I(x)) - \mu \right)^2} \quad (6)$$

$$\mu = \frac{1}{k} \sum_{j=1}^k A_L^{\Gamma_j}(I(x))$$




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

Then, we find the orientations of interest by solving:

$$\underset{P \in S}{\text{minimize}} \quad \sigma_P + \sigma_{\mathcal{C} \setminus P} \quad (7)$$

For this, we could compute the intra-class standard deviation of all of the combination of S (63 combinations in 3D), but it is better to use the information coming from the point-wise rank filter previously computed for the RORPO intensity feature, to reduce the number of possible combinations. 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 curvilinear structure is detected in at most 3 orientations, so if a curvilinear 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 groups: $\{O_L^1\}$, $\{O_L^1, O_L^2\}$ or $\{O_L^1, O_L^2, O_L^3\}$. Therefore, for each pixel, we only have to compute the sum of the standard deviations of these 3 possible groups with their complements and choose the one with the lowest.

Note that the first group has only one orientation and thus, only the standard deviation of its complement is considered. The direction of the structure is estimated by the mean of the orientations in the selected group. Because only a small number of path openings are used, abrupt direction changes can appear in the resulting directional feature, for example in areas of high curvature. To reduce that effect, directions are averaged using a $7 \times 7 \times 7$ window.

Orientation vs. direction: An important point to note is that we encode a curvilinear 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_\star and d_\star , we must ensure that all the vectors are set to the same right half-space. As an example, the mean of 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 curvilinear structure analysis is dealing with multiple scales. Real applications need to cope with varying diameters, lengths and curvatures, which are generally highly correlated. To tackle this issue, 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 curvilinear object analysis, the scale parameter is usually a diameter. However, in the case of RORPO, only a length parameter is available. To vary the diameter, a combination of RORPO with a top-hat filter could be envisaged. Fortunately, in many applications, the diameter, curvature and length of curvilinear 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, e.g. 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 curvilinear 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 complementary explanations about path length.

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

$$\Phi(I) = \bigcup_{L \in l} \Phi_L(I) \quad (8)$$

This multiscale paradigm also applies to the directional feature. Indeed, a direction can be computed for each scale. The final direction is the one associated with the highest RORPO response and 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, in particular for high values of L . Several contributions were devoted to propose various noise resistant versions of path operators [52], [54], in particular the Robust Path Opening (RPO) of Cokelaer et al. [55]. RPO allows for K disconnections (noisy pixels) between two successive path pixels. However this improved robustness comes at the price of increased CPU and memory usage.




Fig. 8. Illustration of our proposed simplified robust path opening (see Sec. 5.1). The disconnected curvilinear structure is fully preserved while the top noise structure is removed.

We propose an alternative to RPO which yields similar results but is less time consuming. Our method relies on a mask-based second-generation connectivity strategy [57] in order to reconnect the noisy parts of curvilinear 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. (9)).

$$A_{L,N}^c = \bigwedge \{I, A_L^c(\delta_N(I))\} \quad (9)$$

An illustration of this strategy is shown in Fig. 8.

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.

5.2.1 Path length

Path length 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 curvilinear 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 overlap as well (see Fig. 9).

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 image (N should be an odd number). 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. In the following experiments, we set $N = 3$.




Fig. 9. 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.

5.2.3 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 curvilinear 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 Section 3.2.3, a quantitative analysis of the path openings responses should be sufficient to decide whether a point belongs to a curvilinear structure. In order to verify this hypothesis experimentally, we computed the number of high responses within the 7 path opening orientations for 362 (resp. 325) synthetic binary tubes (resp. planes). The results of this study (Fig. 10) validate this conjecture. Indeed, nearly all curvilinear structures are detected in at most 3 orientations whereas all planar structures are detected in at least 4 orientations. In addition, it will be observed from the experimental results of Sec. 6 that using more than 7 orientations – in particular the variant of RORPO based on the 13 principal orientations of the $3 \times 3 \times 3$ local neighborhood – does not provide substantial improvements both in terms of intensity and direction feature determination.

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:




Fig. 10. Number of path opening orientations detecting (a) tubes and (b) planes, within the set of 7 orientations illustrated in Fig. 5.




Fig. 11. Computational cost of RORPO, with respect to $|\Omega|$ (left) and L (right, log scale).

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 neighboring 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. 5. The drawback of any covering policy is the existence of limit cases, corresponding to the paths that lie exactly at the frontier between more than three 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 A.

5.3 Computational cost

The computation cost of RORPO is dominated by the path openings. Indeed, the robustness step (Eq. (9)), 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-Appleton (TA) algorithm [54] and the Luengo algorithm [53]. The TA algorithm has a $\mathcal{O}(|\Omega| \log(L))$ complexity. We verified experimentally on synthetic images [58] (see Fig. 11) that the Luengo algorithm is equivalent, albeit slightly slower and not natively robust. However it is more generic for defining orientations, and with our simplified robustness (see Sec. 5.1), native robustness does not matter. Consequently we used the Luengo algorithm.

6 EXPERIMENTS AND RESULTS

6.1 Evaluation framework

6.1.1 Compared methods

6.1.1.1 Frangi Vesselness [3] (FV): Initially developed for the detection of vascular structures, Frangi vesselness is a local, linear, multiscale measure that relies on the analysis of the second derivatives of the image. More precisely, the eigenvalues, λ_1 , λ_2 and λ_3 ($|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$) of the Hessian matrix are involved in the definition of a function \mathcal{V} that quantifies the degree of “vesselness” of any given point x of an image:

$$\mathcal{V}(x) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ (1 - \exp(-\frac{\mathcal{R}_A^2}{2\alpha^2})) \exp(-\frac{\mathcal{R}_B^2}{2\beta^2})(1 - \exp(-\frac{\mathcal{S}^2}{2c^2})) & \text{otherwise} \end{cases} \quad (10)$$

$$\text{with } \mathcal{R}_A = \frac{|\lambda_2|}{|\lambda_3|} \quad \mathcal{R}_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2\lambda_3|}} \quad \mathcal{S} = \sqrt{\lambda_1 + \lambda_2 + \lambda_3} \quad (11)$$

The version of Frangi Vesselness we used is from ITK¹.

6.1.1.2 Optimally Oriented Flux [4] (OOF): Similarly to FV, the OOF is also a local, linear, multiscale operator; it is also based on the analysis of eigenvectors and eigenvalues of a matrix locally computed at each voxel. The main difference is that OOF relies on first-derivatives of the image, namely the image gradient observed on a local sphere around the point of interest. The spectrum of this gradient is modeled as a 3×3 matrix, defined with respect to the direction that minimizes the inward orientation flux. The “vesselness” measure is then obtained similarly to the one proposed by Sato et al. [8] in their pioneering works:

$$\mathcal{M}(x) = \begin{cases} \sqrt{|\lambda_1\lambda_2|} & \text{if } \lambda_1 \leq \lambda_2 < 0 \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

Multiscale analysis can be achieved by varying the radius of the sphere. The version of OOF we used is from a Matlab toolbox².

6.1.1.3 HDCS: Hybrid Diffusion using a Continuous Switch (HDCS) [5] is a recent method extending Edge-Enhancing Diffusion (EED) and Coherence Enhancing Diffusion (CED), both proposed by Weickert [59]. It is based on the eigen-analysis of the structure tensor $J_\rho(\nabla u_\sigma) = G_\rho \star (\nabla u_\sigma \nabla u_\sigma^\top)$, where G_ρ is the Gaussian kernel of standard deviation ρ and ∇u_σ is the gradient of image u at scale σ . HDCS balances EED and CED based on the local anisotropy. Our version of HDCS is provided by TubeTK³.

6.1.1.4 RORPO 13 directions (RORPO 13): In order to assess the relevance of considering 7 orientations in RORPO, we have also developed a version of RORPO involving a denser sampling of the orientation spaces, by considering 13 orientations regularly distributed with respect to the structure of the cubical grid, i.e. the 7 orientations described for RORPO, plus 6 others corresponding to the secondary diagonals joining the centers of opposite edges of the cube. The philosophy of the approach is the same as for the “standard” RORPO operator, and the principal difference lies in the computational cost that is increased by a factor of almost $2 \simeq 13/7$, since the orientation computations are independent.

6.1.2 Intensity feature evaluation criteria

We considered three standard evaluation criteria for the evaluation of intensity feature. The first is the Matthews Correlation Coefficient (MCC) which is well adapted when the structures of interest only represent a small proportion of the total image size:

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

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

The second and third criteria are the true positive and false positive rates (TPR, FPR), respectively. 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%.

6.1.3 Direction feature evaluation criteria

We evaluate directions by estimating the angle in degrees between the ground-truth and the direction provided by the considered operators. The mean and standard deviation of these values over the image constitute the considered evaluation criteria. We note that the maximum error between two angles is 90° , as we compute an orientation difference, i.e. an angle between lines.

6.2 Evaluation of the intensity features

In this section, we evaluate the performance of RORPO for preserving curvilinear features. We do this on both synthetic and real data.

6.2.1 Synthetic data

Since publicly available images of curvilinear features with ground-truth are rare, we used synthetic data produced by the free software package VascuSynth [58]. This package models vessels in 3D in the presence of various amounts of noise. We produced 10 images of

1. <http://itk.org>
 2. <http://www.mathworks.com/matlabcentral/fileexchange/25449-image-edge-enhancing-coherence-filter-toolbox>
 3. <http://www.tubetk.org>




Fig. 12. First row: 2D slices of one of the synthetic grey-level image generated with VascuSynth [58], with background inhomogeneity, and various levels of Gaussian white noise ($\sigma = 0, 10$ and 18). Second row: The associated “vascular model” ground-truth.

size $100 \times 100 \times 100$ and each image is degraded by 7 levels of additive white Gaussian noise (70 images in total). On all images, we added a Gaussian Random Field background (the same for all images) to simulate undesired, non-homogeneous, smooth, blob-like features (see Fig. 12). We compared the results of RORPO 7 and RORPO 13 with OOF, HDCS and FV. The parameters of all these methods were independently optimized for each method and each noise level on the first VascuSynth image using a grid search. These parameters were then used to process the 9 remaining images.

All these methods filter curvilinear structures and return grey-level images. To compute the TPR and FPR, we thresholded each result with the optimal threshold (i.e. the one maximizing the MCC). Table 1 shows the mean MCC of each best thresholded result and Table 2 shows their associated mean TPR and mean FPR. The ROC curve of one of the 10 images is shown in Fig. 13.

We can see that RORPO generally performed better at detecting curvilinear structures, as it detects far fewer false positives than the other methods, while preserving a high TPR. It is only with a high level of Gaussian noise that OOF performs slightly better than RORPO. However, this degree of artificial noise is not representative of the noise in real images. The intensity feature results of RORPO13 are slightly better than RORPO7 ones, but at a cost of a computation time increased by almost 2.

6.2.2 Real data

An experimental comparison of the intensity feature from RORPO versus FV was already conducted in [1]




Fig. 13. The ROC curves comparing the filtering of one of the ten VascuSynth images with a noise level of $\sigma = 10$.

TABLE 1

Filtering performances on synthetic images, for various levels of Gaussian white noise – MCC scores.

Noise (σ)	0	5	8	10	12	15	18
RORPO 7	0.869	0.873	0.844	0.835	0.829	0.771	0.743
RORPO 13	0.871	0.875	0.848	0.838	0.832	0.775	0.746
FV [3]	0.197	0.192	0.193	0.197	0.199	0.187	0.185
OOF [4]	0.825	0.829	0.818	0.813	0.811	0.772	0.755
HDCS [5]	0.798	0.807	0.792	0.777	0.780	0.725	0.700

using CTA images of coronary arteries of the Rotterdam Repository [60]. However this database only provides the skeleton ground-truth of the coronaries, which did not allow for a comparison as accurate as with a full segmentation ground-truth. In this article, we use a coronary CT scan image of size $507 \times 476 \times 469$ voxels, with its associated high-quality, manually defined full segmentation ground-truth. This data was kindly provided by the HeartFlow company [61].

We performed the same experiment as on the synthetic data, except that we did not use HDCS. Indeed, HDCS is more a noise filtering method than a true curvilinear feature detector. It is relevant to compare HDCS with RORPO, OOF and Frangi on images containing only curvilinear structures such as VascuSynth images; but on a cardiac CT scan presenting other structures than blood vessels, HDCS does not perform well enough.

The evaluation of each method was performed in an area of 70 pixels around the coronaries ground-truth to avoid the other curvilinear structures like the ribs. The quantitative results are shown in Table 3 and the volume rendering of each result is presented in Fig. 14.

The Frangi filter gives the worst results as it detects many more false positives. OOF and RORPO both present good results but RORPO again performs the best as it detects fewer false positives than OOF for a similar TPR, which is confirmed by the better MCC value.

6.3 Evaluation of the directional feature

TABLE 2
Filtering performances on synthetic images, for various levels of noise (Gaussian) – TPR / FPR scores.

Noise (σ)	0	5	8	10	12	15	18
RORPO 7	0.784 / 0.026	0.795 / 0.031	0.753 / 0.038	0.743 / 0.043	0.732 / 0.04	0.646 / 0.052	0.613 / 0.062
RORPO 13	0.790 / 0.029	0.798 / 0.030	0.756 / 0.035	0.744 / 0.040	0.739 / 0.045	0.657 / 0.057	0.622 / 0.066
FV [3]	0.682 / 8.766	0.713 / 9.981	0.567 / 6.577	0.613 / 7.238	0.624 / 7.238	0.580 / 7.092	0.641 / 8.763
OOF [4]	0.733 / 0.053	0.745 / 0.058	0.727 / 0.059	0.716 / 0.055	0.714 / 0.057	0.654 / 0.058	0.626 / 0.056
HDSC [5]	0.684 / 0.045	0.699 / 0.046	0.677 / 0.047	0.656 / 0.050	0.659 / 0.050	0.589 / 0.065	0.554 / 0.065




Fig. 14. The coronary ground-truth (a) and the filtering results with RORPO (b), OOF (c), FV (d) (volume rendering).

TABLE 3
Comparison of RORPO, OOF and FV results on a coronary CT scan.

	MCC	TPR	FPR	direction error in degrees: mean (std)
RORPO 7	0.541	0.557	0.475	16.9 (13.3)
FV [3]	0.405	0.492	0.911	14.8 (13.4)
OOF [4]	0.529	0.582	0.594	21.4 (24.6)

To the best of our knowledge, the precision of the directional feature of vesselness-like filters has not been studied extensively. As far as we know, no publicly available ground-truth for directions exists. As above, we perform an evaluation on synthetic and real data.

6.3.1 Evaluation on synthetic data

Here we use again the same synthetic data based on VasculSynth as for the intensity feature. Since the VasculSynth ground-truth consists of a union of cylinders of various orientations and diameter, its local directions can be computed easily.

We computed the directional feature of RORPO 7 and RORPO 13 as presented in Section 4.3. As HDSC does not provide directional information; we compare only with OOF and FV. The local direction of OOF and FV is given by the eigenvector associated to the lowest absolute eigenvalue at each pixel. For a fair comparison, and to avoid errors due to false positives, we only compare directions at pixels that are detected as belonging to a curvilinear structure by all three of OOF, FV and RORPO. The false positives were already taken into account in the evaluation of the intensity feature.

The results are shown in Table 4. The best directions are provided by RORPO 7 almost at all noise levels. We note that the RORPO results are very stable with respect to noise. Indeed, RORPO uses paths, which are semi global structures. The RORPO directions integrate the orientation information over the whole path which are, by their non-locality and anisotropy, a more meaningful neighborhood for the analysis of curvilinear structures than the isotropic neighborhood used by Frangi and OOF. We also note that augmenting the number of sampling directions with RORPO 13 does not improve the results. We experimentally observed that the best selected scale is usually larger for RORPO 13 than RORPO 7, which induces worse results for RORPO 13. Indeed, a longer path length implies a poorer local adaptation in curvilinear structures

6.3.2 Evaluation on real data

We evaluated the directional feature on the CTA exam provided by HeartFlow. The ground-truth direction is given by reference to that of the centerline, also provided by HeartFlow. With optimal parameters, we computed the RORPO, OOF and FV directional features on this exam and evaluated the angular error. As before, we only compare directions for pixels detected as curvilinear structures by the three methods. A selected example of results is shown in Fig. 15 and quantitative results of the comparison are shown in Table. 3.

We see that FV performs slightly better on the accuracy of directions. However it is important to note that this

TABLE 4

Direction error (in degrees) on synthetic images, for various levels of noise (Gaussian) – Mean values and standard deviation into brackets.

Noise (σ)	0	5	8	10	12	15	18
RORPO 7	15.23 (13.83)	15.86 (13.67)	16.16 (14.15)	14.91 (13.40)	15.04 (13.35)	15.13 (13.55)	15.66 (13.66)
RORPO 13	20.86 (14.38)	22.75 (14.93)	21.90 (15.20)	19.86 (14.06)	19.77 (13.86)	19.98 (13.88)	21.05 (14.23)
FV [3]	16.26 (13.28)	16.44 (13.64)	16.95 (13.93)	17.64 (14.85)	17.93 (15.14)	18.16 (15.19)	19.91 (16.96)
OOF [4]	15.42 (13.51)	16.86 (15.64)	15.14 (13.37)	15.92 (14.21)	16.14 (14.54)	17.00 (15.31)	17.87 (15.89)

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 about two times more false directions. In contrast, RORPO provides a slightly worse error (17° vs. 15°) for the orientation of blood vessels but also computes far fewer false positives directions.

The directions provided by OOF are generally worse than those given by Frangi and RORPO, mostly because of errors on the contours of the vessels. It seems that the OOF directions are very sensitive to the chosen scale. In this experiment, the scales were optimized to provide the best OOF intensity feature, which results in directions errors on the contours of larger vessels. Nonetheless, the OOF directions seem, qualitatively, better for small vessels than the Frangi ones.

In real image processing applications, a directional feature can usually be used at two different stages: 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 many fewer false positives, and so resulting in generally more accurate directions. In the second case, when the segmentation is already available, the Frangi Vesselness can 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 curvilinear structures. It consists of an intensity feature which measures the presence of a curvilinear structure, and a directional feature providing at each point the orientation of this curvilinear structure. We have shown that RORPO is a convincing alternative to derivative-based operators which constituted, until now, the reference frameworks proposing both intensity and directional features for grey level images. Indeed, FV, OOF and HDCS rely on local and isotropic eigen-analysis, whereas curvilinear structures are intrinsically anisotropic. RORPO, by using path operators, considers non-local and anisotropic neighborhoods which are well adapted to curvilinear structures.




Fig. 15. Illustration of the RORPO (a), FV (b) and OOF (c), directional feature on the coronary CT scan provided by the HeartFlow company [61].

Experimental quantitative and qualitative evaluations show that both of our features yield state-of-the-art results for the analysis of curvilinear structures. The intensity feature is indeed reliable, as it detects many fewer false positives than other state of the art methods, and the directional feature is more robust as it relies on non-local SE.

As a low level operator for the characterization of curvilinear structures, RORPO can be used as prior information in segmentation frameworks in the same way as the Hessian [62]. 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>). The CTA and ground-truth for our real-data experiments was provided by HeartFlow (<http://www.heartflow.com>).

REFERENCES

- [1] O. Merveille, H. Talbot, L. Najman, and N. Passat, "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.
- [2] —, "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] A. Frangi, W. Niessen, K. Vincken, and M. 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.
- [4] 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.
- [5] A. M. Mendrik, E.-J. Vonken, A. Rutten, M. A. Viergever, and B. van Ginneken, "Noise reduction in computed tomography scans using 3-d anisotropic hybrid diffusion with continuous switch," *IEEE Trans. Med. Imag.*, vol. 28, no. 10, pp. 1585–1594, 2009.
- [6] H. Talbot, "Étude des directions en analyse d'image," Habilitation Thesis, Université Paris-Est, 2013, (In English).
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] H. E. Bennink *et al.*, "A novel 3D multi-scale liness 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.
- [14] 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.
- [15] J. Bigun and G. Granlund, "Optimal orientation detection of linear symmetry," in *Proc. First int. Conf. on Computer Vision (ICCV)*, London, 1987, pp. 433–438.
- [16] M. Krause, J. M. Hausherr, B. Burgeth, C. Herrmann, and W. Krenkel, "Determination of the fibre orientation in composites using the structure tensor and local x-ray transform," *Journal of Materials Science*, vol. 45, no. 4, pp. 888–896, 2009.
- [17] R. Moreno and Ö. Smedby, "Gradient-based enhancement of tubular structures in medical images," *Med. Image Anal.*, vol. 26, no. 1, pp. 19–29, 2015.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] M. Orkisz, 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.
- [25] 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.
- [26] B. E. 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.
- [27] L. Najman and H. Talbot, Eds., *Mathematical morphology: From theory to applications*. ISTE/John Wiley & Sons, 2010.
- [28] 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.
- [29] O. Tankyevych, H. Talbot, and P. Dokládál, "Curvilinear morpho-Hessian filter," in *Proc. IEEE Int. Symp. Biomed. Imaging*. IEEE, 2008, pp. 1011–1014.
- [30] 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.
- [31] P. Soille and H. Talbot, "Directional morphological filtering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1313–1329, 2001.
- [32] P. Dokládál and D. Jeulin, "3D extraction of fibres from micro-tomographic 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.
- [33] 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.
- [34] 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.

- [35] 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.
- [36] 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.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] V. Morard, E. Decencière, and P. Dokládál, "Efficient geodesic attribute thinnings based on the barycentric diameter," *J. Math. Imaging Vision*, vol. 46, no. 1, pp. 128–142, 2013.
- [43] 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.
- [44] 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.
- [45] A. Dufour, O. Tankyevych, B. Naegel, H. Talbot, C. Ronse, J. Baruthio, P. Dokládál, 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.
- [46] 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.
- [47] 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.
- [48] 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.
- [49] 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.
- [50] J. Stawiński, "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.
- [51] 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.
- [52] 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.
- [53] C. L. Luengo Hendriks, "Constrained and dimensionality-independent path openings," *IEEE Trans. Image Process.*, vol. 19, no. 6, pp. 1587–1595, 2010.
- [54] H. Talbot and B. Appleton, "Efficient complete and incomplete path openings and closings," *Image Vis. Comput.*, vol. 25, no. 4, pp. 416–425, 2007.
- [55] 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.
- [56] V. Morard, P. Dokládál, and E. Decencière, "Parsimonious path openings and closings," *IEEE Trans. Image Process.*, vol. 23, no. 4, pp. 1543–1555, 2014.
- [57] G. K. Ouzounis and M. Wilkinson, "Mask-based second-generation connectivity and attribute filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 6, pp. 990–1004, 2007.
- [58] 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.
- [59] J. Weickert, *Anisotropic diffusion in image processing*. Teubner Stuttgart, 1998, vol. 1.
- [60] 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.
- [61] 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. Cardiology*, vol. 61, no. 22, pp. 2233–2241, 2013.
- [62] O. Merveille, O. Miraucourt, S. Salmon, N. Passat, and H. Talbot, "A variational model for thin structure segmentation based on a directional regularization," in *Proc. Int. Conf. Image Process. IEEE*, 2016, pp. 4324–4328.



Odysée Merveille received the engineering degree from the Institut Supérieur de Bio-Sciences de Paris (Université Paris-Est Créteil / ESIEE Paris) and the master degree "Signal and Image in Medicine" from Université Paris-Est in 2013. She received the PhD in computer science from Université Paris-Est in 2016 and is currently a postdoc at Université Strasbourg. 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 ESIEE Paris. He is the co-author of 6 books and over 180 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 professor at Université de Reims Champagne-Ardenne. His scientific interests include mathematical morphology, discrete topology, medical imaging and remote sensing.




Fig. 16. The red tube lies within the border between 5 (a) or 4 (b-c) 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 curvilinear 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. Curvilinear 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 \mathcal{C}_{e_*} plus the 4 cones \mathcal{C}_{d_*} (an example is shown in Fig. 16.(a)) so the intersection between the 4 cones \mathcal{C}_{d_*} , denoted Γ_5 , detects them:

$$\Gamma_5(I) = \min_{i \in \llbracket 1,4 \rrbracket} \{A_L^{\mathcal{C}_{d_i}}(I)\} \quad (13)$$

There are 2 pattern of 4-orientation tubes: tubes lying in 2 cones \mathcal{C}_{e_*} and 2 cones of \mathcal{C}_{d_*} (an example is shown in Fig. 16.(b)) and tubes lying in 3 cones of \mathcal{C}_{e_*} and 1 cone of \mathcal{C}_{d_*} . The union of every combinations \mathcal{O}_i of such 4 cones, denoted Γ_4 , detects the 4-orientations tubes:

$$\Gamma_4(I) = \max_{i \in \llbracket 1,10 \rrbracket} \min_{c \in \mathcal{O}_i} \{A_L^c(I)\} \quad (14)$$

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-curvilinear 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-curvilinear structures than those present in $\Gamma_i(I)$. Consequently, the




Fig. 17. 2D orientations. The blue arrows represent the adjacency relation.

geodilation [27, chapter 1], $\rho(\Delta_i(I), \gamma_L^4(I))$, of Δ_i in $\gamma_L^4(I)$ is used instead of Δ_i . Finally, the removal of remaining non-curvilinear structures is performed as follows:

$$LC_i(I) = \Gamma_i(I) - \min \{ \Gamma_i(I), \rho(\Delta_i(I), \gamma_L^4(I)) \} \quad (15)$$

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. (15)):

$$\Phi^{post}(I) = \max \{ \Phi(I), LC_4(I), LC_5(I) \} \quad (16)$$

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. 17.

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. Curvilinear 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.