

Tubular Structure Filtering by Ranking Orientation Responses of Path Operators[★]

Odyssée Merveille^{1,2}, Hugues Talbot¹, Laurent Najman¹, and Nicolas Passat²

¹ Université Paris-Est, LIGM, UPEMLV-ESIEE-CNRS, France

² Université de Reims Champagne-Ardenne, CReSTIC, France

Abstract. Thin objects in 3D volumes, for instance vascular networks in medical imaging or various kinds of fibres in materials science, have been of interest for some time to computer vision. Particularly, tubular objects are everywhere elongated in one principal direction – which varies spatially – and are thin in the other two perpendicular directions. Filters for detecting such structures use for instance an analysis of the three principal directions of the Hessian, which is a local feature. In this article, we present a low-level tubular structure detection filter. This filter relies on paths, which are semi-global features that avoid any blurring effect induced by scale-space convolution. More precisely, our filter is based on recently developed morphological path operators. These require sampling only in a few principal directions, are robust to noise and do not assume feature regularity. We show that by ranking the directional response of this operator, we are further able to efficiently distinguish between blob, thin planar and tubular structures. We validate this approach on several applications, both from a qualitative and a quantitative point of view, demonstrating noise robustness and an efficient response on tubular structures.

Keywords: mathematical morphology, non-linear filtering, path operators, thin structures, 3D imaging.

1 Introduction

Thin structures can be hard to detect in images. The difficulties stem from spatial sparsity, small size leading to partial volume effects, similarity with textured noise, potential tortuosity, the complexity of their topology when organised as networks and other problems. These difficulties are magnified in 3D applications (*e.g.*, in materials or angiographic imaging), due to structural issues (thin structures are no longer of codimension 1 in 3D) or computational burden.

Some solutions exist for filtering thin objects in 3D images, *i.e.*, for enhancing their signal and/or removing noise and artifacts that limit their detection. Most are based on scale-space and the differential properties of images, mainly via linear approaches (Sec. 2.1). This has led to the development of filters characterizing thin structures by

[★] This research was funded by *Agence Nationale de la Recherche* (Grant Agreement ANR-12-MONU-0010).

considering local information at various scales. Non-linear strategies have also been investigated, particularly in mathematical morphology (Sec. 2.2). To cope with the sparsity or thin structures, both linear and non-linear approaches either use a steerable approach or an orientation sampling procedure. Whereas directional sampling approaches are computationally feasible in 2D, in 3D they are often prohibitively expensive.

Recently, path operators have opened a promising way to take into account the strong anisotropy and non-local properties of thin structures. They are tailored to cope with the challenges of detecting thin, anisotropic, non-necessarily locally regular, and noisy structures. They are available in 3D, but by themselves they do not allow to distinguish between the various classes of thin objects.

In this article, we propose a 3D tubular structure filtering method based on path operators (Sec. 3). We describe this filter, formalised in the mathematical morphology framework (Sec. 4). In order to assess the behaviour of this filter, we compare it versus a gold-standard Hessian-based filter and a more traditional morphological approach. This comparative study is carried out both on synthetic and real data (Sec. 5). We finally summarise our contributions, and discuss further improvements (Sec. 6).

2 Related Works

The methods devoted to filtering thin structures can be divided – among many other classifications – into two categories, namely those relying on differential (mostly linear) operators (Sec. 2.1), and those relying on non-linear (often mathematical morphology) operators (Sec. 2.2).

2.1 Differential Filters

Thin structures filtering and detection based on second-order derivative properties of the image were first proposed in [29,21]. In these, the eigenvectors of multiscale Hessian matrices and their associated eigenvalues can be analysed to characterise isotropic (blobs), planar and tubular structures as well as their scale and orientation.

This strategy has led to the proposal of several “vesselness” measures, that combine differential information into heuristic formulations. The vesselness proposed by Frangi *et al.* in [13] 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 [18,8], for instance for guidance of a diffusion framework [23]. In [2] the second derivatives were associated to first derivatives and a Canny filter, while in [40], a strain energy function used a stress tensor computed from the Hessian tensor.

To achieve multiscale detection, derivative operators are typically combined with a differentiable convolution kernel. In order to avoid the induced blurring effects, these computations may be replaced by a gradient vector flow [1]. The use of a bi-Gaussian kernel was also proposed to better take into account the bimodal nature of tubular structures versus background [41]. This issue was also dealt with by considering optimally oriented flux [19].

Steerable filters [14] are anisotropic filters that can be expressed in term of a linear combination of basis filters. As such they are often used to detect oriented features

such as tube-like structures. In [17], the convolution between a bar profile and the second derivative is used in a non-linear combination to introduce a multiscale approach. A framework for 3D steerable filters was first proposed in [9], using a n^{th} Gaussian derivative basis filter. Then, it was proposed in [15] to use 3D steerable filters based on the 2^{nd} and 4^{th} Gaussian derivatives to detect dendritic profiles and used them as input for a classification approach.

Discrete gradients are also used in tube-like detection. In [27], tube-like orientation is first estimated using a set of discrete orientations called “sticks”. Then instead of a classical low pass filter, an anisotropic non-linear filter, the L-filter, is used to enhance tube-like features. The Maximum Curvature [11,5] of tube-like structure is also computed by the 2^{nd} derivative operator along the thirteen discrete lines of a $3 \times 3 \times 3$ kernel.

2.2 Non-linear Filters

Non-linear approaches include those based on mathematical morphology [26]. A common notion is the structuring element (SE), a geometric pattern from which basic operators (erosions, dilations, openings, closings, *etc.*) can be defined. Efficient operators using line segments in 2D were proposed in [31,32] and extended to 3D in [10]. Grey-level hit-or-miss transforms were explored as a way for detecting 3D vessels in angiographic data in [25]. Spatially-variant mathematical morphology also led to the development of approaches for anisotropic SE-based linear structure filtering, with reconnection purposes [42,36,12]. In these works, both linear and non-linear techniques were used in synergy.

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. The resulting region-based approaches have been involved in the design of hierarchical filtering techniques, for instance via the notion of component-tree [39], and their asymmetrical variants [34].

The SE- and connectivity-based approaches present dual intrinsic strengths and weaknesses. The SE-based approaches naturally handle anisotropy – which is highly desirable for linear structure filtering – but they still lead to mostly local filters. This also a weakness of differential operators. In comparison, the connectivity-based approaches lead to more global descriptors; but the underlying notion of adjacency remains too generic to efficiently model the anisotropy of linear structures.

To answer this problem, geodesic paths [6] have been introduced as a solution 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 [28] using geodesic voting, similar to path density. Polygonal path images [3] extended this idea allowing for better regularization and fewer artifacts. However these solutions are currently very costly in 3D.

The notion of morphological path operator, introduced in [4,16], constitutes a way to merge the strengths of these approaches, in the context of thin structure filtering. We present these operators in Sec. 4, however we now present our main strategy.




Fig. 1. When sampling orientations from a point, in an isotropic structure a.k.a. a blob (a), oriented operators all respond nearly identically (green arrows). In a plane (b), some proportion respond positively. In a tube (c), only a few orientations respond.

3 General Strategy

Our strategy for filtering tubular structures derives from the simple observation of Fig. 1. Let $\mathcal{F}_{\lambda,\phi}$ be the response of an oriented filter on image I depending on an orientation sampling of the unit sphere by longitude λ and latitude ϕ . Without loss of generality, we assume a bright feature on a dark background. For any orientation-dependent filter, be it linear or non-linear, sampling based or steerable, we will have that in any blob features of larger dimension than the considered scale, the response will be positive and similar in all directions. Conversely, assuming a horizontal planar feature, a positive response will only be observed in the directions where $\phi \equiv 0$. Finally, if the local feature is tubular, then there is only a narrow range of directions λ and ϕ where a positive response will be observed.

Therefore, if we consider a discretization $\{(\lambda_1, \phi_1), (\lambda_2, \phi_2), \dots, (\lambda_n, \phi_n)\}$ of all possible orientations on the unit sphere, we could distinguish blobs, planes and tubes by counting the number of orientations where a positive response is observed. However this would not provide us with a filtering method. Instead, we can *rank* the responses pixelwise from highest to lowest, *i.e.*, at a particular position (x, y, z) , we would have for $i, j, m \in [1, \dots, n]$:

$$\mathcal{F}_{\lambda_i, \phi_i}[I](x, y, z) \geq \mathcal{F}_{\lambda_j, \phi_j}[I](x, y, z) \geq \dots \geq \mathcal{F}_{\lambda_m, \phi_m}[I](x, y, z) \quad (1)$$

This pixelwise ranking allows us to construct n response images $\mathcal{F}_1, \dots, \mathcal{F}_n$ such that, for all points of I ,

$$\mathcal{F}_1[I] \geq \mathcal{F}_2[I] \geq \dots \geq \mathcal{F}_n[I] \quad (2)$$

Here, $\mathcal{F}_1[I]$ is the image of maximum response and $\mathcal{F}_n[I]$ is the minimum response. In 2D, the oriented difference residual $\text{ODR} = \mathcal{F}_1[I] - \mathcal{F}_n[I]$ has been used before to distinguish between isotropic and oriented areas in image [32]. In 3D, this residual performs the same function, *i.e.*, it would provide a high response for Fig. 1(b) or (c) but not (a), and a fortiori the response would also be low in dark areas of the volume. Now, if we assume there is some level m such that in bright planar features, the response \mathcal{F}_m is always high and always low in bright tubular features, then the tubular difference residual $\text{TDR} = \mathcal{F}_1[I] - \mathcal{F}_m[I]$ offers a high response only in tubular structures.

In many ways, this approach is similar to the Hessian analysis used in various vesselness measures. Principal feature directions are obtained by eigen-analysis instead of sampling directions, and response is replaced by rapid variation, modeled by curvature, *i.e.*, second derivatives of the intensities. Blobs are objects that have little variation along the three eigenvectors, planes have high variation in one direction, and tubes in two. However performing Hessian analysis imposes to consider a smoothed, isotropic, local neighborhood, at various scales. To improve on this, a non-local, anisotropic operator may be preferable. For efficiency, it is also crucial to consider operators that do not require sampling orientations with a high density. Amongst the possible tools satisfying these requirements, those based on paths are the most promising. However, most of them are currently too costly to be used in 3D image analysis. In the following, we introduce path operators which seems to be the most effective choice to date.

4 Ranking Orientation Responses of Path Operators

4.1 Path Operators

Morphological path operators lie at the convergence of mathematical morphology and graph-based path optimization. They include in particular the local minimal path approach [38], the polygonal path image [3] (both available only in the 2D case so far), and the path openings/closings, that we describe hereafter, and which constitutes the basis of our proposed filter. Without loss of generality we consider here the path opening for filtering bright objects on a dark background. Path closing would simply do the converse. Path openings were introduced in [4,16]. As stated, they derive from SE-based approaches, as they consider as structuring elements families of curvilinear paths. However, they also derive from connectivity-based approaches, since these paths are intrinsically linked to the topological structure mapped onto the image support.

Adjacency The support X of an image can be equipped with adjacency links, *i.e.*, a binary relation \rightarrow that is irreflexive and possibly non-symmetric. Practically, the existence of a link $a \rightarrow b$ between two points a and b of X means that one can go locally from a to b in X . In other words, the couple (X, \rightarrow) constitutes a (directed) graph. A path \mathbf{a} of length L is then constituted of L points of X successively adjacent. More generally, we denote as $\Pi_L(X)$ the set of all the paths of length L on X , with respect to an adjacency relation. We note $\sigma(\mathbf{a})$ the set of points successively constituting the path \mathbf{a} . More formally, we have $\sigma(\mathbf{a}) = \{a_1, a_2, \dots, a_L\}$ such that $a_i \rightarrow a_{i+1}$ for $i = 1$ to $L - 1$. The choice of the adjacency then controls the orientation, anisotropy degree and potential tortuosity of the induced paths.

Binary path opening The binary path opening $\alpha_L(X)$ of length L of a set X – viewed as a binary image – is defined as the union of all the paths of length L in X , *i.e.*

$$\alpha_L(X) = \bigcup \{\sigma(\mathbf{a}) \mid \mathbf{a} \in \Pi_L(X)\} \quad (3)$$

In other words, it preserves in $\alpha_L(X)$ all the points of X which belong to at least one path of length L . As a combination of openings, it is an (algebraic) opening.

Grey-level path opening Extension of path openings to the case of grey-level images is straightforward. Specifically the path opening of a grey-level image I , defined on X and taking its values in a finite set $\Lambda \subseteq \mathbb{Z}$ corresponds to stacking the results of the binary path openings at each grey-level value $\lambda \in \Lambda$, carried out on the subgraphs associated to the threshold sets $X_\lambda(I) = \{x \in X \mid I(x) \geq \lambda\}$. More formally, the grey-level path opening $A_L(I)$ is the grey-level image defined on X as

$$A_L(I)(x) = \max\{\lambda \in \Lambda \mid x \in \alpha_L(X_\lambda(I))\} \quad (4)$$

Path opening computation and robust path opening The adjacency considered for defining a given path opening operator is generally periodic and spatially invariant over the space X . This regularity allowed the development of efficient algorithms [33], and in any dimensions [22]. Since long paths can be sensitive to noise, an improved algorithm was also proposed, for coping with paths that include up to a maximum ratio of noisy nodes [33]. A simpler, more efficient approach, based on a local tolerance to noise led to the definition of Robust Path Openings (RPO) [7].

4.2 RPO-Based Filtering

If the considered adjacency is the standard 6- or 26-adjacency, the grey-level path opening simply reduces to a standard, isotropic connected opening [26]. Path opening however authorizes the use of anisotropic adjacencies. These can model orientation in the considered space X . In the 3D discrete space \mathbb{Z}^3 , 3, 7 or 13 principal orientations can naturally be defined. These rely on the grid's principal orientations, and its principal and the secondary diagonals, respectively. Adjacency configurations with 7 orientations are sufficient to accurately quantify the orientation space in a path-based paradigm. In this multi-directional framework it is then possible to develop an oriented version of path-based filtering for thin structure filtering.

4.3 Orientation Space Sampling

As stated above, we consider the 7 main orientations of the discrete space \mathbb{Z}^3 , composed by the three principal orientations corresponding to the 3 vectors of the orthogonal basis $\{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$, denoted $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$, plus the 4 vectors composing the principal diagonals $\{(1, 1, 1), (1, 1, -1), (1, -1, 1), (1, -1, -1)\}$, denoted $\{\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3, \mathbf{d}_4\}$. We note \mathcal{D} the set of these 7 vectors. In order to cover the whole space around these 7 axes, we consider a 3D discrete cone around each of them. In \mathbb{R}^3 , such cover is obtained by associating to \mathbf{e}_\star (resp. \mathbf{d}_\star) the cone $C_{\mathbf{e}_\star}$ (resp. $C_{\mathbf{d}_\star}$) bounded by 4 (resp. 3) edges composed by the vectors $\{\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3, \mathbf{d}_4\}$ (resp. $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$) that includes \mathbf{e}_\star (resp. \mathbf{d}_\star). We note C_e the set of cones generated by \mathbf{e}_\star and C_d the set of cones generated by \mathbf{d}_\star . Then C is the set of the 7 cones. These cones – which are unambiguously modelled in \mathbb{Z}^3 due to their discrete boundaries – are illustrated in Fig. 2.

The space \mathbb{R}^3 is fully covered by these 7 cones of C . More precisely, each point lies into exactly 2 cones over the 7, except points that lie into cone boundaries, that lie in at most 4 cones. Side effects of this property are discussed at the end of the section.

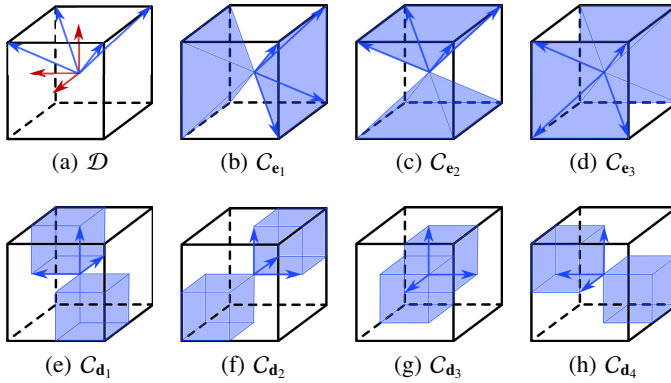


Fig. 2. (a) The 7 vectors of \mathcal{D} . (b–h) The 7 cones (in blue) of \mathcal{C} . These cones form a cover of \mathbb{R}^3 (see text).

4.4 Cone-Oriented Robust Path Opening

For each cone C_c of \mathcal{C} , it is possible to carry out a path opening, based on the graph induced by C_c on the support X of the image. This elementary opening operation is parametered by C_c and the length L of the paths that are preserved in the filtered image (see Eqs. (3–4)). These are the only two parameters that govern the behaviour of the path-opening. This small set of parameters have a physical meaning, since they refer to the orientation and size of the structures that the opening preserves. Consequently, they are easy to tune, based on application context.

In [7], noise-robust path opening consists of locally authorizing disconnections of up to K successive noise pixels. K is usually set to 1, which gives good results in practice. In the sequel, for a given cone C_c , we note $R_c^L(I)$ such robust path opening of the image I , for the path length L , and local noise tolerance $K = 1$.

4.5 Pointwise Rank Filtering

For an image I , we can compute the 7 RPO-filtered images $R_c^L(I)$. These images then provide orientation information at every point $x \in I$. As seen in Sec. 3 the distribution of the values of these 7 images can provide an accurate characterisation of the image anisotropy at x . For instance, if all the values $R_c^L(I)(x)$ are similar, the image intensity is necessarily homogeneous around x , characterising an isotropic structure. Contrariwise, if these values are high in some number of directions, and low in the others, the structure around x is anisotropic, thus characterising a thin structure.

Based on these remarks, we can build the 7 rank-filtered images $\phi_i^L(I)$, obtained by pointwise rank filters of the $R_c^L(I)$. More formally, for $i \in [1, 7]$, we set

$$\phi_i^L(I)(x) = \text{RF}_i \{ R_c^L(I)(x) \mid \mathbf{c} \in \mathcal{C} \} \quad (5)$$

where RF_i is the rank-filter of order i , *i.e.*, that yields the i^{th} highest value within a set. In particular, $\phi_1^L(I)$, $\phi_4^L(I)$ and $\phi_7^L(I)$, are respectively the pointwise max, median and min filtered images obtained from the $R_c^L(I)$ RPOs.

4.6 RORPO: A Filter Based on Ranking Orientations Responses Path Operator

As stated above, the 7 values of $R_c^L(I)(x)$ provide some information about the anisotropy of the structure around x . We can see that any line-like structure, excluding some limit cases we will deal with later, can only be contained in up to three cones at once. Conversely, any plane-like structure intersects at least five cones. We confirmed this geometric result by exhaustively analysing a dense sampling of the orientation space using randomly generated tube-like and plane-like discrete structures.

These validated results motivate our definition of the RPO-based filtering $\phi^L(I)$ as the tubular difference operator residual

$$\phi^L(I) = \phi_1^L(I) - \phi_4^L(I) \quad (6)$$

Less formally, the RPO-based filtering $\phi^L(I)$ is computed as the difference between the max and the median values of the pointwise rank-filters of the RPO images. We term our novel filter RORPO: Ranking Orientations Responses Path Operator.

4.7 Suppressing Artifacts Generated by Limit Cases

As stated above, tubular structures should be detected in at most three RPO orientations, but some limit cases occur. Indeed, the definition of orientations as discrete closed cones implies an overlap between neighboring RPO orientations. When one of these overlapping zones concerns more than two orientations, a tubular structure lying into this overlapping zone may be detected in more than three orientations and so is not considered tubular anymore. Experimental evidence suggests that tubes concerned by these limit orientations represent less than 4% of all tubes and can be classified into 2 patterns: 4-orientation tubes and 5-orientation tubes. 4-orientation tubes are detected in 2 of the 3 orientations defined by C_e plus 2 of the 4 orientations defined by C_d . 5-orientation tubes are detected in 1 of the 3 orientations defined by C_e plus the 4 orientations defined by C_d .

In order to keep 4 and 5-orientation tubes in the final result, we need to detect these special cases using extra processing. As 4 and 5-orientation tubes are defined as lying into overlapping of a finite set of RPO orientations, they can easily be isolated using intersection of orientations (see Eqs. (14) and (16)). We define the following cone sets, that correspond to relevant special cases:

$$O_1 = \{C_{e_1}, C_{e_2}, C_{d_2}, C_{d_3}\} \quad (7)$$

$$O_2 = \{C_{e_1}, C_{e_2}, C_{d_1}, C_{d_4}\} \quad (8)$$

$$O_3 = \{C_{e_2}, C_{e_3}, C_{d_1}, C_{d_2}\} \quad (9)$$

$$O_4 = \{C_{e_2}, C_{e_3}, C_{d_3}, C_{d_4}\} \quad (10)$$

$$O_5 = \{C_{e_1}, C_{e_3}, C_{d_2}, C_{d_4}\} \quad (11)$$

$$O_6 = \{C_{e_1}, C_{e_3}, C_{d_1}, C_{d_3}\} \quad (12)$$

We then detect the missing tubular structure by the following algorithm:

$$\Gamma_5 = \min_{c \in C_d} (R_c^L(I)(x)) \quad (13)$$

$$LC_5 = \Gamma_5 - \min(\Gamma_5, \gamma(\Phi_6^L, \Phi_4^L)) \quad (14)$$

$$\Gamma_4 = \max_{i \in \{1, \dots, 6\}} \min_{c \in O_i} (R_c^L(I)(x)) \quad (15)$$

$$LC_4 = \Gamma_4 - \min(\Gamma_4, \gamma(\Phi_5^L, \Phi_4^L)). \quad (16)$$

where $\gamma(I, M)$ is the morphological reconstruction operator of image I under mask M (see [37]). Finally, we have

$$\text{RORPO}(I) = \max\{\phi^L(I), LC_4, LC_5\} \quad (17)$$

Thus, handling the limit cases simply consists of isolating the 4-orientation and 5-orientation tubes and to compose them to the result of RORPO. We note that this requires negligible additional computation as only simple operators are involved.

5 Experiments and Results

We now experimentally evaluate the RORPO filter. These evaluations are quantitatively carried out on synthetic data with ground truth, allowing to compute standard quality scores. The RORPO filter is then applied on challenging 3D images in the context of angiography.

5.1 Compared Methods and Quality Scores

In order to assess the efficiency of the RORPO filter, we consider two other methods, namely Frangi's vesselness, which is the gold-standard within the family of differential filters, and a classical morphological filter based on path-openings and a top-hat.

Frangi's vesselness [13]. As evoked in Sec. 2.1, Frangi's vesselness is a multiscale tubularity measure that relies on the eigenvalues, λ_1 , λ_2 and λ_3 , of the Hessian matrix. The values are then involved in the definition of a function \mathcal{V}_0 that quantifies the degree of vesselness of any given point x of an image, with respect to the local intensity profile in its neighbourhood:

$$\mathcal{V}_0(x) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ (1 - \exp(-\frac{\mathcal{R}_A^2}{2\alpha^2})) \exp(-\frac{\mathcal{R}_B^2}{2\beta^2})(1 - \exp(-\frac{\mathcal{S}^2}{2c^2})) & \text{otherwise} \end{cases} \quad (18)$$

$$\text{with } \mathcal{R}_A = \frac{|\lambda_2|}{|\lambda_3|} \quad \mathcal{R}_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2\lambda_3|}} \quad \mathcal{S} = \sqrt{\lambda_1 + \lambda_2 + \lambda_3} \quad (|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|) \quad (19)$$

The parameters α , β and c are set as originally proposed in [13]. A thorough optimization using an iteratively restarted Nelder-Meade simplex algorithm of the scale parameters (minimum, maximum, and number) was performed, in order to fairly compare this measure to the proposed RORPO filter.

A standard non-linear tubular structure filtering. We also compare the RORPO filter with another standard filter obtained by composing the Robust Path Opening (RPO) with a top-hat operator. This filter is applied for several values of lengths and noise tolerance, that are the same involved in the RORPO filter, thus leading to a multiscale approach.

Quality scores. In order to compare the results obtained by these three methods on synthetic data (Sec. 5.2), we considered ROC curves and two similarity criteria: Matthews Correlation Coefficient (MCC) and Dice Coefficient, defined as

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (20)$$

$$Dice = \frac{2TP}{TP + FN + TN + FP} \quad (21)$$

where TP , TN , FP , FN are true/false positives/negatives, respectively.

As we deal with sparse features, the set of pixels belonging to the ground truth object (TP_{GT}) is always much smaller than the set of pixels belonging to its background (TN_{GT}). The considered ROC curves consider the true positive rate (TPR) as a function of the false positive rate (FPR), where both the TPR and the FPR are defined with respect to the ground truth object, *i.e.*, $TPR = \frac{TP}{TP_{GT}}$ and $FPR = \frac{FP}{TP_{GT}}$. Consequently, the FPR could well exceed 100%, but such values are uninteresting.

5.2 Synthetic Images

In these first experiments, we use a synthetic image containing both tubular, planar and blob structures (see Fig. 3(a,b)). A tubular structure is represented by a 3D simple helix with decreasing intensity, diameter and thickness. As a consequence, tube-like structures are represented in many orientations and with varying dimensions and degrees of tortuosity. Blobs were added using an additive Gaussian random field. The ground truth is a binary representation of the helix (see Fig. 3(c)).

The results of the three filters on this image are shown in Fig. 4. Qualitatively, we observe that RORPO is the only filter that fully removes the plane structure. The result of the vesselness is also much “noisier” than the other two. Indeed, these two morphological filters have the property that they always reduce the grey-level of the input image. This suppressive behaviour leads to cleaner images as seen on Fig. 4(a,c). In contrast, the vesselness does not share such a mathematical property. The visual analysis, that emphasises the good performance of RORPO filter is confirmed by the CCM and Dice measures.

Then, ROC curves have been computed from these filtered images by successive thresholding. To account for the 1-dimensional nature of the structures of interest, a skeletonization of the thresholded result was performed. The results, illustrated in Fig. 5(a), confirm that the RPO-top-hat is hindered by its inability to remove the planar structures. In this specific case, our RORPO filter outperforms Frangi’s vesselness,

The last experiment quantifies the robustness of the RORPO filter with respect to noise. Various amounts of white Gaussian noise were added to the synthetic image,

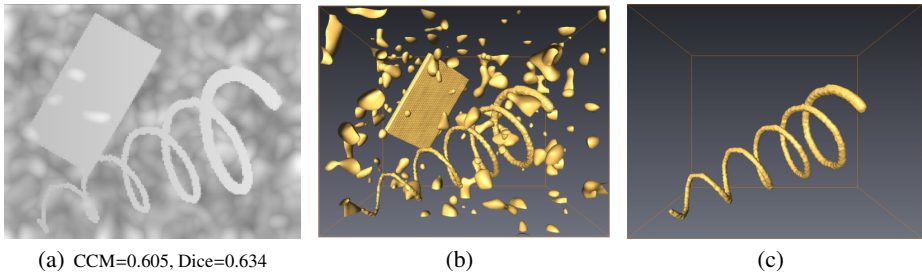


Fig. 3. Synthetic image: (a) maximum intensity projection and (b) isosurface. (c) Ground truth. Optimal CCM and Dice scores obtained by thresholding are indicated for (a).




Fig. 4. Filtered synthetic image: maximum intensity projection. (a) RORPO. (b) Frangi's vesselness. (c) and RPO-top-hat. Optimal CCM and Dice scores obtained by thresholding are indicated for (a–c).




Fig. 5. ROC curves on synthetic data (a,b) and on real data (c). (a) Comparison of the three filters, plus the native image. (b) Noise robustness of the RORPO filter. (c) Comparison between RORPO and Frangi's vesselness on the Rotterdam repository. For both filtering, the central curve is the mean ROC curve and the other two are the mean plus or minus one standard deviation curves.

and ROC curves were computed in each case. The results are shown in Fig. 5(b), and partially illustrated in Fig. 6. In order to better take into account the quantitative effects of the noise, no skeletonization was performed in this evaluation. The most noticeable observation is the high robustness of the RORPO filter for low level noise. In particular, up to a standard deviation of 5 almost no difference can be observed.




Fig. 6. (a–c) Synthetic image for various amounts of Gaussian noise: maximum intensity projection. (d–f) Corresponding filtered images.

5.3 Real Images

Finally we assessed the behaviour of the RORPO filter versus Frangi’s vesselness, on real 3D images. To this end, we consider Magnetic Resonance Angiography (MRA) and Computed Tomography Angiography (CTA) images, that are among the most challenging for tubular pattern detection, due to the presence of acquisition noise, physiological artifacts, neighbouring anatomical structures, *etc.* The actual difficulties induced by vessel enhancement and detection from such data have motivated many contributions until now [20,35].

Brain MRA images. We first consider Time-of-Flight MRAs of the cerebral arterial tree, containing topologically and geometrically complex linear structures with varying diameter and intensity. Moreover, they are corrupted by various artifacts, and present a low SNR, since sequences capture the signal of moving structures, instead of stationary tissues [24]. The images considered here were acquired on a whole-body scanner (Siemens Magnetom Verio 3.0 T, gradient slope = 200 T/m/s, flow encoding sequence, TR = 42.7 ms, TE = 6.57 ms, resolution = $0.4 \times 0.4 \times 0.8 \text{ mm}^3$). A 3D view of such MRA image is given in Fig. 7(a). The two filtering results are shown in Fig. 7(b–c).

Visually, Frangi’s vesselness seems richer, and RORPO more accurate. This is a consequence of Frangi’s enhancing behaviour, which in this case leads to false positives. Conversely RORPO suppressive behaviour may lead to potential false negatives. As well, the size of the vessels is overestimated by the vesselness (due to the Gaussian kernel convolution), while RORPO preserves the geometry of the vessels.




Fig. 7. Brain arteries. (a) Time-of-flight magnetic resonance angiography, viewed in volume rendering (from low intensities, in red, to high intensities, in yellow). (b) RORPO filter: 8 lengths from $L = 30\text{px}$ to 130px ; computation time: 32mn. (c) Frangi's vesselness: 4 scales from 1 to 3; computation time: 13mn.




Fig. 8. Coronary arteries. (a) Computed tomography angiography, viewed in volume rendering (from low intensities, in red, to high intensities, in yellow); the coronary medial axes are outlined in white. (b) RORPO filter: 4 scales from $L = 20\text{px}$ to 50px ; computation time: 50mn. (c) Frangi's vesselness: 4 scales from 0.4 to 1.8; computation time: 6mn.

Heart CTA images. We then consider CTAs of the coronary arteries of the heart from the Rotterdam repository [30]. These vascular structures are much simpler than cerebral MRAs. Coronary arteries are composed of only a few branches organised as a tree structure with homogeneous contrast due to contrast agent injection. Similar remarks can be derived as with the above MRA data. We observe that the RORPO filter leads to a noisier result than before, since it detects some parts of the origin of the aorta. The results obtained from Frangi's vesselness are also less satisfactory than for MRAs, essentially due to the presence of dense irregular tissues of the heart generating many false positives.

Coronary centerline ground truth is available in the Rotterdam database. ROC curves were computed on the 17 images of the training set of the Rotterdam repository. Two ground-truth images did not overlap a large part of the coronary tree for two patients so we chose to remove these images from the comparison study. We show in Fig. 5(c) results of the 15 ROC curves comparing RORPO and Frangi's Vesselness filtering. For each

filtering method, the mean ROC curve was computed after normalizing each curve to 100% of true positive (TPR). Results on the 15 patients are almost identical. RORPO filter always gives better results for lower values of false positive than Frangi's Vesselness. The latter, in some cases, better recovers coronaries but only for unusably large values of false positive.

6 Discussion and Conclusion

In this article, we propose a filter aimed at removing non-tubular structures. Leveraging from eigen-analysis principles, we make use of one important idea: instead of relying on *local* features (the three eigenvalues of the Hessian) to distinguish tubular structures from planes or blobs, we use *semi-global* features: paths. As information are no longer local, we are now able to analyze the responses corresponding to the seven cones partitioning the 3D space. By ordering these responses from highest to lowest, we observe that tubular structures are absent from the fourth rank down. The difference between the maximum response and this fourth rank yields a resulting image where only tubular structures are present. Both qualitative and quantitative analysis demonstrate the promising capacities of our novel RORPO filter for tubular structures, as well as its noise robustness. In particular, it discriminates tubes from planar objects much better than previous top-hat-based approaches. Compared to Frangi's vesselness, the RORPO image contains noticeably fewer false positives and does not overemphasize structures.

However, the vesselness response is not limited to tubular structures and provides more information, such as orientation and scale. The RORPO scale we used in this paper is related to the length of the structures and not to the scale of observation as in the Gaussian scale-space. It would be interesting to complete the RORPO filter with scale and orientation information. This will be the topic of further research.

While our filter preserves tubular structures, the resulting image sometimes present some disconnections at junctions (see for example Fig. 8(b)). This is because a junction is not tubular but more like a blob. A future research direction is to cope with this problem, for example with a finer study of the direction behaviour at such connection points. The RORPO filter should be useful as a low level input in any pipeline for network extraction (such as coronary arteries segmentation), as a complement to other features, including vesselness measures. In future work, we aim to build such a complete segmentation algorithm, in the case of the cerebral vascular network.

A C++ implementation (with a test image and some usage explanations), as well as videos showing both Frangi's Vesselness and RORPO filtering on both the Heart and Brain images, are available at <http://path-openings.github.io/RORPO/>.

References

1. Bauer, C., Bischof, H.: A novel approach for detection of tubular objects and its application to medical image analysis. In: Rigoll, G. (ed.) DAGM 2008. LNCS, vol. 5096, pp. 163–172. Springer, Heidelberg (2008)
2. Bennink, H.E., van Assen, H.C., Streekstra, G.J., ter Wee, R., Spaan, J.A.E., ter Haar Romeny, B.M.: A novel 3D multi-scale liness filter for vessel detection. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part II. LNCS, vol. 4792, pp. 436–443. Springer, Heidelberg (2007)

3. Bismuth, V., Vaillant, R., Talbot, H., Najman, L.: Curvilinear structure enhancement with the polygonal path image—application to guide-wire segmentation in X-ray fluoroscopy. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part II. LNCS, vol. 7511, pp. 9–16. Springer, Heidelberg (2012)
4. Buckley, M., Talbot, H.: Flexible linear openings and closings. In: ISMM. Computer Imaging and Vision, vol. 18, pp. 109–118. Kluwer (2000)
5. Chapman, B.E., Parker, D.L.: 3D multi-scale vessel enhancement filtering based on curvature measurements: Application to time-of-flight MRA. *Med. Image Anal.* 9, 191–208 (2005)
6. Cohen, L., Deschamps, T.: Segmentation of 3D tubular objects with adaptive front propagation and minimal tree extraction for 3D medical imaging. *Comput. Meth. Biomech. Biomed. Eng.* 10, 289–305 (2007)
7. Cokelaer, F., Talbot, H., Chanussot, J.: Efficient robust d-dimensional path operators. *J. Sel. Top. Signal* 6, 830–839 (2012)
8. Danielsson, P.E., Lin, Q., Ye, Q.Z.: Efficient detection of second-degree variations in 2D and 3D images. *J. Vis. Commun. Image R* 12, 255–305 (2001)
9. Derpanis, K.G., Gryn, J.M.: Three-dimensional nth derivative of Gaussian separable steerable filters. In: ICIP, vol. 3, pp. 553–556 (2005)
10. Dokládal, P., Jeulin, D.: 3-D extraction of fibres from microtomographic images of fibre-reinforced composite materials. In: Wilkinson, M.H.F., Roerdink, J.B.T.M. (eds.) ISMM 2009. LNCS, vol. 5720, pp. 126–136. Springer, Heidelberg (2009)
11. Du, Y.P., Parker, D.L., Davis, W.L.: Vessel enhancement filtering in three-dimensional MR angiography. *JMRI-J Magn. Reson. Im.* 5, 353–359 (1995)
12. Dufour, A., Tankyevych, O., Naegel, B., Talbot, H., Ronse, C., Baruthio, J., Dokládal, P., Passat, N.: Filtering and segmentation of 3D angiographic data: Advances based on mathematical morphology. *Med. Image. Anal.* 17, 147–164 (2013)
13. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A.C.F., Delp, S.L. (eds.) MICCAI 1998. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
14. Freeman, W.T., Adelson, E.H.: The design and use of steerable filters. *IEEE T. Pattern Anal.* 13, 891–906 (1991)
15. González, G., Aguet, F., Fleuret, F., Unser, M., Fua, P.: Steerable features for statistical 3D dendrite detection. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) MICCAI 2009, Part II. LNCS, vol. 5762, pp. 625–632. Springer, Heidelberg (2009)
16. Heijmans, H.J.A.M., Buckley, M., Talbot, H.: Path openings and closings. *J. Math. Imaging Vis.* 22, 107–119 (2005)
17. Koller, T., Gerig, G., Székely, G., Dettwiler, D.: Multiscale detection of curvilinear structures in 2D and 3D image data. In: ICCV, pp. 864–869 (1995)
18. Krissian, K., Malandain, G., Ayache, N., Vaillant, R., Troussot, Y.: Model-based detection of tubular structures in 3D images. *Comput. Vis. Image Und.* 80, 130–171 (2000)
19. Law, M.W.K., Chung, A.C.S.: Three dimensional curvilinear structure detection using optimally oriented flux. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part IV. LNCS, vol. 5305, pp. 368–382. Springer, Heidelberg (2008)
20. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. *Med. Image Anal.* 13, 819–845 (2009)
21. Lorenz, C., Carlsen, I.C., Buzug, T.M., Fassnacht, C., Weese, J.: Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2D and 3D medical images. In: Troccaz, J., Mösges, R., Grimson, W.E.L. (eds.) CVRMed-MRCAS 1997, CVRMed 1997, and MRCAS 1997. LNCS, vol. 1205, pp. 233–242. Springer, Heidelberg (1997)
22. Luengo Hendriks, C.L.: Constrained and dimensionality-independent path openings. *IEEE T. Image Process.* 19, 1587–1595 (2010)

23. Manniesing, R., Viergever, M.A., Niessen, W.J.: Vessel enhancing diffusion: A scale space representation of vessel structures. *Med. Image Anal.* 10, 815–825 (2006)
24. Miyazaki, M., Lee, V.S.: Nonenhanced MR angiography. *Radiology* 248, 20–43 (2008)
25. Naegel, B., Passat, N., Ronse, C.: Grey-level hit-or-miss transforms—Part II: Application to angiographic image processing. *Pattern Recogn.* 40, 648–658 (2007)
26. Najman, L., Talbot, H. (eds.): *Mathematical morphology: From theory to applications*. ISTE/John Wiley & Sons (2010)
27. Orkisz, M., Hernández-Hoyos, M., Douek, P., Magnin, I.: Advances of blood vessel morphology analysis in 3D Magnetic Resonance Images. *Machine Graph. Vis.* 9, 463–472 (2000)
28. Rouchdy, Y., Cohen, L.D.: Image segmentation by geodesic voting. Application to the extraction of tree structures from confocal microscope images. In: *ICPR*, pp. 1–5 (2008)
29. Sato, Y., Nakajima, S., Atsumi, H., Koller, T., Gerig, G., Yoshida, S., Kikinis, R.: 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In: Troccaz, J., Mösges, R., Grimson, W.E.L. (eds.) *CVRMed-MRCAS 1997, CVRMed 1997, and MRCAS 1997*. LNCS, vol. 1205, pp. 213–222. Springer, Heidelberg (1997)
30. Schaap, M., Metz, C.T., van Walsum, T., van der Giessen, A.G., Weustink, A.C., Mollet, N.R.A., Bauer, C., Bogunović, H., Castro, C., Deng, X., Dikici, E., O'Donnell, T., Frenay, M., Friman, O., Hernández Hoyos, M., Kitslaar, P.H., Krissian, K., Kühnel, C., Luengo-Oroz, M.A., Orkisz, M., Smedby, Ö., Styner, M., Szymczak, A., Tek, H., Wang, C., Warfield, S.K., Zambal, S., Zhang, Y., Krestin, G.P., Niessen, W.J.: Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Med. Image Anal.* 13, 701–714 (2009)
31. Soille, P., Breen, E., Jones, R.: Recursive implementation of erosions and dilations along discrete lines at arbitrary angles. *IEEE T. Pattern Anal.* 18, 562–567 (1996)
32. Soille, P., Talbot, H.: Directional morphological filtering. *IEEE T. Pattern Anal.* 23, 1313–1329 (2001)
33. Talbot, H., Appleton, B.: Efficient complete and incomplete path openings and closings. *Image Vision Comput.* 25, 416–425 (2007)
34. Tankyevych, O., Talbot, H., Passat, N.: Semi-connections and hierarchies. In: Hendriks, C.L.L., Borgefors, G., Strand, R. (eds.) *ISMM 2013*. LNCS, vol. 7883, pp. 159–170. Springer, Heidelberg (2013)
35. Tankyevych, O., Talbot, H., Passat, N., Musacchio, M., Lagneau, M.: Angiographic image analysis. In: *Medical Image Processing: Techniques and Applications*, ch. 6, pp. 115–144. Springer (2011)
36. Tankyevych, O., Talbot, H., Dokládál, P.: Curvilinear morpho-Hessian filter. In: *ISBI*, pp. 1011–1014 (2008)
37. Vincent, L.: Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE T. Image Process.* 2, 176–201 (1993)
38. Vincent, L.: Minimal path algorithms for the robust detection of linear features in gray images. In: *ISMM. Computer Imaging and Vision*, vol. 12, pp. 331–338. Kluwer (1998)
39. Wilkinson, M.H.F., Westenberg, M.A.: Shape preserving filament enhancement filtering. In: Niessen, W.J., Viergever, M.A. (eds.) *MICCAI 2001*. LNCS, vol. 2208, pp. 770–777. Springer, Heidelberg (2001)
40. Xiao, C., Staring, M., Shamonin, D.P., Reiber, J.H.C., Stolk, J., Stoel, B.C.: A strain energy filter for 3D vessel enhancement with application to pulmonary CT images. *Med. Image Anal.* 15, 112–124 (2011)
41. Xiao, C., Staring, M., Wang, Y., Shamonin, D.P., Stoel, B.C.: Multiscale bi-Gaussian filter for adjacent curvilinear structures detection with application to vasculature images. *IEEE T. Image Process.* 22, 174–188 (2013)
42. Zana, F., Klein, J.C.: Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE T. Image Process.* 10, 1010–1019 (2001)