



HAL
open science

RORPO : A morphological framework for curvilinear structure analysis; Application to the filtering and segmentation of blood vessels

Odysée Merveille

► **To cite this version:**

Odysée Merveille. RORPO : A morphological framework for curvilinear structure analysis; Application to the filtering and segmentation of blood vessels. Computation and Language [cs.CL]. Université Paris-Est, 2016. English. NNT : 2016PESC1056 . tel-01531160

HAL Id: tel-01531160

<https://pastel.hal.science/tel-01531160>

Submitted on 1 Jun 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.

UNIVERSITÉ — — PARIS-EST

Université Paris-Est
Laboratoire d'Informatique Gaspard Monge (LIGM)

Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Computer Science

Odyssée MERVEILLE

**RORPO: A morphological framework for
curvilinear structure analysis.**

**Application to the filtering and segmentation of
blood vessels**

Hugues TALBOT	<i>Director</i>
Nicolas PASSAT	<i>Co-Director</i>
Isabelle BLOCH	<i>Reviewer</i>
Grégoire MALANDAIN	<i>Reviewer</i>
Laurent COHEN	<i>Examinator</i>
Nelly PUSTELNIK	<i>Examinator</i>
Régis VAILLANT	<i>Examinator</i>

21 November 2016

ABSTRACT

The analysis of curvilinear structures in 3D images is a complex and challenging task. Curvilinear structures are thin, easily corrupted by noise, and can have a complex geometry. Despite the numerous applications in material science, remote sensing and medical imaging, and the large number of dedicated methods developed the last few years, the detection of such structures remains a difficult problem.

In this work, we provide an analysis of curvilinear structures. We first propose a new framework, called RORPO, to characterize such structures via two features: an intensity feature, which preserves the intensity of curvilinear structures, while decreasing the intensity of other structures; and a directional feature, providing at each point the direction of the curvilinear structure. RORPO, unlike state-of-the-art methods, is a non-local and non-linear framework, that is better adapted to the intrinsic anisotropy of curvilinear structures. RORPO is based on recent advances in Mathematical Morphology: the path operators.

We provide a full description of the structural and algorithmic details of RORPO, and we also conduct a quantitative comparative study of our features with three popular curvilinear structure analysis filters: the Frangi Vesselness, the Optimally Oriented Flux, and the Hybrid Diffusion with Continuous Switch.

Besides the straightforward filtering applications, both RORPO features can be used as priors to characterize curvilinear structures. We propose a regularization term for variational segmentation which embed these features. Classical regularization terms are not adapted to curvilinear structures, and usually lead to the loss of most of the low contrasted ones. Based on the RORPO features, we propose to only regularize curvilinear structures along their main axis. This directional regularization better preserves curvilinear structures, but also reconnects some of the parts of these structures that may have been disconnected by noise.

We present results of the segmentation of retinal images with the Chan *et al.* model either with the classical total variation, or with our directional regularization. This confirms that our regularization term is better suited for images with curvilinear structures.

RÉSUMÉ

L'analyse des structures curvilignes en 3 dimensions est un problème difficile en analyse d'images. En effet, ces structures sont fines, facilement corrompues par le bruit et présentent une géométrie complexe. Depuis plusieurs années, de nombreuses méthodes spécialement dédiées au traitement d'images contenant des structures curvilignes ont vu le jour. Ces méthodes concernent diverses applications en science des matériaux, télédétection ou encore en imagerie médicale. Malgré cela, l'analyse des structures curvilignes demeure une tâche complexe.

L'objet de ce manuscrit est la caractérisation des structures curvilignes pour l'analyse d'images. Nous proposons en premier lieu une nouvelle méthode appelée RORPO, à partir de laquelle deux caractéristiques peuvent être calculées. La première est une caractéristique d'intensité, qui préserve l'intensité des structures curvilignes tout en réduisant celle des autres structures. La deuxième est une caractéristique de direction, qui fournit en chaque point d'une image, la direction (sous la forme d'un vecteur) d'une structure curviligne potentielle.

RORPO, contrairement à la plupart des méthodes de la littérature, est une méthode non locale, non linéaire et mieux adaptée à l'anisotropie intrinsèque des structures curvilignes. Le fondement de notre méthode repose sur une notion récente de Morphologie Mathématique: les opérateurs par chemins.

Dans cette thèse, nous proposons une description complète de RORPO et de ses propriétés algorithmiques. Nous menons aussi une étude quantitative de ses deux caractéristiques en les comparant avec trois autres filtres populaires de détection de structures curvilignes: La Vesselness de Frangi, les Optimally Oriented Flux (OOF) et l'Hybrid Diffusion with Continuous Switch (HDSCS).

RORPO peut directement servir au filtrage d'images contenant des structures curvilignes, afin de spécifiquement les préserver, mais aussi de réduire le bruit. De plus, les deux caractéristiques de RORPO peuvent aussi être utilisées comme information a priori sur les structures curvilignes, afin d'être intégrées dans une méthode plus complexe d'analyse d'image. Nous proposons une telle application dans la deuxième partie de cette thèse. Nous concevons un terme de régularisation destiné à la segmentation variationnelle, qui est mieux adapté aux structures curvilignes que la plupart des autres termes de la littérature. Pour ce faire, nous intégrons les deux caractéristiques de RORPO dans cette régularisation, ce qui nous permet de régulariser les structures curvilignes seulement dans la direction de leur axe principal. Nous préservons ainsi mieux ces structures mais nous pouvons aussi reconnecter certaines structures curvilignes déconnectées par le bruit.

Les résultats de cette nouvelle régularisation sont présentés sur la segmentation de vaisseaux sanguins du fond d'oeil. Nous comparons aussi ces résultats avec le modèle de segmentation de Chan *et al.* ce qui nous permet de montrer que notre terme de régularisation est en effet bien mieux adapté aux structures curvilignes.

REMERCIEMENTS

Je tiens tout d'abord à remercier Nelly Pustelnik, Laurent Cohen et Régis Vaillant qui ont accepté de faire partie de mon jury, ainsi que Grégoire Malandain et Isabelle Bloch pour avoir également pris le temps de rapporter mon manuscrit de thèse.

Je voudrais remercier mes deux directeurs de recherche, Hugues et Nicolas pour m'avoir donné l'opportunité de faire cette thèse. J'aurais difficilement pu trouver de meilleures personnes pour m'initier à la recherche. Merci à Nicolas d'avoir pu se rendre disponible même à distance et à Hugues pour m'avoir guidée et (surtout) supportée au quotidien.

J'aimerais aussi remercier Laurent Najman pour son aide ponctuelle et ses conseils (presque) toujours à propos, qui m'ont permis de prendre du recul sur mon travail.

J'ai eu la chance de pouvoir travailler quelques mois dans deux laboratoires durant ma thèse: Tout d'abord dans le laboratoire ICube à Strasbourg, au sein de l'équipe MIV. Je voudrais ainsi remercier Benoît Naegel pour cette opportunité et l'ensemble de l'équipe pour m'avoir accueillie chaleureusement ; si chaleureusement d'ailleurs que j'y reste en post-doc. I also want to thank Leo Grady for welcoming me at HeartFlow for a month and Michiel Schaap who helped and advised me, always with patience.

Je remercie également l'ensemble des membres du projet VIVABRAIN. Faire partie d'un tel projet, mettant en jeu de nombreux acteurs et domaines de recherche différents, a été très enrichissant. Je remercie particulièrement Stéphanie Salmon et Olivia Miracourt, les vraies matheuses avec lesquelles j'ai eu la chance de collaborer directement.

Je voudrais remercier Sylvie Cach et Caroline Farmouza pour leur aide très précieuse concernant toutes les questions administratives au cours de ma thèse. Je remercie aussi l'ensemble des personnes du laboratoire à l'ESIEE pour la super ambiance au quotidien. C'était très agréable de venir tous les jours au labo avec le sourire. Un grand merci à Diane, Clara, Julie, Ali, Bruno, Stéphane, Kacper, Vincent, Vence et Yukiko pour les pauses café, les jeux, les pères Noël et surtout tous les gâteaux :). Merci à l'équipe jeux de rôles, Thibault, Laurent, Geoffrey, Mathieu, Timothée et Élodie pour toutes ces soirées qui me manquent déjà. Un merci spécial à Tim et Élo pour m'avoir recueillie et remonté le moral lorsque mon chéri n'était pas là. Je voudrais également remercier Clara, qui m'a hébergée pendant tout un mois dans son appartement à San Francisco, et a tout organisé pour que mon séjour se passe idéalement. Un énorme merci à mes deux copines de bureau, Élodie et Éloïse. Ça n'aurait jamais été pareil sans vous. Merci pour toutes les discussions, les pauses thé, le shopping, Disneyyyyyyy, et tous les bons souvenirs que je garderai même si je n'ai plus les surprises kinder.

Je voudrais enfin remercier toute ma formidable famille que je n'ai pas pu voir autant que je l'aurais voulu. Merci à mes parents, qui ont toujours été là, m'ont soutenue et montré leur fierté à tout moment. C'était très important. Merci à ma sœur Carla qui, même si elle est

loin et complètement surbookée, a toujours pris le temps de skyper avec moi. Merci à mon petit frère Loïc d'avoir pris son courage à deux mains et bravé les transports pour venir tout seul passer des supers week-end chez nous. Une petite pensée pour mon beau-père Philippe, qui a vécu en parallèle les galères de la thèse. Merci à mes grands parents, Alice, Gilbert, Henriette et René pour tout leur amour. Je remercie les meilleurs beaux-parents qui soient, Patricia et Christian, pour leur support, leur bonne humeur et l'intérêt qu'ils ont toujours porté à mon travail.

Et finalement, je remercie l'amour de ma vie, Florian, (et oui j'ai mis ton nom !). Tu as toujours su me soutenir, m'aider, me faire rire et surtout me rassurer comme personne. Tout aurait été beaucoup moins bien sans toi.

TABLE OF CONTENTS

1	INTRODUCTION	13
2	CURVILINEAR STRUCTURE FEATURES IN THE LITERATURE	19
2.1	Local approaches	19
2.1.1	Derivative-based	19
2.1.2	Filter banks	23
2.2	Non-local approaches	24
2.2.1	SE-based approaches	24
2.2.2	Connectivity-based approaches	25
2.2.3	Path-based approaches	26
I	TUBULAR FEATURES EXTRACTION	27
3	CONTEXT	29
3.1	Motivation	29
3.2	General Strategy	31
3.2.1	Geometric observation and oriented filters	31
3.2.2	The oriented filter choice	32
3.2.3	Alternative mathematical morphology based approach	33
4	PREVIOUS WORK ON PATH OPERATORS	35
4.1	Mathematical morphology basic operators	36
4.1.1	Structuring elements	36
4.1.2	Dilation and erosion	36
4.1.3	Opening and closing	39
4.2	Filtering curvilinear structures with line segments as structuring elements	41
4.3	Path Operators	42
4.3.1	Paths	43
4.3.2	Binary and grey-level path opening	44
4.3.3	Space discretization	45
4.3.4	Path opening algorithms	46
4.4	Robust to noise path opening	47
4.4.1	Incomplete path opening (IPO)	48
4.4.2	Robust path opening (RPO)	49
5	THE RORPO FRAMEWORK	53
5.1	Ranking orientations	54
5.1.1	Pointwise rank filter (PRF)	54
5.1.2	Counting high path opening responses	55
5.2	Intensity feature	55
5.2.1	The generic RORPO intensity operator	55
5.2.2	Setting the i_t threshold	57
5.2.3	Limit orientations in 3D	61
5.3	Directional feature	64
5.3.1	Finding the orientations of interest	64
5.3.2	Combining the orientations	66

5.4	Multiscale approach	66
5.4.1	Length based multiscale approach	67
5.4.2	Methodology	68
6	ALGORITHMIC CONSIDERATIONS	69
6.1	RORPO algorithm	70
6.2	Simplified robust path opening	70
6.3	Parameters	72
6.3.1	Path length	73
6.3.2	Noise robustness parameter	74
6.3.3	Orientations	76
6.4	Computational cost	77
7	RESULTS AND COMPARISONS	79
7.1	Comparison framework	80
7.1.1	Compared methods	80
7.1.2	Evaluation criteria	84
7.1.3	Parameters Optimization	86
7.2	Synthetic images	86
7.2.1	Intensity feature comparison	88
7.2.2	Directional feature comparison	89
7.3	Real images	91
7.3.1	Intensity feature comparison	91
7.3.2	Directional feature comparison	91
II	VARIATIONAL REGULARIZATION FOR CURVILINEAR STRUCTURE FILTERING AND SEGMENTATION	95
8	INTRODUCTION	97
9	PREVIOUS WORK	99
9.1	Inverse problem and the need of a regularization term	100
9.2	Previous work on total variation in image analysis	101
9.2.1	Tikhonov regularization	101
9.2.2	Total variation	102
9.2.3	TV <i>vs.</i> Tikhonov regularization	103
9.2.4	Solving the total variation	104
9.3	Segmentation model	106
9.3.1	The Chan <i>et al.</i> model	106
9.3.2	Proximal splitting algorithms	106
10	DIRECTIONAL TOTAL VARIATION	111
10.1	Motivation	112
10.2	The directional total variation	112
10.2.1	Directional gradient	113
10.2.2	Embedding the curvilinear structure features	115
10.3	Solving the segmentation model	117
10.4	Results and Comparisons	119
10.4.1	Compared methods	119
10.4.2	Synthetic image segmentation	119
10.4.3	Segmentation of the Drive database	122
11	CONCLUSION	125

Appendix	127
A NOTATIONS	129
A.1 Image	129
A.2 Dot product of images	129
A.3 Norms of an image	129
A.4 Gradient of an image	129
A.5 Divergence of an image	130
A.6 Norms of the gradient of an image	130
B SUBGRADIENT AND PROXIMITY OPERATOR	131
B.1 Subgradient	131
B.2 Proximity operator	133
B.3 Link between the subgradient and the proximity operator	134
C PROOF OF SOME STATEMENTS	137
C.1 TV of a differentiable function	137
C.2 TV of the indicator function	137
LIST OF PUBLICATIONS AND COMMUNICATIONS	139
ACRONYMS	141
BIBLIOGRAPHY	143

INTRODUCTION

Image analysis consists of extracting information from image data, usually by means of image processing techniques. With digital imaging becoming more prevalent, so has image analysis over the last few decades. Image analysis is used in many fields involving imaging technologies, including astronomy, medicine, robotics, biology and defense. Most image analysis methods require, at some point, the detection of various structures in images. These structures may be complex, like objects or passerby in a scene, or simpler, like round cells in microscopy images. In any case, suitable prior knowledge, characterizing the structure of interest, is usually required to be embedded in the detection method. Common priors describe the shape, color or pose of the structure.

One type of useful prior characterizes the geometry of the structure. In this thesis, we focus on the detection of structures with a specific geometry: the curvilinear structures. A curvilinear structure, in a n -dimensional (nD) image, is locally significantly longer in one dimension than in the $n - 1$ remaining dimensions (see Figure 1).







	Blob	Curvilinear	Plane
2D	 <p>$d_1 \simeq d_2$</p>	 <p>$d_1 \gg d_2$</p>	
3D	 <p>$d_1 \simeq d_2 \simeq d_3$</p>	 <p>$d_1 \gg d_2 \simeq d_3$</p>	 <p>$d_1 \simeq d_2 \gg d_3$</p>

Figure 1: Illustration of the different type of structures in the 2D and 3D space, according to their sizes in the different dimensions. The reader should keep in mind that these illustrations are simplified drawings based on the geometric properties of each structure. In real applications, curvilinear structures and plane-like structure are usually not straight.

Many image analysis applications involve the detection of curvilinear structures. For instance, one may want to extract roads in remote sensing images, or fibers in material science

images. In medical image analysis, the detection of curvilinear structures is often a common but crucial task: *e.g.* the segmentation of the respiratory tree or blood vessels both in 2D or 3D (see Figure 2).




Figure 2: Example of images involving curvilinear structures. (a) Glass fibers in CT-scan, (b) roads in remote sensing, (c) the airways in CT-scan, and (d) blood vessels in retinal images. (a) and (c) are slices of 3D images while (b) and (d) are pure 2D images.

Despite a large field of applications, curvilinear structure characterization remains a difficult task due to their specific geometry. First of all, the structures of interest are very sparse, which means that the ratio of pixels containing information on curvilinear structures is very small. This obviously complicates the task, in particular for statistical methods and learning approaches.

Moreover, this sparsity makes the evaluation of curvilinear structure detection algorithms very difficult. This is a frequently underestimated, but crucial, consequence. Indeed, without good performance measures, comparing and improving curvilinear structure detection methods is impossible.

Due to their low thickness, curvilinear structures are also more sensitive to noise and artifacts than other structures (see Figure 3). Even a small amount of noise may be sufficient to disrupt their contours, leading to disconnections or misconnections, which make the detection task even more challenging. Once again, classical quality scores do not measure the preservation of connectivity. Finally, curvilinear structures in real applications generally present complex geometries and topology. They can be more or less tortuous, present differ-

ent orientations and scales inside the same image, and also form a network (possibly with cycles), making geometric priors very difficult to use.




Figure 3: A few examples of challenging curvilinear structure detection. (a) An electron micrograph of glass fibers. The difficulty comes from their low-contrast and thinness. Illustration from [86].(b) an X-ray fluoroscopy taken during an angioplasty. The detection of the catheter, which is highlighted in green in (c), is interesting to track the surgery procedure.

In this work, we address the general problem of curvilinear structure characterization and detection, both through methodological developments and their applications to medical image analysis. We only focus on solid curvilinear structures, *i.e.* curvilinear structures which are not hollow.

In the first part of this work, we propose a new non-linear operator, called the Ranking Orientation Responses of Path Operators (RORPO), built upon the notion of path operator in mathematical morphology. From this new operator, two curvilinear characteristics can be computed: an intensity feature, that can be seen as a curvilinearity measure, and a directional feature, providing a local orientation estimation of the curvilinear structures. Both features are low-level curvilinear structure characteristics, and are essentially devoted to be embedded in more sophisticated image processing methods, like filtering or segmentation.

Secondly, we propose a way to use these features in a variational segmentation framework, where both the intensity and directional features are embedded in a novel directional regularization term. The results constitute promising perspectives to address the specific problem of blood vessel segmentation in brain MRA images, which is an essential part of the VIVABRAIN project.

THE VIVABRAIN PROJECT

Vascular diseases such as stenoses, aneurysms or arterio-venous malformations, affect millions of people worldwide. Their diagnosis is of utmost importance, and often involves medical imaging as angiography. Nonetheless, angiographic images are complex and time consuming to analyze manually. Since the mid 1990, various tools have been developed to assist physicians with this task [30, 88]. However this still remains an open problem, particularly when dealing with 3D images.

The VIVABRAIN project¹ is a pluridisciplinary project involving several research fields like computer science, mathematics, physics and medicine. The VIVABRAIN project aims at providing tools and data, allowing for a better understanding of Magnetic Resonance Angiography (MRA) images of vascular networks, especially in the human brain. This project is articulated around five main steps:

- The acquisition of MRA images;
- The extraction of the vascular volume;
- The design of a vascular model from the volume;
- The simulation of the blood flow inside the model;
- The simulation of MRA acquisitions from the model and the flow simulations.

In addition to the tools developed during this five-year project, the final goal of VIVABRAIN is to provide to the scientific community a set of MRA images along with their ground truth (see Figure 4). Indeed, the provided MRA will be the simulated images generated from the vascular model (the ground truth) and the flow simulations. With these data, scientists will be able to validate and compare their algorithms on real complex images. The Brainweb project [19] proposed a similar framework for standard T1 or T2 Magnetic Resonance Imaging (MRI). To the best of our knowledge, there does not exist publicly available ground truth of the brain vascular network.




Figure 4: The VIVABRAIN Project.

The work presented in this thesis was funded by the Agence Nationale de la Recherche (ANR) that supports the VIVABRAIN project. It is a subset of the second step of the project that is dealing with the filtering and segmentation of the brain vascular network.

¹ <http://vivabrain.fr>

THESIS CONTEXT

This PhD work started in October 2013, and took place under the supervision of Hugues Talbot and Nicolas Passat. Most of the time I worked at ESIEE Paris in the Laboratoire d'Informatique Gaspard Monge (LIGM). However I had the opportunity to spend a few months in other laboratories.

In June 2014, Leo Grady invited me to spend a month at HeartFlow² in California, where I worked under the supervision of Michiel Schaap on the segmentation of the coronaries on CT-scan images.

I also spent two months, in November 2014 and June 2015, in Strasbourg at the ICube laboratory to work on the VIVABRAIN project. I collaborated with Benoît Naegel on the filtering and segmentation of the MRA images of the project, using the tools I developed.

During these three years, with collaborators, I published four articles in international conferences (ECCV 2014, ISMM 2015, ICIIP 2016 and MICCAI 2016) and a journal article has undergone a first round of revision with IEEE PAMI. The reader may find a more exhaustive list in the List of Publications and Communications, at the end of this manuscript.

OUTLINE

In the next chapter, we review the state-of-the art on the filtering and segmentation of curvilinear structures. Then, this manuscript is divided into two parts.

The first part presents our RORPO operator. We start by introducing the general context and motivation of this new curvilinear structure detector in Chapter 3, followed by a detailed presentation of the previous work related to path operators in Chapter 4. Chapter 5 presents the framework we developed, and the two features derived from RORPO. Algorithmic considerations on the parameters, robustness and computation time are discussed in Chapter 6. Finally we present a validation and comparison study in Chapter 7.

The second part of the manuscript presents the regularization term we propose for variational segmentation of curvilinear structures. After a short introduction in Chapter 8, Chapter 9 presents the previous work on the total variation that is the basis of our regularization term. We present our methodology and segmentation results in Chapter 10. Finally a conclusion and the future work are presented in Chapter 11.

² <https://www.heartflow.com/>

CURVILINEAR STRUCTURE FEATURES IN THE LITERATURE

Curvilinear structure filtering and segmentation in 3D images is an active research area, which has led to the development of many methods over the last two decades. Most methods use specific information to detect curvilinear structures. Such information can be an assumption on the shape, geometry or intensity of curvilinear structures, which is directly used to filter or segment the images. Alternatively, it can correspond to features provided by some other methods embedded into a higher-level filtering or segmentation framework. In this chapter, we present an overview of the different assumptions and features used in the literature to specifically deal with curvilinear structures in images. The reader may refer to [51] [85] for exhaustive reviews on the subject.

Classifying the various approaches that deal with curvilinear structures is not an easy task. Most methods combine different strategies to achieve their goal, which renders a clear and consistent classification of these methods difficult. Nonetheless, we decided in this chapter to divide the different approaches in two classes. The first section presents methods that focus on local hypotheses characterizing curvilinear structures. These methods exploit local analyses of the intensity profile and orientation to detect curvilinear structures. The second section presents non-local methods, usually derived from mathematical morphology. They aim at detecting curvilinear structures by analyzing the image in larger neighborhoods.

2.1 LOCAL APPROACHES

2.1.1 *Derivative-based*

The vast majority of works dealing with curvilinear structures are based on the derivatives of the image. Indeed, one of the first straightforward assumption one can make to detect a structure in an image is to use the variation of its intensity profile.

The second-order derivatives represent the local curvature of the intensity of the image and the magnitude of each derivative is the degree of the intensity curvature along its axes. By studying the intensity curvature of a structure, one is able to distinguish between plane-like, blob-like and curvilinear structures. Indeed, a blob-like structure, at the right scale, has three high intensity curvatures, while a plane-like structure has two high intensity curvatures, and a curvilinear structure only has one high intensity curvature (see Figure 5.(a-c)).

However, structures are usually not oriented along one of the 3D image axes. In such general conditions, an approach consists of performing a spectral analysis of the Hessian

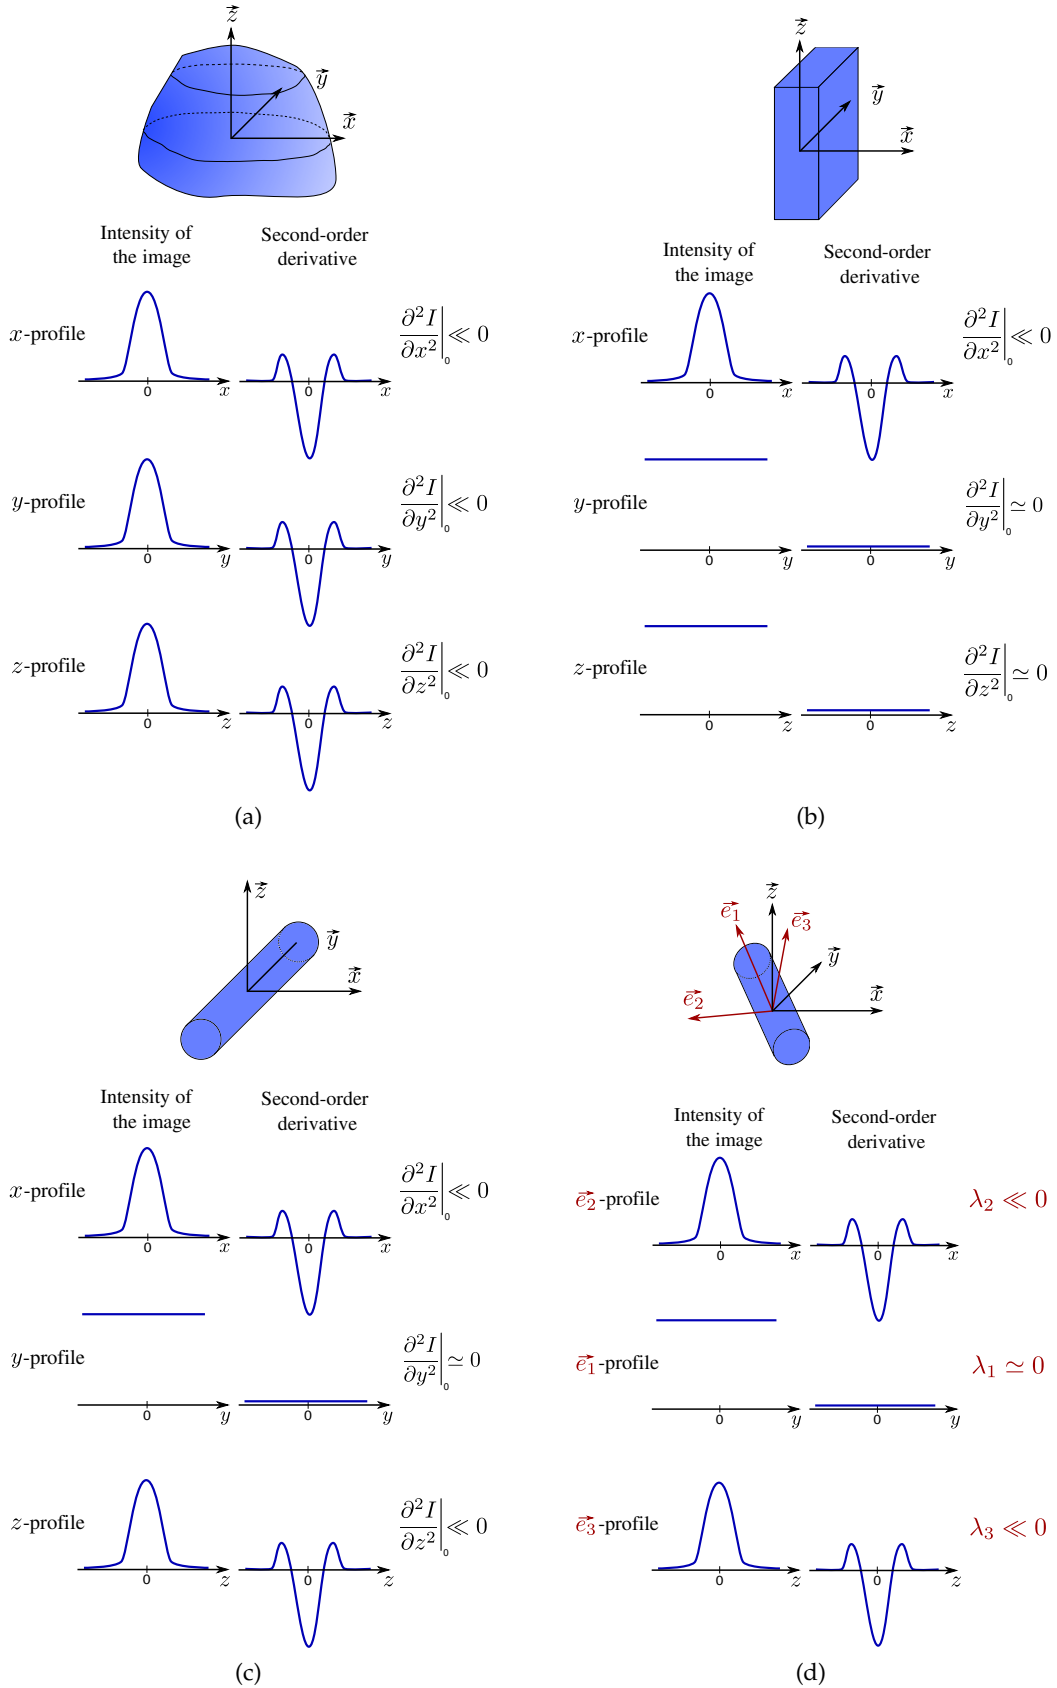


Figure 5: Second-order derivative pattern of a blob (a), a plane (b) and a curvilinear structure (c) whose main intensity curvatures are oriented in the direction of the image axes. Each structure is characterized by a different set of second-order derivative magnitudes. (d) presents a curvilinear structure in an arbitrary orientation. The three eigenvalues of its Hessian matrix correspond to the magnitude of the second-order derivatives in the direction of the associated eigenvectors.

matrix (see Figure 5.(d)). The Hessian matrix, \mathcal{H} , is the matrix of all the second derivatives of the image:

$$\mathcal{H}(I_\sigma) = \begin{bmatrix} \frac{\partial^2 I_\sigma}{\partial x^2} & \frac{\partial^2 I_\sigma}{\partial x \partial y} & \frac{\partial^2 I_\sigma}{\partial x \partial z} \\ \frac{\partial^2 I_\sigma}{\partial y \partial x} & \frac{\partial^2 I_\sigma}{\partial y^2} & \frac{\partial^2 I_\sigma}{\partial y \partial z} \\ \frac{\partial^2 I_\sigma}{\partial z \partial x} & \frac{\partial^2 I_\sigma}{\partial z \partial y} & \frac{\partial^2 I_\sigma}{\partial z^2} \end{bmatrix} \quad (1)$$

where I_σ is the convolution of image I with a Gaussian of variance σ^2 , which introduces the notion of scale.

Numerous methods use the eigenvalues and eigenvectors of this Hessian matrix to compute measures characterizing curvilinear structures. These measures are often called *vesselness*, since a common application is the filtering of blood vessels. Lorenz *et al.* [53] proposed the first vesselness based on the two first eigenvalues (*i.e.* those with the highest magnitude). In [80], Sato *et al.* proposed another vesselness also based on the first two eigenvalues, but which has the ability to discriminate between curvilinear and plane-like structures. The most commonly used vesselness was proposed by Frangi *et al.* [33], and is based on the three eigenvalues of the Hessian matrix. More recently, Li *et al.* [52] proposed another vesselness that claims a better specificity than the Frangi's.

The first-order image derivatives (*i.e.* the gradient) are also used in edge detection, as in the pioneer work of Canny [12]. Based on the same principle as the Hessian matrix, several works use the structure tensor, J , that is built from the first-order derivatives:

$$J(I_\sigma) = \begin{bmatrix} \left(\frac{\partial I_\sigma}{\partial x}\right)^2 & \frac{\partial I_\sigma}{\partial x} \frac{\partial I_\sigma}{\partial y} & \frac{\partial I_\sigma}{\partial x} \frac{\partial I_\sigma}{\partial z} \\ \frac{\partial I_\sigma}{\partial y} \frac{\partial I_\sigma}{\partial x} & \left(\frac{\partial I_\sigma}{\partial y}\right)^2 & \frac{\partial I_\sigma}{\partial y} \frac{\partial I_\sigma}{\partial z} \\ \frac{\partial I_\sigma}{\partial z} \frac{\partial I_\sigma}{\partial x} & \frac{\partial I_\sigma}{\partial z} \frac{\partial I_\sigma}{\partial y} & \left(\frac{\partial I_\sigma}{\partial z}\right)^2 \end{bmatrix}$$

where I_σ is still the convolution of image I with a Gaussian of variance σ^2 .

The spectral analysis of the structure tensor also provides information on the geometry of the structure at each point of the image. However this analysis is different from the Hessian's. Instead of measuring curvature, the structure tensor measures the intensity variations along the three principal axes of a structure. In this context, a curvilinear structure is characterized by two large eigenvalues and a small one, whereas a plane-like and blob-like structure have one and three large eigenvalues respectively. The eigenvalues of the structure tensor are always positives both for bright and dark structures, which makes the distinction impossible unlike the Hessian eigenvalues. Nonetheless, the structure tensor based approaches are usually considered less noise sensitive than the Hessian based approaches, as they used lower-order derivatives.

Methods based on the structure tensor for curvilinear structure detection in 3D are typically more recent. In 2005, Agam and Wu [1] used the structure tensor to estimate the direction of blood vessels. A modified version of the structure tensor, called the *radial structure tensor*, was introduced by Wiemker [101] to better discriminate the various structures based on the correlation between the gradient and radial vectors. More recently, Moreno and

Smedby [63] also used the structure tensor and radial vectors, but instead of weighting the gradient using the radial vectors, they used it to remove the gradients which are not likely to belong to a curvilinear structure.

Many works combine several derivative-based approaches. Koller *et al.* [43] proposed an oriented filter based on the first derivatives. However, the direction of the structure is computed beforehand using the eigenvectors of the Hessian matrix. In [73] and [3], the curvilinear structure extraction is performed by analyzing the curvature and directions of structures computed from the Weingarten matrix. This matrix combines information from the first and second-order derivatives of the image.

An interesting approach was proposed by Bauer *et al.* [5] to tackle the problem of the blurring effect induced by the linear scale space (this point will be discussed in the Chapter 3). They proposed to compute the Hessian matrix based on the gradient vector field (GVF) [106] instead of the classical gradient. The Frangi vesselness is then computed, but requires only a single scale. Moreover, it also avoids the fusion of adjacent curvilinear structures, unlike classical Hessian-based vesselness (see Figure 6). Another contribution by Krissian *et al.* [46], proposed to use the eigenvalues of the Hessian matrix to perform a pre-selection of points likely to be located near the vessel axes.




Figure 6: A slice of a synthetic image containing two small close curvilinear structures on the left and one large curvilinear structure on the right (a). The results of a Gaussian filter with a small (b) and a larger (c) scale. The magnitude of the GVF (d). Without careful scale selection, applying a Gaussian filter may cause the two close curvilinear structures to merge. Only one curvilinear structure would be distinguishable from the resulting derivatives. Using the GVF dispenses with scale section, and allows the distinction between close curvilinear structures. (Images from [5])

Another class of approaches use the gradient within a diffusion framework to handle curvilinear structures. The first anisotropic diffusion filter for images was proposed by Perona and Malik [70]. They proposed to control the degree of diffusion by a function of the gradient magnitude. The idea was to apply a high degree of diffusion where the gradient is low, to smooth/reduce the noise, while reducing the diffusion in the presence of high gradient (*i.e.* edges) to preserve the contours. Based on similar principles, specific anisotropic diffu-

sion filters were proposed to enhance curvilinear structures. Weickert [100] proposed the Coherence-Enhancing Diffusion (CED) filter. In this method, the diffusion tensor is built using the structure tensor eigenvalues, so that the diffusion in a curvilinear structure preferably takes place in the direction of this structure. In 1997, Krissian *et al.* [45] also proposed an anisotropic diffusion filter where the diffusion tensor is built upon the direction of the main curvatures, and the magnitude of the smoothed gradient. Later, Krissian [44] proposed to replace the magnitude of the smoothed gradient by the first derivatives along each main intensity curvature and he added a data fidelity term. This results in a more specific filter which better preserves small curvilinear structures.

The Hessian matrix was also used to build the diffusion tensor. Cañero *et al.* [13] proposed such a filter, which also takes into account the image coherence by using both eigenvectors and eigenvalues of a normalized Hessian matrix. Manniesing *et al.* [56] designed a similar method extended to 3D and proposed a smooth version of the Frangi vesselness, which is required to have good scale-space properties.

A few methods propose to compute a flux on vector fields, based on the image derivatives, to specifically deal with curvilinear structures. Vasilevskiy and Siddiqi introduced an image gradient flux to drive the evolution of an active contour model [97]. Maximizing this flux on the curve makes it converge to locations where the curve normals are aligned with the gradient field. Later, Descoteaux *et al.* [27] computed a flux on a vector field which embed the Frangi vesselness directions in addition to the gradient field. Later, Law and Chung [49] developed the Optimally Oriented Flux (OOF). They defined the oriented flux as the flux of the image gradient projected along a direction onto a local sphere. The direction maximizing the outward oriented flux is the curvilinear structure direction. They showed that this minimization can be solved by an eigenvalue analysis of a specific matrix, and used these eigenvalues inside vesselness formulations to filter curvilinear structures. In 2010, Law and Chung [50] adapted their OOF to drive an active contour with better properties than the Vasilevskiy *et al.* model, in particular with regards to flow leaks. The OOF is also used as curvilinear structure features. For example, Benmansour and Cohen [8] used the eigenvalues of the OOF matrix to design an anisotropic metric better suited for the curvilinear structure extraction within a minimal path framework.

All of these methods, based on the derivatives, share the same drawbacks which we discuss in detail in the next chapter.

2.1.2 Filter banks

Filter banks are a general framework where the output results from the combination of a set of filters sharing some properties. In general, the filter banks designed for curvilinear structures are oriented, *i.e.* each filter of the bank is sensitive to a different orientation. By combining all the oriented filters, filter banks are able to detect structures in arbitrary orientations. Even if some filter banks are derivative-based, the unifying idea remains the combination of oriented filters. That is why we chose to dedicate a specific section to this class of filters.

Most filter banks dedicated to curvilinear structure detection were proposed for 2D images [4, 47, 92]. Indeed, they usually require to discretize space in many orientations, which can quickly become unmanageable in 3D. Nonetheless, Rangayyan *et al.* proposed a 3D filter

based on the Gabor wavelet to detect blood vessels [74].

Steerable filters [35], are another important type of filter banks. They are characterized by the structure of their responses. The response at any given orientation of a steerable filter is a linear combination of a small number of rotated versions of itself. For instance, the first-order derivatives of a Gaussian are steerable. The first-order Gaussian derivative along orientation θ , G_1^θ , can be written as a linear combination of the first-order Gaussian derivatives along axis x , G_1^0 , and along axis y , $G_1^{\frac{\pi}{2}}$:

$$G_1^\theta = \cos(\theta)G_1^0 + \sin(\theta)G_1^{\frac{\pi}{2}}$$

The Hessian filter and the OOF were also proved to be steerable [41, 8]. Steerable filters are most often used to detect 2D curvilinear structures, however a few contributions can handle 3D curvilinear structures. Amongst them, Gonzales *et al.* [36] proposed a steerable filter based on the second and fourth Gaussian derivatives, and use it to train a SVM classifier to detect dendrites.

Wavelet based filters, and in particular those designed to handle curvilinear structures, can also be seen as filter banks. Again, only a few of them can handle 3D images. Beamlets were proposed by Donoho *et al.* [28] to detect curvilinear structures in large 3D images. Later, Woiselle *et al.* [105] extended Beamlets to BeamCurvelets, which are able to deal with curvilinear structures of different sizes, unlike Beamlets.

2.2 NON-LOCAL APPROACHES

In this section, we discuss other types of methods, generally less known, based on non-local approaches. Non-local approaches focus on more global curvilinear structure characteristics. A first class of methods assume that curvilinear structures are composed of a basic shape called *Structuring Element (SE)*, whereas a second class of methods focus on characteristics that describe curvilinear structures more globally. A third and last type of approach models curvilinear structures with paths.

2.2.1 SE-based approaches

Many basic operators in mathematical morphology (erosions, dilations, openings, closings, *etc.*) use SEs (see Chapter 4 for a presentation). The advantage of such operators is that the SE shape can be precisely adapted to the structure one wants to preserve. Several approaches have been proposed to specifically handle curvilinear structures. A first approach consists of modeling curvilinear structures by small straight lines, and performs openings or closings with these SE in sufficiently many different orientations [90, 108]. The efficiency and robustness to noise of these approaches was studied in [82]. Another approach is to use a SE modeling both the curvilinear structure and its local background, and perform an hit-or-miss transform [69, 64].

An SE-based approach was also investigated in order to reconnect curvilinear structures which may have been disconnected by noise [87] (see Figure 7). This approach, called *Morpho-*

Hessian, computes the direction of blood vessels with an Hessian-based approach and performs modified openings or closings with segments oriented in these directions.




Figure 7: 2D image of a neurite (a) and the result from filtering with the Morpho-Hessian (b). Images from [87].

Even though these approaches showed good results, the fixed straight geometry of such SE remains a limitation for handling curvilinear structures which may present local tortuosities. This leads to the development of more flexible SEs [10], which we extensively present in Chapter 4.

2.2.2 Connectivity-based approaches

A few methods in mathematical morphology are based on the notion of connectivity on graphs. Instead of working with fixed-shape SE, these methods focus on connected components. This is the case of component-trees [78], which losslessly represent an image as a tree. Each node is a connected component and each level in the tree correspond to all the connected components of the image thresholded at a given intensity value (see Figure 8).

In this framework, it is possible to compute various attributes on the connected components, and further remove all the components of the image which do not reach a given value in an attribute (under certain conditions). Several attributes were investigated to characterize curvilinear structures. The majority considers scalar attributes; one can cite the ratio inertia/volume [103] or an elongation criteria [107]. Geodesic attributes were also proposed such as the geodesic tortuosity [62] or the geodesic elongation [48]. More recently, vectorial attributes were proposed, which consist of a vector of scalar attributes. A connected component is removed if its attribute vector significantly differs from a reference vector according to a designed distance measure [94]. Vector attributes were applied, for instance, on 3D PC-MRA of the brain by Caldairou *et al.* [11].

By construction, connected filters cannot split connected components, which may result in erroneous connections between curvilinear structures and artifacts, or at the junctions of

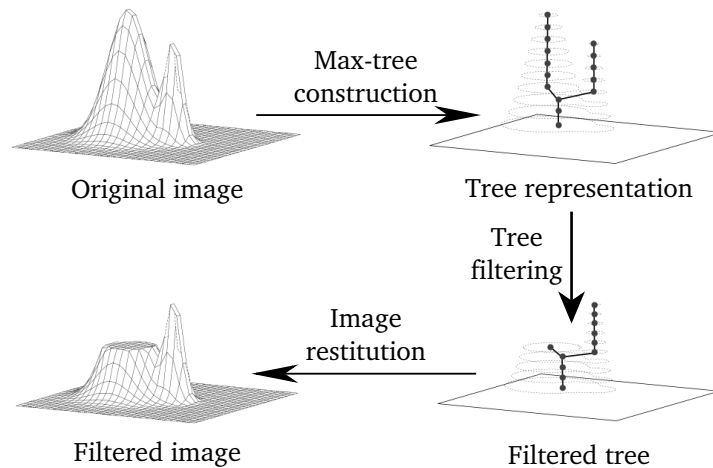


Figure 8: Filtering principle on component tree. Illustration from [79].

curvilinear structures. Some attempts to minimize these drawbacks were proposed, either via tiling approaches [11], or with asymmetric notions of connectivity [71].

2.2.3 Path-based approaches

The links between connectivity-based and path-based approaches (described below) were also investigated in [102].

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. By contrast, connectivity-based approaches lead to more global descriptors; unfortunately the anisotropy of curvilinear structures is not really taken into account. To address this problem, geodesic paths [20] were introduced to consider long-range, anisotropic neighborhoods, while still being flexible. A few approaches used the notion of geodesic paths on curvilinear structures [26, 104]. In these methods, metrics are designed so that a curvilinear structure corresponds to a geodesic path. These methods generally require some user interaction as at least the starting point must be set by the user.

Recently, Rouchdy and Cohen introduced a new curvilinear structure detector called *geodesic voting* [75]. The idea is to compute several geodesic paths from a starting point to different points in the image. The curvilinear structures are those which present a high geodesic path density. Geodesic voting was initially proposed in 2D, but recent applications were extended to 3D for blood vessel segmentation [76]. Inspired by the same idea, Bismuth *et al.* proposed the Polygonal Path Image (PPI) [9]; however, due to computational cost, only 2D implementations were proposed.

In 1998, Vincent proposed the pioneering notion of local optimal path [98]. He proposed to compute minimal paths inside local oriented cones. This led, to the development of flexible linear operators [10], and later path operators [39]. The reader may find an extensive presentation of the work on path operators in Section 4.3.

In the next chapter, we present the context and motivation of our work.

Part I

TUBULAR FEATURES EXTRACTION

CONTEXT

3.1 MOTIVATION

As part of the VIVABRAIN project, the guiding application for this work is the filtering and segmentation of 3D brain MRA.

In preliminary work, Dufour *et al.* [30] proposed a first segmentation method, based on a morpho-Hessian and component-trees approach. This method is able to reconnect vessels that appear disconnected due to noise or signal loss. However, some problems remain: (1) curved plane-like structures, probably from the scalp vascularization, are not removed during the filtering/segmentation process (see Figure 9); (2) the component-tree segmentation is not fully automated and requires the user to provide markers to initialize the segmentation.

Based on these observations, the first motivation of this work was to propose a solution to the above mentioned problems and improve the segmentation results.




Figure 9: Illustration of a plane-like structure (red ellipse) disturbing the segmentation on a brain MRA from the VIVABRAIN database.

In the previous chapters, we have seen that the analysis of curvilinear structures, like blood vessels, has attracted a lot of attention over the past few years, and that most of these works require a curvilinear prior. We call *low-level curvilinear feature* all information specifically characterizing curvilinear structures, such as their direction, diameter or position in the image. In particular, we call:

- *Intensity feature*, the map associating high values for pixels belonging to a curvilinear structure, and low values for pixels belonging to any other structure;
- *Directional feature*, the map that associates a vector providing the direction of a potential curvilinear structure to each point of an image.

The *vesselness* computed by most recent works on filtering and segmentation of blood vessels may be assimilated to an intensity feature. The main difference is that vesselness, additionally to the detection of curvilinear structures, also aims at detecting the junctions between vessels, which are not curvilinear but blob-like structures. This makes these filters specific to blood vessels and not general low-level curvilinear structure filters as we want to propose.

Among all the methods dealing with curvilinear structures, only a few provide low-level features, and even fewer work in practice in 3D. We have seen that the majority of 3D image filters able to characterize curvilinear structures while removing plane-like and blob-like structures are derivative-based filters. Even though they have interesting properties, they also present a few drawbacks. These may perturb any curvilinear structure related application like segmentation, tracking or classification. Among these drawbacks, we can list the followings:

- *The local neighborhood problem*

Derivative-based filters distinguish between structures based on a local neighborhood analysis. This analysis is often sufficient to make a good decision, but in some tricky cases, it fails. Particularly, the derivative-based directional feature and, to a lesser extent the intensity feature, may not be sufficiently reliable on the edge of curvilinear structures and lead to a lack of consistency in the extracted features (see Figure 10).




Figure 10: Illustration of false detections (a) and direction error estimation (b) coming from a local neighborhood analysis. The red squares represent the considered neighborhood centered on the small red dot. In figure (b), each blue square represent a pixel and the red arrow is the estimated direction based on the local neighborhood. The bottom arrow should be vertical but the local neighborhood analysis yields an incorrect direction estimation.

- *The linear scale-space multiscale problem*

The classical multiscale approach of most derivative-based filters relies on a linear scale-space paradigm. At each scale, a Gaussian filter is first applied, its standard deviation parameter representing the scale, then the derivative-based filter is computed. This approach allows these filters to detect curvilinear structures at multiple scales, but from a blurred version of the initial image. On this blurred image, the curvilinear structure

edges have been moved, resulting in possible incorrect detections. When curvilinear structures are brighter than their background, derivative-based filters typically tend to detect wider curvilinear structures than they actually are. Moreover, this approach can also connect structures which are not actually connected (see Figure 11).

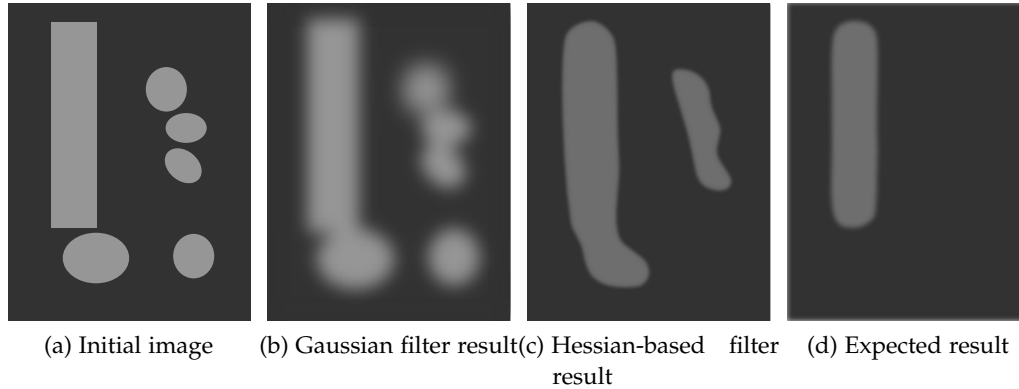


Figure 11: Illustration of the linear scale-space problem typical of many derivative-based filters. (b) the result of applying a Gaussian filter on (a). (c) the result of applying a derivative-based filter on (b). (d) the expected result of applying a curvilinear structure filter on (a). Because of the multiscale approach, structures that were not connected on the initial image become connected on the derivative-based filter result. Moreover, the curvilinear structure on the Hessian-based filter result is slightly wider than the expected result.

For all these reasons, the first motivation of our work was to propose a new curvilinear structure feature detector that could be an alternative to the derivative-based filters. This alternative should share the same advantages of derivative-based filters and avoid their drawbacks. More specifically, we aimed at:

- Enhancing curvilinear structures, while removing plane-like and blob-like structures;
- Extracting both intensity and directional features;
- Preserving curvilinear structure edges;
- Proposing a method working both in 2D and 3D;
- Relying on a non-local approach.

3.2 GENERAL STRATEGY

In this section, we explain our general strategy for distinguishing between curvilinear, plane-like and blob-like structures and motivate the choice of path operators.

3.2.1 Geometric observation and oriented filters

As presented in Chapter 1, the difference between curvilinear, plane-like and blob-like structures is geometric. It is based on the size of their main dimensions. A curvilinear structure is

longer in one of its dimensions, whereas plane-like structures are larger in two dimensions and blob-like structures have similar sizes in all dimensions. By counting the number of large sizes in all dimensions, one should be able to distinguish between the different structure types. To do so, we chose an approach based on oriented filters. We call *oriented filter* any filter which response depends on the orientation in which it is computed.

Let \mathcal{F} be an oriented filter and \mathcal{O} be a set of chosen orientations in the 3D space such that $\mathcal{F}^o(I)$, $o \in \mathcal{O}$, is the response of filter \mathcal{F} in the orientation o on image I . If a bright structure lies in orientation o , then $\mathcal{F}^o(I)$ will output a high value for this structure. Let us now consider a blob, a plane and a curvilinear structure, processed by such filter. The blob-like structure, as it lies in all orientations, presents a high response for all $\mathcal{F}^o(I)$. The plane-like structure will respond highly in almost all orientations except for those along its small size. In the same way, the curvilinear structure will only respond highly in a few orientations corresponding to those along its only large size (see Figure 12).

Our strategy to only preserve curvilinear structures relies on the geometric difference between structures. We propose to filter an image by an oriented filter and to classify each structures based on the number of high responses.




Figure 12: (a) A blob-like structure, (b) a plane-like structure and (c) a curvilinear structure in blue. Arrows represent examples of orientations $o \in \mathcal{O}$. A green (resp. red) arrow represents a high (resp. low) response of the oriented filter along this orientation. The blob-like structure responds in all orientations whereas the plane-like structure only responds in 4 out of the 8 orientations, and the curvilinear structure only responds in 1 orientation.

3.2.2 The oriented filter choice

Among the large choice of oriented filters in the literature, we selected the path operators. A detailed presentation of path operators is proposed in Chapter 4, but we focus, in this section, on our motivation to chose such an operator.

The major criterion we considered was the non-locality of path operators. Most curvilinear structure filters compute the response of a structure in an isotropic neighborhood, the size of which depends on a given scale, chosen according to the size of the structure of interest. This approach is not optimal for curvilinear structures, since they are highly anisotropic. In particular, it may lead to false detections (see Figure 10.(a)) and incorrect orientation estimations (see Figure 10.(b)), especially near the structure edges. Path operators, by computing

the response along an anisotropic neighborhood fitting in the curvilinear structure, avoid this pitfall.

Path operators are also edge-preserving. Edges are either preserved or removed but in any case they are neither moved nor blurred. This is a useful property for a filter designed to provide features for segmentation applications.

Contrary to the major part of oriented filters, path operators also handle local tortuosity. From a given orientation, path operators are able to detect structures even if they locally deviate from this orientation. This is especially useful as curvilinear structures are rarely perfectly straight in real applications.

Recent advances in path operators have proposed versions that are robust to noise which is also essential when dealing with real images. This noise robustness, combined with the non-locality approach, allows path operators to reconnect curvilinear structures that may appear disconnected due to noise.

Last but not least, path operators are almost parameter free. Indeed, the only real parameter is the path length, which is related to the length of the structures of interest. No other weight or parameter tuning is required.

3.2.3 *Alternative mathematical morphology based approach*

In a preliminary work, partially presented in [59], we first explored an alternative approach to distinguish between curvilinear structures and plane-like structures, based on mathematical morphology tools. This approach also relies on path operators but involves, in addition, radial openings. Figure 13 illustrates this approach. In part due to the incoherence between the radial and path operators, we showed that this type of approach was not as effective as purely path-based approaches to preserve curvilinear structures while removing other structures.




Figure 13: A path opening with path length L efficiently preserves a curvilinear structure (a) but also a plane-like structure (b). To remove the latter, a radial opening with a line of size $\frac{L}{k}$ is performed in the orthogonal plane P of the structure. This radial opening preserves the structures which are larger than $\frac{L}{k}$ i.e. the plane-like structures. A top-hat is then performed to only preserve curvilinear structures.

PREVIOUS WORK ON PATH OPERATORS

In this chapter, we present a brief overview of the basic operators and notions of mathematical morphology, that we will require in subsequent sections. For an up to date presentation of the topic, the reader may refer to [65]. Then, we propose an exhaustive review on the path operators.

Contents

4.1	Mathematical morphology basic operators	36
4.1.1	Structuring elements	36
4.1.2	Dilation and erosion	36
4.1.3	Opening and closing	39
4.2	Filtering curvilinear structures with line segments as structuring elements .	41
4.3	Path Operators	42
4.3.1	Paths	43
4.3.2	Binary and grey-level path opening	44
4.3.3	Space discretization	45
4.3.4	Path opening algorithms	46
4.4	Robust to noise path opening	47
4.4.1	Incomplete path opening (IPO)	48
4.4.2	Robust path opening (RPO)	49

4.1 MATHEMATICAL MORPHOLOGY BASIC OPERATORS

In the following, we consider an image I as a function $I : X \rightarrow G$, where X is the set of points of I and G is its set of possible values (grey-levels).

4.1.1 Structuring elements

In mathematical morphology, operators are often associated with SEs. A SE, B_p , is a set of generally connected image points, which is known *a priori* and whose origin point is p . This origin point does not necessarily belong to B_p .

SEs may present various shapes and sizes, depending on the application. If a structure with specific geometric properties is sought, one can adapt the shape, spatial orientation and size of the SE. For example, in road detection applications, using a set of line-like SEs may be a good choice.

The rotational symmetric of SE is often used in mathematical morphology in association with the notion of operator duality. The center of symmetry is the origin of the SE. In the following, we call \check{B}_p the symmetric B_p . Figure 14 shows examples of SE and their symmetric.




Figure 14: Examples of a few 2D structuring elements, B_p , and their symmetric \check{B}_p . Each square represents an image point, and the square with the red cross is the origin point p of the SE.

4.1.2 Dilation and erosion

4.1.2.1 Binary operators

A binary image is an image that only has two grey-levels: black pixels with value 0 and white pixels with value 1. A binary image can also be considered a subset of the image support. Let I_b be a binary image and X its support, if I_b is an entirely white image, then $I_b = X$ otherwise $I_b \subset X$.

The dilation of image I_b by the structuring element B_p , noted $\delta_{B_p}(I_b)$, may be defined as the union of all the B_p such that p belongs to I_b :

$$\delta_{B_p}(I_b) = \bigcup_{p \in I_b} \{B_p\} \quad (2)$$

The erosion, noted $\varepsilon_{B_p}(I_b)$ may be defined as the set of all the origin points p such that B_p is entirely included in I_b :

$$\varepsilon_{B_p}(I_b) = \{p, B_p \subseteq I_b\} \quad (3)$$

However, the erosion may also be seen as an intersection:

$$\varepsilon_{B_p}(I_b(x)) = \bigcap_{b \in B_p} X(x + b) \quad (4)$$

where $x + b$ is a vector addition. Figure 15 illustrates and compares the binary dilation and erosion.

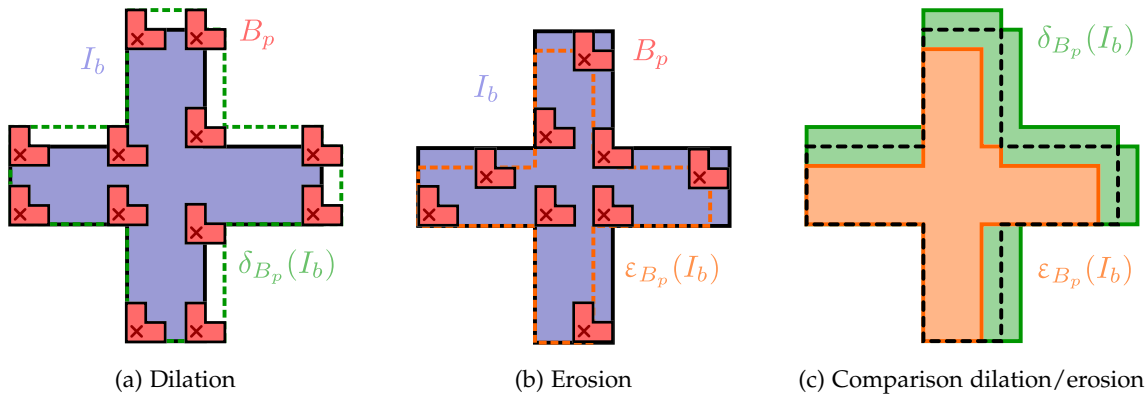


Figure 15: Illustration of (a) a dilation (green), and (b) an erosion (orange) with the same SE B_p (red). The comparison between both results is shown in (c) where the dotted black line is the initial image.

4.1.2.2 Grey-level operators

One way to easily extend binary operators to grey-scale is to consider the set of all the successive thresholds of a grey-scale image. Let I_t be the result of the thresholding of image I at level t :

$$I_t = \{x \in X, I(x) \geq t\} \quad (5)$$

Then, the grey-level dilation (resp. erosion) of I by B_p may be seen as the result of applying the same binary dilation (resp. erosion) at each threshold of I (see Eqs. (6) and (7)).

$$\delta_{B_p}(I(x)) = \max \{t \in G, x \in \delta_{B_p}(I_t)\} \quad (6)$$

$$\varepsilon_{B_p}(I(x)) = \max \{t \in G, x \in \varepsilon_{B_p}(I_t)\} \quad (7)$$

This way of defining the dilation and erosion is not efficient. In practice, the erosion and dilation are computed by taking a minimum or maximum over the structuring element, considered as a sliding window, as follows:

$$\delta_{B_p}(I(x)) = \max_{b \in B_p} I(x - b) \quad (8)$$

and

$$\varepsilon_{B_p}(I(x)) = \min_{b \in B_p} I(x + b) \quad (9)$$

Following these definitions, the erosion and dilation are computed by taking respectively a minimum or maximum over the structuring element, considered as a sliding window. Figure 16 shows an example of grey-level dilation and erosion.




Figure 16: An image (a) and its dilation (b) and erosion (c) by a square of size 4 pixels

4.1.2.3 Properties

- Duality

Dilation and erosion are dual operators by complementation. Let $A \subseteq X$, be a set of image points, we refer to A^c as the complement of A such that $A^c = X \setminus A$, then:

$$[\varepsilon_{B_p}(I_b)]^c = \delta_{\check{B}_p}(I_b^c) \quad (10)$$

$$[\delta_{B_p}(I_b)]^c = \varepsilon_{\check{B}_p}(I_b^c) \quad (11)$$

- Extensivity

The dilation is extensive whereas the erosion is anti-extensive:

$$I_b \subseteq \delta_{B_p}(I_b) \quad (12)$$

$$\varepsilon_{B_p}(I_b) \subseteq I_b \quad (13)$$

- Commutativity

The dilation is commutative whereas the erosion is not:

$$\delta_{B_p}(I_b) = \delta_{I_b}(B_p) \quad (14)$$

The dilation also commutes with the union while the erosion commutes with the intersection:

$$\delta_{B_p} \left(\bigcup I_b \right) = \bigcup \delta_{I_b}(B_p) \quad (15)$$

$$\varepsilon_{B_p} \left(\bigcap I_b \right) = \bigcap \varepsilon_{I_b}(B_p) \quad (16)$$

- Associativity

The dilation is associative whereas the erosion is not. Let $A, B, C \subseteq X$:

$$\delta_A(\delta_B(C)) = \delta_{\delta_A(B)}(C) \quad (17)$$

- Monotonicity

The dilation and erosion are both increasing. If $A \subseteq C$, then:

$$\delta_{B_p}(A) \subseteq \delta_{B_p}(C) \quad (18)$$

$$\varepsilon_{B_p}(A) \subseteq \varepsilon_{B_p}(C) \quad (19)$$

- Adjunct

The dilation and erosion, using the same structuring element, are adjunct operators:

$$X \subseteq \varepsilon_B(Y) \iff \delta_B(X) \subseteq Y \quad (20)$$

4.1.3 Opening and closing

4.1.3.1 Binary operators

Opening and closing are two dual operators resulting from the combination of erosion and dilation. The opening $\gamma_{B_p}(I_b)$ of a binary image I_b with the SE B_p is defined as the composition of an erosion followed by an adjunct dilation. However, it may also be seen as the union of the SE which are included in the binary image:

$$\gamma_{B_p}(I_b) = \delta_{B_p}(\varepsilon_{B_p}(I_b)) \quad (21)$$

$$= \cup \{B_p, B_p \subseteq I_b\} \quad (22)$$

The closing $\varphi_{B_p}(I_b)$ of a binary image I_b with the SE B_p is defined as the composition of a dilation followed by an adjunct erosion:

$$\varphi_{B_p}(I_b) = \delta_{B_p}(\varepsilon_{B_p}(I_b)) \quad (23)$$

Intuitively, for white objects on dark background, an opening tends to remove objects smaller than the chosen SE, whereas a closing tends to fill holes smaller than the chosen SE. Figure 17 shows a comparison of binary dilation, erosion, openings and closings.

4.1.3.2 Grey-level operators

Similarly to grey-level dilation and erosion, grey-level opening and closing are defined by applying the same binary opening or closing on each threshold of the grey-level image I (see Eqs. (24) and (25)). We note that in practice, more efficient algorithms were developed to compute these operators. An illustration of such operators is presented in Figure 18.

$$\gamma_{B_p}(I(x)) = \max \{t \in G, x \in \gamma_{B_p}(I_t)\} \quad (24)$$

$$\varphi_{B_p}(I(x)) = \max \{t \in G, x \in \varphi_{B_p}(I_t)\} \quad (25)$$

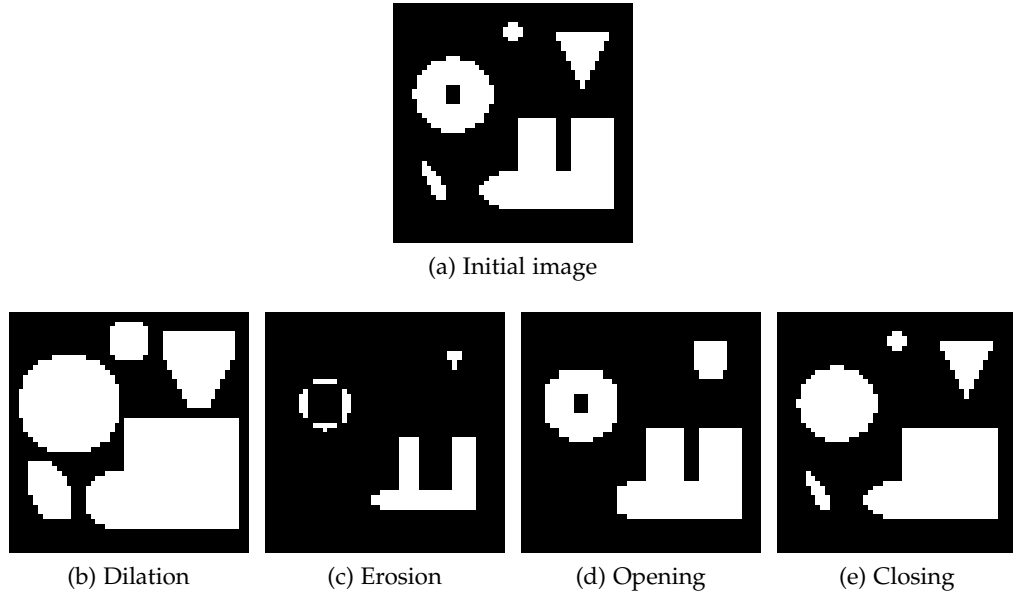


Figure 17: Illustration of basic binary mathematical morphology operators on a 50×50 pixel image by a square of size 4 pixels.

4.1.3.3 Properties

- Duality

Opening and Closing are dual operators:

$$[\varphi_{B_p}(I_b)]^c = \gamma_{\check{B}_p}(I_b^c) \quad (26)$$

$$[\gamma_{B_p}(I_b)]^c = \varphi_{\check{B}_p}(I_b^c) \quad (27)$$

- Extensivity

The closing is extensive whereas the opening is anti-extensive:

$$I_b \subseteq \varphi_{B_p}(I_b) \quad (28)$$

$$\gamma_{B_p}(I_b) \subseteq I_b \quad (29)$$

- Idempotence

The closing and opening are both idempotent:

$$\gamma_{B_p}(\gamma_{B_p}(I_b)) = \gamma_{B_p}(I_b) \quad (30)$$

$$\varphi_{B_p}(\varphi_{B_p}(I_b)) = \varphi_{B_p}(I_b) \quad (31)$$

- Monotonicity

The closing and opening are both increasing. If $A \subseteq C$, then:

$$\gamma_{B_p}(A) \subseteq \gamma_{B_p}(C) \quad (32)$$

$$\varphi_{B_p}(A) \subseteq \varphi_{B_p}(C) \quad (33)$$

- Composition

The maximum of several openings is an opening, and the minimum of several closings is a closing.



Figure 18: An image (a) and its opening (b) and closing (c) by a square of size 4 pixels

4.1.3.4 Algebraic opening and closing

Based on the previous properties, more general definitions of the opening and closing may be given. Any filter which is increasing, idempotent and anti-extensive is an algebraic opening. In the same way, any filter which is increasing, idempotent and extensive is an algebraic closing.

4.2 FILTERING CURVILINEAR STRUCTURES WITH LINE SEGMENTS AS STRUCTURING ELEMENTS

A common approach to filter curvilinear structures in mathematical morphology is to use openings or closings with line segments as SE. For the sake of clarity, we only present this approach with openings as we consider bright structures of interest on a dark background. However, it is easily transposable to dark structures on a bright background using closings instead of openings.

Filtering an image by an opening with a line segment preserves curvilinear structures, but only those that everywhere include the SE. The SE length is a parameter that is chosen according to the application, but the orientation issue remains. In order to filter curvilinear structures with a arbitrary orientations, the classical approach is to use a bank of line segments with the same length, but with varying orientations. Openings with each of these SEs are applied and the final result is the union of all the openings (see Figure 19). This approach is sometimes called *Radial Opening*.

Radial opening requires the computation of many erosions/dilations with line segments. Van Herk *et al.* [96] proposed in 1992 an efficient algorithm to perform 2D erosions and dilations with a line segment of arbitrary size, but only in the main natural orientations of a 2D image (horizontal, vertical and the two diagonals). In 1996, Soille *et al.* [83] extended this work by proposing an algorithm to compute erosions and dilations along discrete lines at arbitrary angles. Nevertheless, this algorithm results in erosions and dilations which are not translation invariant. Finally, a translation invariant algorithm, with the same complexity as the latter, was proposed in 2001 [82]. Recently, Van Droogenbroeck and Buckley [95]

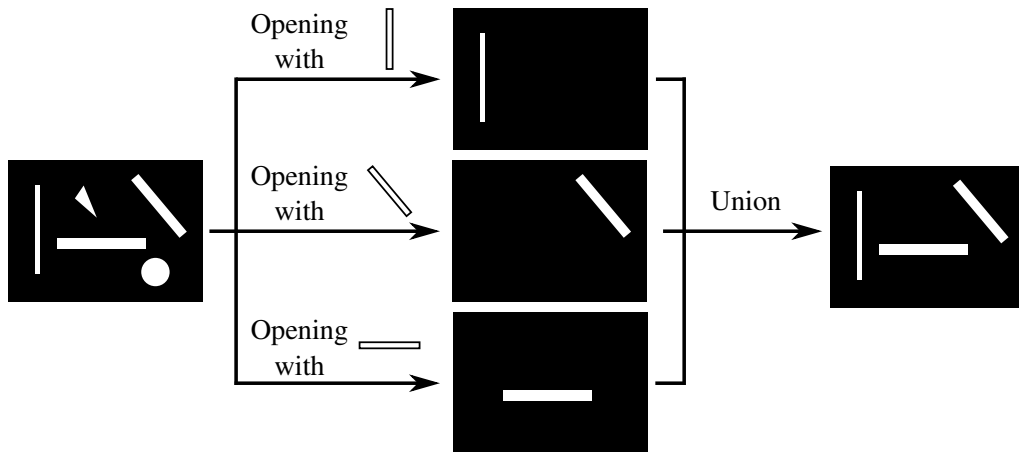


Figure 19: Illustration of a radial opening. Only the three useful SE are shown here.

proposed a novel algorithm to compute erosions but also openings with line segments. This is currently the fastest known algorithm.

Radial opening presents many advantages. First, as the opening is anti-extensive, this approach only decreases the signal of non-curvilinear structures. This means that it neither creates false positives nor moves edges of the sought structures. Moreover, this approach is non-local since it looks into a neighborhood of the size of the SE. This certainly avoids false detections, which may be caused by local analysis. Finally, the only parameters of a radial opening are the SE length and the number of different orientations.

However, this approach assumes that curvilinear structures are locally straight. This is a restrictive hypothesis. Indeed, to filter tortuous parts of curvilinear structures, the SE length should be reduced to a length small enough for the straight SE to fit the tortuous parts of the curvilinear structure. Nevertheless, all the non-curvilinear structures into which the line SE can fit are also preserved by openings. Consequently, the smaller the SE length, more both tortuous curvilinear structures and non-curvilinear structures are preserved (see Figure 20).

4.3 PATH OPERATORS

We have seen that openings (or closings) may be used with line segment SE to filter straight curvilinear structures. In the case of curved curvilinear structures this approach may still be used but tends to also preserve many isotropic structures since the SE length needs to be small.

To address this problem, Buckley and Talbot proposed flexible linear openings and closings [10] which were later better formalized and called path operators [39]. Path operators include two dual operators: path opening and path closing. As in the previous section, without loss of generality, we focus the following explanations on path openings.

Contrary to a classical opening which uses one SE with a fixed shape and length, a path opening uses a set of flexible segments, defined by their length and orientation, called *paths*. If at least one path fits a structure, this structure is preserved by the path opening. Consequently, path openings do not require a strong hypothesis on the shape of the structure of interest as classical openings, only a general orientation is required.

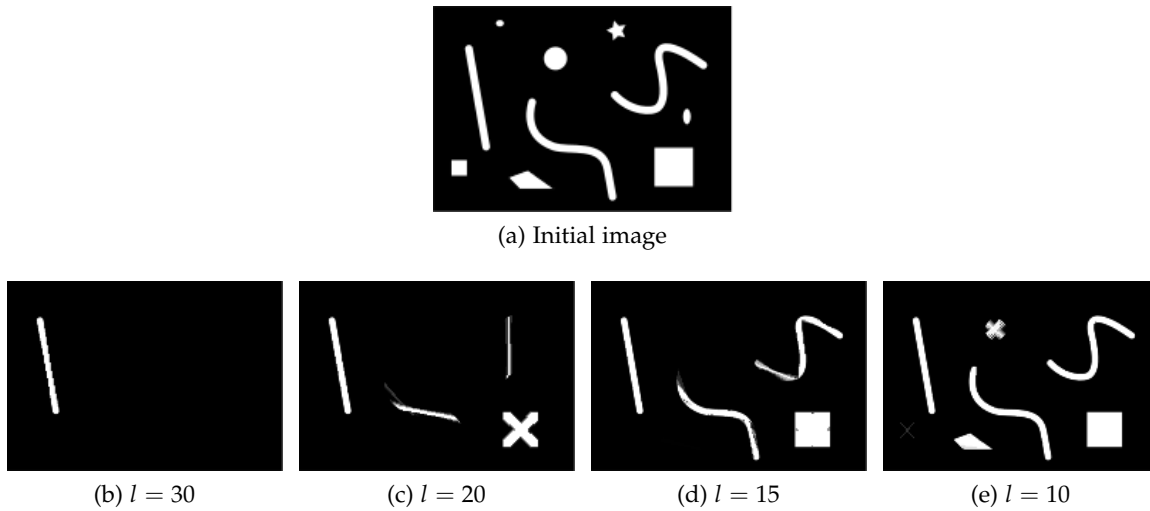


Figure 20: An image (a) and its radial openings for several SE lengths l (b-e). For large values l , only the straight structure is preserved as no line SE fit in the curved curvilinear structures. When l increases, curved structures are better preserved, but so are non-curved curvilinear structures.

In this section, we first define paths and the adjacency relation required to compute a path opening (Section 4.3.1). Then we properly define the path opening operator (Section 4.3.2). The sets of classically chosen orientations are described in Section 4.3.3. Finally we present the different algorithms proposed in the literature to compute path openings in Section 4.3.4.

4.3.1 Paths

A path is a flexible linear structure defined as a set of L connected points on a directed graph $G = (X, E)$. X is the set of pixels of an image and E the set of edges representing the connectivity between pixels. Let \rightarrow be an irreflexive, non-symmetric binary relation. For two points x and y of X , $x \rightarrow y$ means that there is an edge from x to y .

More formally, we define a path π of length L as a set $\sigma(\pi)$ of L successively connected pixels on a graph G :

$$\sigma(\pi) = \left\{ \{x_1, x_2, \dots, x_L\}, x_i \rightarrow x_{i+1} \right\} \quad \forall i \in [1, L-1], i \in \mathbb{N} \quad (34)$$

Practically, the set of edges, E , is generated from an elementary adjacency relation, Γ , which is periodically reproduced over X (see Figure 21). We note $\Pi_L^\Gamma(X)$ the set of all the paths of length L on X with adjacency relation Γ .

The adjacency relation, Γ , defines a global orientation for all the paths $\Pi_L^\Gamma(X)$. For example, Γ_1 (see Figure 21.(a)) only define globally vertical paths, whereas Γ_2 (see Figure 21.(c)) only defines globally diagonal paths. In the following, we call Γ the (global) orientation of the path opening.




Figure 21: Illustration of two 2D adjacency relations, resp. (a) and (c), and their associated graph, resp. (b) and (d) on the same image support E . E_i is the set of edges generated from Γ_i . Examples of paths of length 3 are represented in red and paths of length 4 in blue.

Adjacency relation and symmetry

We defined \rightarrow as a non-symmetric binary relation. Indeed, it is important that paths be defined with a non-symmetric relation, otherwise they could present loops or could not all have the same global orientation (see Figure 22.a). Nonetheless, the paths defined with this asymmetric relations are themselves symmetric.

For a given global orientation, two equivalent adjacency relations may be defined. For example, in Figure 21.(a), we chose to represent Γ_1 as 3 vectors from a pixel to its 3 upward neighbors, but exactly the same paths could be defined with the symmetric of Γ_1 (i.e. 3 vectors from a pixel to its 3 downward neighbors), see Figure 22.(b-c).

More formally, let $-\Gamma$ be the symmetric of Γ , then $\Pi_L^\Gamma = \Pi_L^{-\Gamma}$ and $\alpha_L^\Gamma = \alpha_L^{-\Gamma}$.

To represent the symmetry of the adjacency relation, we illustrate a path opening orientation as an hourglass shape (see Figures 25 and 26).

4.3.2 Binary and grey-level path opening

The binary path opening, α_L^Γ , of length L along orientation Γ , is defined as the union of all the paths of length L on the binary image I_b (see Eq. (35)). This operator preserves each point of I_b belonging to at least one path of $\Pi_L^\Gamma(I_b)$ and removes the others. An example is shown in Figure 23.

$$\alpha_L^\Gamma(I_b) = \bigcup \{ \sigma(\pi), \pi \in \Pi_L^\Gamma(I_b) \} \quad (35)$$

Path openings may be easily extended to grey-levels in the same manner as classical openings. Let I be a grey-level image, $I : X \rightarrow G$, where X is the set of points of I and G is its set of grey-levels. As previously defined, I_λ is the thresholding of I at grey-level λ , ($\lambda \in G$). The grey-level path opening, A_L^Γ with length L along orientation Γ is defined by:

$$A_L^\Gamma(I(x)) = \max_{\lambda \in G} \{ \lambda, x \in \alpha_L^\Gamma(I_\lambda) \} \quad (36)$$

An illustration of a grey-level path opening compared with a classical opening is shown in Figure 24. The main advantage of path operators is that the SE may deviate from the global




Figure 22: Example of 2 paths (red and blue) with a symmetric binary relation (a). Both paths do not have the same global orientation and one presents a loop. On graphs defined with a non-symmetric binary relation, the same paths can be defined on a graph generated from Γ_1 (b) or its symmetric $-\Gamma_1$ (c).

orientation Γ . Consequently, path openings are more flexible than the classical openings as they can handle locally tortuous curvilinear structures.

4.3.3 Space discretization

We have seen that a path opening is computed according to a global orientation Γ . In order to recover curvilinear structures in all orientations, several path openings must be applied with various global path orientations. Since the more path opening orientations are applied, the longer the computation time, one generally choose the sparsest but isotropic discretization, which preserves all curvilinear structures. For 2D images, a discretization with 4 path opening orientations is generally chosen [10] whereas for 3D images, 7 path opening orientations are usually enough [22]. These two space discretizations are shown in Figures 25 and 26. Each path opening orientation (both in 2D and 3D) consists of a cone centered on a main vector (represented in red on the Figures). This vector represents the global orientation of any path lying in its path opening orientation.

Let \mathcal{O} be the set of the chosen 2D or 3D orientations, the final result of a path opening is the union of all the path openings along the different orientations of \mathcal{O} . The union of these path openings with a path length L , $A_L(I)$ is then defined as follows:

$$A_L(I) = \bigcup \{A_L^o, o \in \mathcal{O}\} \quad (37)$$

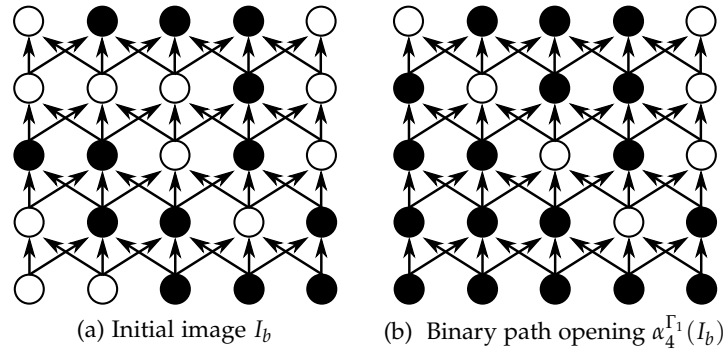


Figure 23: Example of a vertical path opening of length 4. Each disc represents a pixel of X . The white discs are the white pixels and the set of white discs is the binary image I_b . The arrows represent the connectivity of the image generated from Γ_1 . For example, the white pixel at row 2 and column 3 does not belong to any path of at least 4 pixels. Consequently, it is not part of the path opening result.




Figure 24: Comparison of a radial opening (b) and a path opening (c) on an image (a). The length of the SE is 10 pixels. Forty orientations are considered for the radial opening versus 4 for the path opening.

4.3.4 Path opening algorithms

Several algorithms have been proposed in the last few years to compute path openings. The first algorithm was proposed by Heijmans *et al.* in 2005 [39]. It is based on a recursive implementation. In 2007, Talbot *et al.* proposed an efficient path opening algorithm, based on an ordered implementation, which reduces the complexity and memory cost of the algorithm [86]. Both algorithm implementations are only available for 2D images. A dimensionality independent path opening algorithm was first proposed by Luengo Hendriks in 2010. More recently, a fast path opening algorithm was proposed by van de Gronde *et al.* [37] based on the Talbot *et al.* algorithm, but improved to make it dimensionality-independent.. A summary of the different algorithms and their properties is presented in Table 1.

Other algorithms have also been proposed to compute modified version of path openings. Luengo Hendriks [54] proposed a constrained path opening, where paths can no longer zig zag inside a given orientation. Indeed, the zig zag artificially increases the length of a path




Figure 25: The 4 path opening orientations in 2D. The arrows represent the adjacency relation. A path in a given orientation, going through the center point of the square, lies in the light blue shape. The red arrow is also the main vector (*i.e.* the global orientation of any path lying in this path opening orientation).

Algorithm	Complexity	Running time	Memory cost	Dimensionality
Heijmans <i>et al.</i> [39]	$o(NL)$	$\gg 4t$	$o(NL)$	2D
Talbot <i>et al.</i> [86]	$o(N \log(L))$	$4t$	$o(N)$	2D
Luengo Hendriks [54]	$o(N \log(L))$	$> 4t$	$o(N)$	nD
Van de Gronde <i>et al.</i> [37]	$o(N \log(L))$	t	$o(NL)$	nD

Table 1: Comparison of the different path opening algorithms. N is the number of pixels of the image and L the path length.

inside a structure. This improvement slightly increases the running time and the memory cost.

In 2013, Morard *et al.* proposed a parsimonious path opening [61]. Instead of computing paths in all the image, this approach cuts images in non-overlapping blocks, and pre-computes shortest paths from block borders to their opposite sides; then a 1D opening is applied only along these pre-computed paths. Pixels that are not involved in a path are set to the minimal grey-level value of the image. The final result is given by a morphological reconstruction under the initial image. Parsimonious path openings considerably reduce the computation time (the complexity is in $O(N)$). However, the result is similar but not the same as classical path openings.

4.4 ROBUST TO NOISE PATH OPENING

We have seen that the path opening is a powerful tool for curvilinear structure filtering, since it presents all advantages of a radial opening, but also allows flexible structuring elements more suitable for real curvilinear structures. Nonetheless, the sensitivity to noise remains. As seen in Chapter 1, curvilinear structures are highly sensitive to noise as a few noise pixels are often sufficient to disconnect such structures. As path openings assume connected SEs, a path opening with a path length equal to the whole curvilinear structure length does not preserve this structure. The only way to still preserve it is to reduce the path length in order

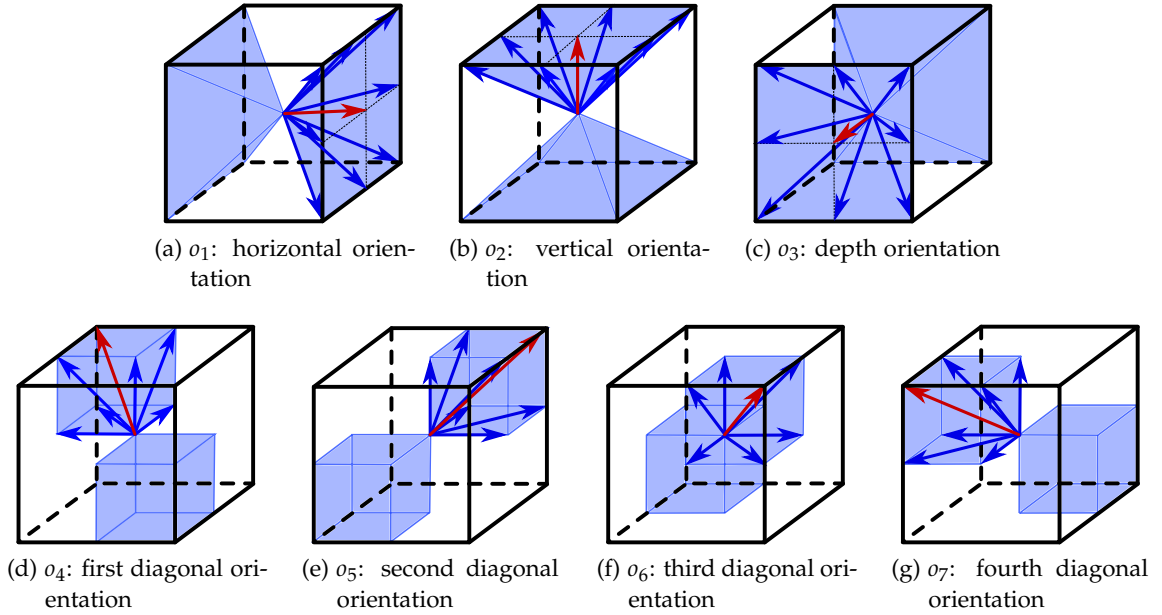


Figure 26: The 7 path opening orientations in 3D. The arrows represent the adjacency relation. A path in a given orientation, going through the center point of the cube, lies in the light blue shape. The red arrow is also the main vector (*i.e.* the global orientation of any path lying in this path opening orientation).

to fit each disconnected part of the curvilinear structure. However, doing so reduces the discriminative power of path opening between isotropic and curvilinear structures, since it preserves all the isotropic structures for which the longest size lies in between the reduced length and the true one (see Figure 27). This problem is particularly important as the longer the curvilinear structure, the higher the probability of disconnections.

To cope with this problem, several solutions were proposed to make path openings robust to noise.

4.4.1 Incomplete path opening (IPO)

Incomplete Path Opening (IPO) was first proposed by Heijmans *et al.* in 2005 [39]. They introduced a new parameter K which is the number of allowed disconnections inside a path of length L . Let I_b be a binary image on the image support X ($I_b \subset X$) and $\Pi_L(X)$ be the set of paths on X . The set of paths of an IPO with K accepted noise pixels on image I_b , $\Pi_{L,K}^\Gamma(I_b)$, is defined as follows:

$$\Pi_{L,K}^\Gamma(I_b) = \{a \in \Pi_L^\Gamma(X), |\sigma(a) \cap I_b^c| \leq K\} \quad (38)$$

where $|\cdot|$ is the cardinality of a set, *i.e.* the number of points, and I_b^c is the complementary set of I_b ($I_b^c = X \setminus I_b$). The IPO, $\alpha_{L,K}^\Gamma$ with path length L , K allowed noise pixels and along orientation Γ is then defined by:

$$\alpha_{L,K}^\Gamma(I_b) = \bigcup \{\sigma(\pi), \pi \in \Pi_{L,K}^\Gamma(I_b)\} \quad (39)$$




Figure 27: The noise problem of path opening. A binary image (a) and the same binary image where the curvilinear structure is disconnected by a noise pixel (c). Lengths of curvilinear structure parts are shown in red. The result of the path opening on (a) and (c) with different length (b,d,e). To preserve the noisy curvilinear structure, a path opening with length L_2 should be applied, but this opening also preserves an isotropic structure.

The IPO may be seen as a path opening which is more flexible as it allows more paths than a regular path opening. If $\Pi_L^\Gamma(I_b)$ is the set of allowed paths of a path opening, then $\Pi_L^\Gamma(I_b) \subseteq \Pi_{L,K}^\Gamma(I_b)$.

The algorithm proposed by Heijmans *et al.* has a complexity in $O(NLK)$ where N is the number of pixels, L the path length and K the number of allowed noise pixels per path. In 2007, Talbot *et al.* proposed an efficient algorithm to compute the IPO. Its complexity is $O(N \log(L))$, which is better; however, the Talbot *et al.* algorithm requires an extra memory cost of about K times the memory cost of the Heijmans *et al.* algorithm. Consequently, both algorithms are too expensive in 3D, and have never been implemented, to the best of our knowledge.

4.4.2 Robust path opening (RPO)

In 2013, Cokelear *et al.* proposed a new method to make the path opening robust to noise and tractable in 3D and called it Robust Path Opening (RPO) [21]. Their idea to reduce the complexity is to introduce a noise parameter, Q ($Q \in \mathbb{N}$), to handle the robustness as for the incomplete paths, but which is independent of the path length.

Instead of allowing K disconnections per path, RPO allows paths to present any disconnections of at most Q noise pixels. Consequently the number of noise pixels in a path of length L could vary between 0 and $Q \times \frac{L-1}{2}$. RPO is even more flexible than the IPO, as the

set of paths of a RPO is larger than the set of incomplete path opening. Let $\Pi_{L,Q}^\Gamma(X)$ be the set of paths of RPO on X , then $\Pi_L^\Gamma(X) \subseteq \Pi_{L,K}^\Gamma(X) \subseteq \Pi_{L,Q}^\Gamma(X)$. Figure 28 shows examples of paths for the path opening, the incomplete path opening and the RPO and Figure 29 shows a comparison of the three algorithms on a grey-level image.




Figure 28: Examples of possible paths of length 8 for the classical path opening (a), the IPO with $K = 2$ (b) and the RPO with $Q = 2$. The black squares represent the noise pixels.

The RPO algorithm of Cokelear *et al.* is based on the path opening algorithm of Luengo Hendriks [54]. The complexity of the RPO was experimentally assessed to be $O(QN \log(L))$. This is significantly better than the IPO but worse than the complexity of the efficient IPO. However, the memory cost of RPO is constant, unlike the efficient incomplete path opening. Moreover, in practice, since the noise parameter Q is more permissive in terms of noise proportion in a path, than the noise parameter K of the IPO, Q is generally set to 1 or 2 whereas K may be much larger. A comparison of both robustness to noise path opening is shown in Figure 29




Figure 29: Comparison of a path opening (b), an IPO with $K = 2$ (c) and a RPO with $Q = 2$ (d) with path length 40 on a synthetic image (a). The classical path opening does not preserve the disconnected curvilinear structures. The curvilinear structures are slightly brighter in the RPO result than in the IPO result as the set of paths is less restrictive, which allows the RPO to find a fitting path in a higher grey-level. However, the RPO also preserves more noise and isotropic structures than IPO.

RPO is the only robust to noise path opening version that is available for 3D images. However, the computational cost and memory cost are still fairly high. As our work lies

on path openings, we have proposed a simplified robust path opening which gives similar results in 3D as RPO, but requires significantly less computational and memory cost (see Chapter 6.2). In the following we shall always use our simplified robust to noise version of the path opening.

THE RORPO FRAMEWORK

So far, we have presented our strategy, which is to distinguish curvilinear structures from plane and blob-like structures, using their number of high responses to path opening. In this chapter, we present the framework we developed to achieve this goal (Section 5.1), as well as how we derive both intensity and directional features from it (Section 5.2 and 5.3). Finally, we present a multiscale approach of this framework (Section 5.4).

Contents

5.1	Ranking orientations	54
5.1.1	Pointwise rank filter (PRF)	54
5.1.2	Counting high path opening responses	55
5.2	Intensity feature	55
5.2.1	The generic RORPO intensity operator	55
5.2.2	Setting the i_t threshold	57
5.2.3	Limit orientations in 3D	61
5.3	Directional feature	64
5.3.1	Finding the orientations of interest	64
5.3.2	Combining the orientations	66
5.4	Multiscale approach	66
5.4.1	Length based multiscale approach	67
5.4.2	Methodology	68

5.1 RANKING ORIENTATIONS

In Section 3.2, we have observed that counting the number of high responses of path openings may allow us to distinguish between curvilinear structures, plane-like and blob-like structures. In order to count these high responses, we use rank filters. In the following, we first define the pointwise rank filter PRF , then we explain how we use it to count high path opening responses.

5.1.1 Pointwise rank filter (PRF)

The rank filter in image processing was first proposed by Heygster [40]. Applying a rank filter RF_k of order k on an image assigns, at each pixel x , the k^{th} smallest grey-level value amongst neighbors pixels of x . The neighborhood is defined by a window of a given size, centered on x . In particular, RF_1 returns the smallest value while RF_n returns the highest value in the window of size n pixels.

Let X be the support of an image and G its set of grey-levels. Instead of ranking grey-level values inside the same image, we define the Pointwise Rank Filter (PRF), $PRF_k : X^n \rightarrow X$, of order k , of a set of images as the filter that assigns, for each pixel x , the k^{th} highest grey-level value amongst the pixels at the same position of every image of the set.

More formally, let $A = \{A_1, A_2, \dots, A_n\}$, $A \in X^n$, be a set of n images of the same size, and $f_k : G^n \rightarrow G$, $k \in [1, n]$, be a function returning the k^{th} highest value amongst n values, then the Pointwise Rank Filter (PRF) of order k is defined as follows:

$$PRF_k(A)(x) = f_k(A_1(x), A_2(x), \dots, A_n(x)) \quad (40)$$

In particular, $PRF_1(A)$ is the pointwise maximum filter (see Eq. (41)), $PRF_n(A)$ is the pointwise minimum filter (see Eq. (42)) and $RF_{\frac{n+1}{2}}$ the pointwise median filter (if n is odd). Figure 30 illustrates the difference between the rank filter and the pointwise rank filter.

$$PRF_1(A)(x) = \max_{i \in [1, n]} A_i(x) \quad (41)$$

$$PRF_n(A)(x) = \min_{i \in [1, n]} A_i(x) \quad (42)$$

We now define the Pointwise Rank Filter (PRF), $PRF(A)$, of a set of n images A ($PRF : X^n \rightarrow X^n$), as the filter returning n ranked images, each resulting of applying one of the n PRF of order k :

$$PRF(A) = \{ PRF_i(A) \in X, \forall i \in [1, n] \} \quad (43)$$

At each point x , we have the relation: $PRF_1(A)(x) \geq PRF_2(A)(x) \geq \dots \geq PRF_n(A)(x)$.




Figure 30: Comparison between the rank filter of order 3 (a) and the pointwise rank filter of order 3 (b). The rank filter RF_3 centered on the grey-level 9 return the 3th smallest value inside the red window, which is 1. The pointwise rank filter PRF_3 returns the 3th largest value amongst the 4 pixels at the same position in each image, which is 2. Note that RF counts from the minimum and PRF counts from the maximum.

5.1.2 Counting high path opening responses

Let \mathcal{O} be a set of n path opening orientations (see Section 4.3.3), then $(A_L^o(I))_{o \in \mathcal{O}}$ is the set of the n path opening responses with path length L of the image I . Applying the PRF on $(A_L^o(I))_{o \in \mathcal{O}}$ provides n ranked path opening images $(\gamma_L^i(I))_{i \in [1, n]}$:

$$(\gamma_L^i(I))_{i \in [1, n]} = PRF\left((A_L^o(I))_{o \in \mathcal{O}}\right) \quad (44)$$

Interpretation in the binary case

Let us consider a binary image I_b , and its n responses to a path opening $(A_L^o(I_b))_{o \in \mathcal{O}}$. To simplify the explanation, we suppose that I_b contains several structures which are either preserved or removed entirely by a path opening. A structure which appears in p out of the n $A_L^o(I_b)$ responses, only appears in the first p ranked path opening images $(\gamma_L^i(I_b))_{i \in [1, p]}$. In this way, $\gamma_L^i(I_b)$ contains all structures preserved in at least i out of the n path opening orientations. An illustration in the 2D case with the 4 orientations of Figure 35 is shown in Figure 31.

5.2 INTENSITY FEATURE

5.2.1 The generic RORPO intensity operator

We have observed that curvilinear structures are preserved in fewer path opening orientations than other structures. We have also proposed, in the previous section, a method for counting the number of orientations detecting each structure. Indeed, γ_L^i contains all structures preserved in at least i path opening orientations.

Let i_{tube} be the maximal number of path opening orientations preserving a curvilinear structure ($i_{\text{tube}} \in [1, n - 1]$) and i_{plane} be the minimal number of path opening orientations




Figure 31: Illustration of the ranking procedure on a binary image I_b . First, the path openings are applied (A_L^o), then the PRF is computed to obtain the ranked path opening imaged $\gamma_L^i(I)$. The curvilinear structures appearing in 2 A_L^o are present in $\gamma_L^1(I)$ and $\gamma_L^2(I)$. The disc appears in all path opening responses because L is smaller than its diameter.

preserving a plane-like structure ($i_{\text{plane}} \in [1, n]$). Let i_t be a threshold value such as $i_{\text{plane}} \geq i_t > i_{\text{tube}}$, $i_t \in [1, n]$. We define the RORPO intensity feature, Φ_L as follows:

$$\Phi_L(I) = \gamma_L^1(I) - \gamma_L^{i_t}(I) \quad (45)$$

Indeed, $\gamma_L^{i_t}$ does not contain any curvilinear structures, as $i_t > i_{\text{tube}}$, but contains all plane-like structures, as $i_t \leq i_{\text{plane}}$. Moreover, γ_L^1 contains all structures detected by at least one path opening orientation. By subtracting $\gamma_L^{i_t}$ from γ_L^1 , RORPO only preserves structures detected in at most $i_t - 1$ orientations, namely, the curvilinear structures.

So far we have presented arguments based on a binary image, however, the same reasoning remains valid for gray-level images. Assuming bright structures on dark background, a grey-scale curvilinear structure presents a high value in γ_L^1 and a smaller one in $\gamma_L^{i_t}$, whereas a plane-like structure presents both high values in γ_L^1 and $\gamma_L^{i_t}$ (not necessarily equal, $\gamma_L^1 \geq \gamma_L^{i_t}$). Consequently, the RORPO intensity response is higher for the curvilinear structures than for the plane-like structures.

Beyond the preservation of curvilinear structures, the RORPO intensity feature also reduces the intensity of non-curvilinear structures and the mean background intensity.

In the following we present how we set the threshold value i_t in the 2D and 3D case.

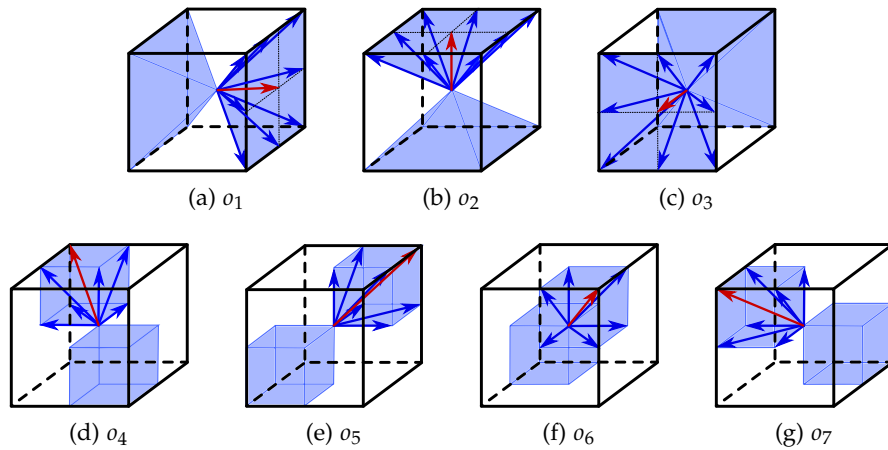


Figure 32: The 7 path opening orientations in 3D.

5.2.2 Setting the i_t threshold

5.2.2.1 The 3D case

In the 3D case, we use the 7 path opening orientations, $(o_i)_{i \in [1,7]}$ of Figure 32. To experimentally determine the i_t threshold value, we have uniformly sampled the 3D space in N different orientations, and generated for each of them, a synthetic straight binary tube of length 40 pixels and width 3 pixels. Using a uniform sampling strategy, we also have generated N' binary plane-like structures of length 40×40 pixels and of width 3 pixels. Examples of these samplings are shown in Figure 33.(a-b).




Figure 33: Illustration of the 3D space sampling with $N = 362$ tubes (a), and with $N' = 34$ plane-like structures (b). The 2D space sampling with $N = 12$ lines is presented in (c).

For each structure, we computed the 7 path openings along the $(o_i)_{i \in [1,7]}$, and we evaluated the number of path opening orientations which detects this structure. A structure may be partially detected by a path opening, so we need to set a threshold on the ratio of pixels required to call the structure "detected" in a given orientation. In these experiments, we

present the results for several thresholds: 99%, 80% and 50%. For example, if the detection threshold is 80%, a structure is said to be detected if at least 80% of its pixels are detected by the path opening in a given orientation.

The results of applying a path opening with path length $L = 40$ are shown in Table 2 for curvilinear structures and Table 3 for planes.

detection threshold	Percentage of tubes detected in i orientations						
	1	2	3	4	5	6	7
99%	0	82.92	17.08	0	0	0	0
80%	0	66.94	29.20	3.31	0.55	0	0
50%	0	54.27	41.32	3.58	0.83	0	0

Table 2: Percentage of tubes detected in i path opening orientations ($i \in [1,7]$) for $N = 363$ tubes, each in a different orientation, and considering several detection thresholds.

detection threshold	Percentage of planes detected in i orientations						
	1	2	3	4	5	6	7
99%	0	0	0	31.13	12.58	56.29	0
80%	0	0	0	0	0	81.45	18.55
50%	0	0	0	0	0	78.93	21.07

Table 3: Percentage of planes detected in i path opening orientations ($i \in [1,7]$) for $N = 323$ planes, each in a different orientation, and with several detection thresholds.

Results with the varying detection thresholds show the same trend: curvilinear structures are mostly detected in 2 or 3 path opening orientations (in 95% of the cases), whereas plane-like structures are detected in at least 4 path opening orientations. In other words, that means that $i_{\text{plane}} = 4$ and $i_{\text{tube}} = 3$. Consequently, we choose $i_t = 4$ which satisfies $i_{\text{plane}} \geq i_t > i_{\text{tube}}$. In 3D, the RORPO intensity feature is then defined by:

$$\Phi_L(I) = \gamma_L^1(I) - \gamma_L^4(I) \quad (46)$$

An illustration of the RORPO intensity feature is presented in Figure 34 on a synthetic image. This image was designed to contain the 3 types of structures in 3D: a plane-like structure, blob-like structures and a spiral (*i.e.* a curvilinear structure) with varying radius and width. We can see that the RORPO intensity feature successfully removes the plane-like structure and most of the blob-like structures, while preserving the spiral. Moreover, the mean background intensity is also decreased, resulting in a more contrasted spiral. More results, validations and comparisons are presented in Chapter 7.

With this definition, up to 5% of curvilinear structures may be lost by RORPO; *i.e.* the curvilinear structures detected in 4 and 5 path opening orientations. In order to preserve these, we propose a post processing procedure, which is described in Section 5.2.3.

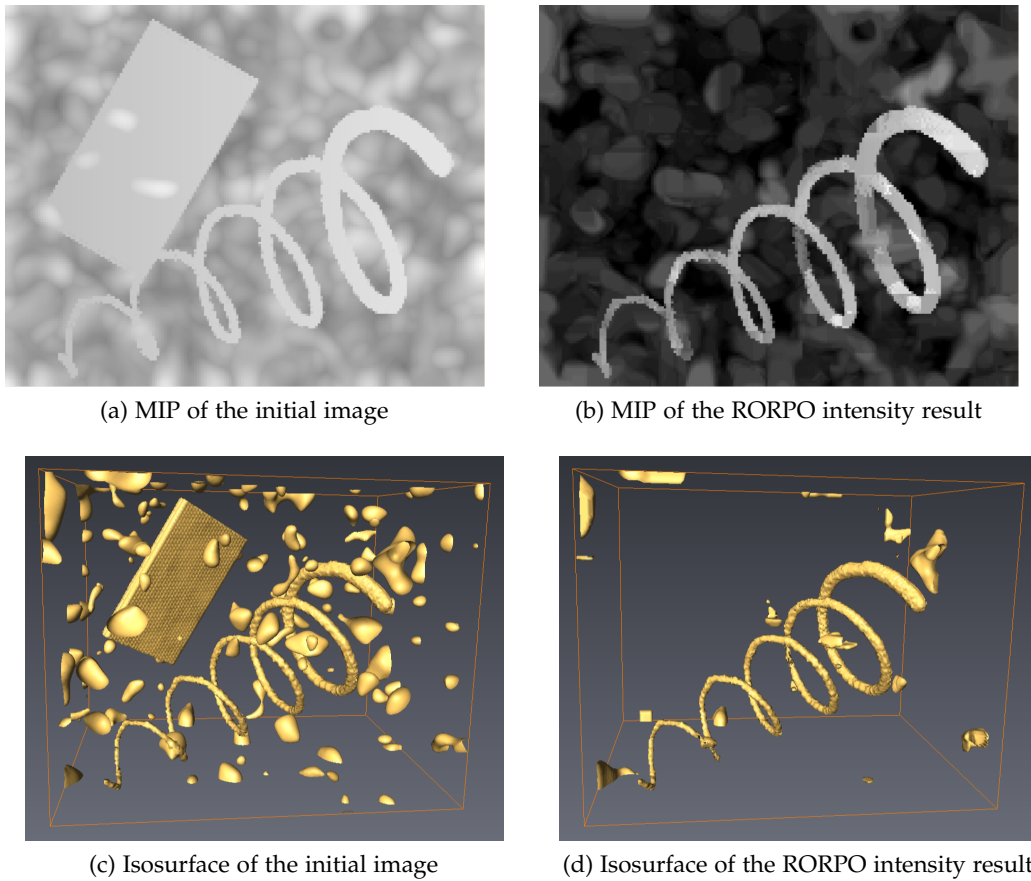


Figure 34: Illustration of the RORPO intensity feature on a 3D synthetic image containing a plane-like structure, blob-like structures and a spiral with varying radius and width. The initial image and the RORPO intensity result are both shown with a projection view (MIP) (a–b) and an isosurface view (c–d).




Figure 35: The 4 path opening orientations in 2D.

5.2.2.2 The 2D case

In the 2D case, there are no plane-like structure. The only condition on the threshold value i_t is: $i_t > i_{tube}$. To determine this threshold in 2D, we performed the same experiment as in the 3D case, but solely on tubes. We used the 4 path opening orientations, $(o_i)_{i \in [1,4]}$ of Figure 35 and also sampled the 2D space to generate synthetic tubes in different orientations (see Figure 33.(c)). The results of the experiment are shown in Table 4.

detection threshold	Percentage of tubes detected in i orientations			
	1	2	3	4
99%	0	98.73	1.27	0
80%	0	98.73	1.27	0
50%	0	93.67	6.33	0

Table 4: Percentage of tubes detected in i path opening orientations ($i \in [1, 4]$) for $N = 157$ tubes, each in a different orientation, and with several detection thresholds.

We can see that curvilinear structures are detected in 2 or 3 path opening orientations, irrespective of the detection threshold. Based on these results, $i_{tube} = 3$. Consequently, we set $i_t = 4$ which satisfies $i_t > i_{tube}$. In 2D, the RORPO intensity feature is then also defined by:

$$\Phi_L(I) = \gamma_L^1(I) - \gamma_L^4(I) \quad (47)$$

The reader should note that the i_t threshold value is the same in 2D and 3D. This is serendipitous as its value only depends on the choice of the set of path opening orientations. With the choice we made, $i_t = 4$ both in 2D and 3D, but with other path opening orientation sets, the value of i_t may be different in 2D and 3D.

An illustration of the RORPO intensity feature in 2D is presented in Figure 36 on a synthetic image.




Figure 36: Illustration of the RORPO intensity feature on a 2D synthetic image.

5.2.3 Limit orientations in 3D

5.2.3.1 Origin

In the previous section we saw that a few curvilinear structures seemed to be detected in 4 or 5 path opening orientations in the 3D case. These curvilinear structures lie in boundary orientations, resulting from the overlap of the path opening orientations we chose.

Path opening orientations need to overlap in order to preserve tortuous structures lying in several orientations. At each point of the image, only two path opening orientations should overlap, which is the case for most of the image points. However, as the path opening orientations are discrete closed cones, two adjacent orientations, which should be distinct, still overlap at their boundaries. The limit orientation problem appears when these boundaries are shared by more than 3 cones. Indeed, curvilinear structures lying in these boundaries yield a high response in more than 3 path opening orientations, and thus, are not preserved by RORPO. In this section we propose a post-processing approach to preserve these curvilinear structures .

5.2.3.2 Post-processing

We call *4-orientation* (resp. *5-orientation*) *tubes*, curvilinear structures that are detected in 4 (resp. 5) path opening orientations. Examples of such tubes are shown in Figure 37. In the following, we call *main orientations* the first three path opening orientations o_1 , o_2 and o_3 corresponding to the horizontal, vertical and depth orientation, and we call *diagonal orientations* the remaining four path opening orientations o_4 , o_5 , o_6 and o_7 . The post-processing procedure is composed of 3 steps:

- Detection of the 4 or 5-orientation tubes;
- Removal of the remaining plane-like structures;
- Supremum with the RORPO intensity feature result.

In the first step, we isolate the 4 and 5-orientation tubes from the other structures.

The 5-orientation tubes are detected in one main orientation and the four diagonal orientations. Consequently, the intersection between the path opening results with the four diagonal orientations preserve these tubes. We call this intersection $\Gamma_L^5(I)$:

$$\Gamma_L^5(I) = \min_{i \in \{4,7\}} A_L^{o_i}(I) \quad (48)$$

The 4-orientation tubes are detected either in two main and two diagonal orientations (pattern 1) or in three main and one diagonal orientations (pattern 2). Consequently, the union of the combinations, \mathcal{C} , of intersections between four of these orientations, detects the

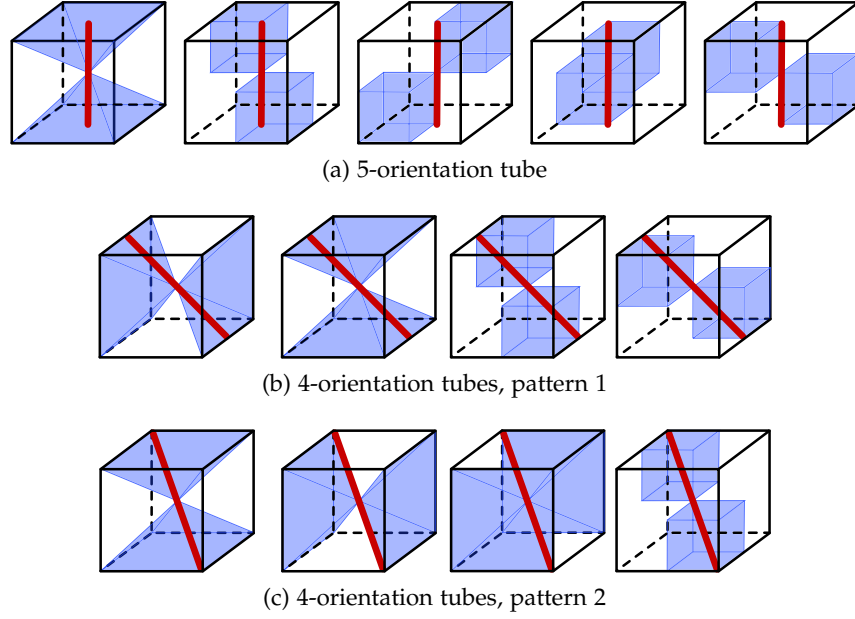


Figure 37: Example of 4 and 5-orientation tubes. The red tube lies in the boundary between 5 (a) or 4 (b) and (c) path opening orientations.

4-orientation tubes. There are ten different combinations of interest: six for pattern 1 and four for pattern 2 (see Eq. (49)). We call this union $\Gamma_L^4(I)$ (see Eq. (50)).

$$\begin{aligned}
 \mathcal{C}_1 &= \{o_1, o_2, o_4, o_7\} \\
 \mathcal{C}_2 &= \{o_1, o_2, o_5, o_6\} \\
 \mathcal{C}_3 &= \{o_1, o_3, o_5, o_7\} \\
 \mathcal{C}_4 &= \{o_1, o_3, o_4, o_6\} \\
 \mathcal{C}_5 &= \{o_2, o_3, o_6, o_7\} \\
 \mathcal{C}_6 &= \{o_2, o_3, o_4, o_5\} \\
 \mathcal{C}_7 &= \{o_1, o_2, o_3, o_4\} \\
 \mathcal{C}_8 &= \{o_1, o_2, o_3, o_5\} \\
 \mathcal{C}_9 &= \{o_1, o_2, o_3, o_6\} \\
 \mathcal{C}_{10} &= \{o_1, o_2, o_3, o_7\}
 \end{aligned} \tag{49}$$

$$\Gamma_L^4(I) = \max_{i \in [1,10]} \min_{c \in \mathcal{C}_i} A_L^c(I) \tag{50}$$

$\Gamma_L^5(I)$ and $\Gamma_L^4(I)$ preserve both 5 and 4-orientation tubes, but they may also preserve plane-like structures that lie in the same limit orientations. The second step of this post-processing removes these remaining plane-like structures. To this end, we need to isolate these plane-like structures. We define Δ_L^i as the first ranked path opening orientation that does not contain i -orientation tubes. In particular, $\Delta_L^4(I) = \gamma_L^5$ and $\Delta_L^5(I) = \gamma_L^6(I)$. Then $\min\{\Gamma_L^i(I), \Delta_L^i(I)\}$ contains only these remaining plane-like structures. However, it may happen that $\Delta_L^i(I)$ only contains subsets of plane-like structures. To recover the totality of these plane-like structures, we also add a geodesic dilation step ρ of $\Delta_L^i(I)$ in γ_L^i ; since γ_L^i is the first ranked path opening orientation containing all plane-like structures. Finally the removal of the remaining plane-like structures is simply performed as follows:

$$RP_L^i(I) = \Gamma_i(I) - \min\{\Gamma_L^i(I), \rho(\Delta_L^i(I), \gamma_L^i(I))\} \tag{51}$$

The final step of this post-processing consists of adding the i -orientation tubes to the RORPO intensity feature result (see Eq. (52)). Even though this post-processing procedure

includes several operators, all of them are easily computed, and the additional computational cost of this post-processing is negligible compared to RORPO. An illustration of the complete post processing procedure in the case of the 5-orientation tubes is shown in Figure 38.

$$\Phi_L^{post}(I) = \max \{ \Phi_L(I), RP_L^4(I), RP_L^5(I) \} \quad (52)$$

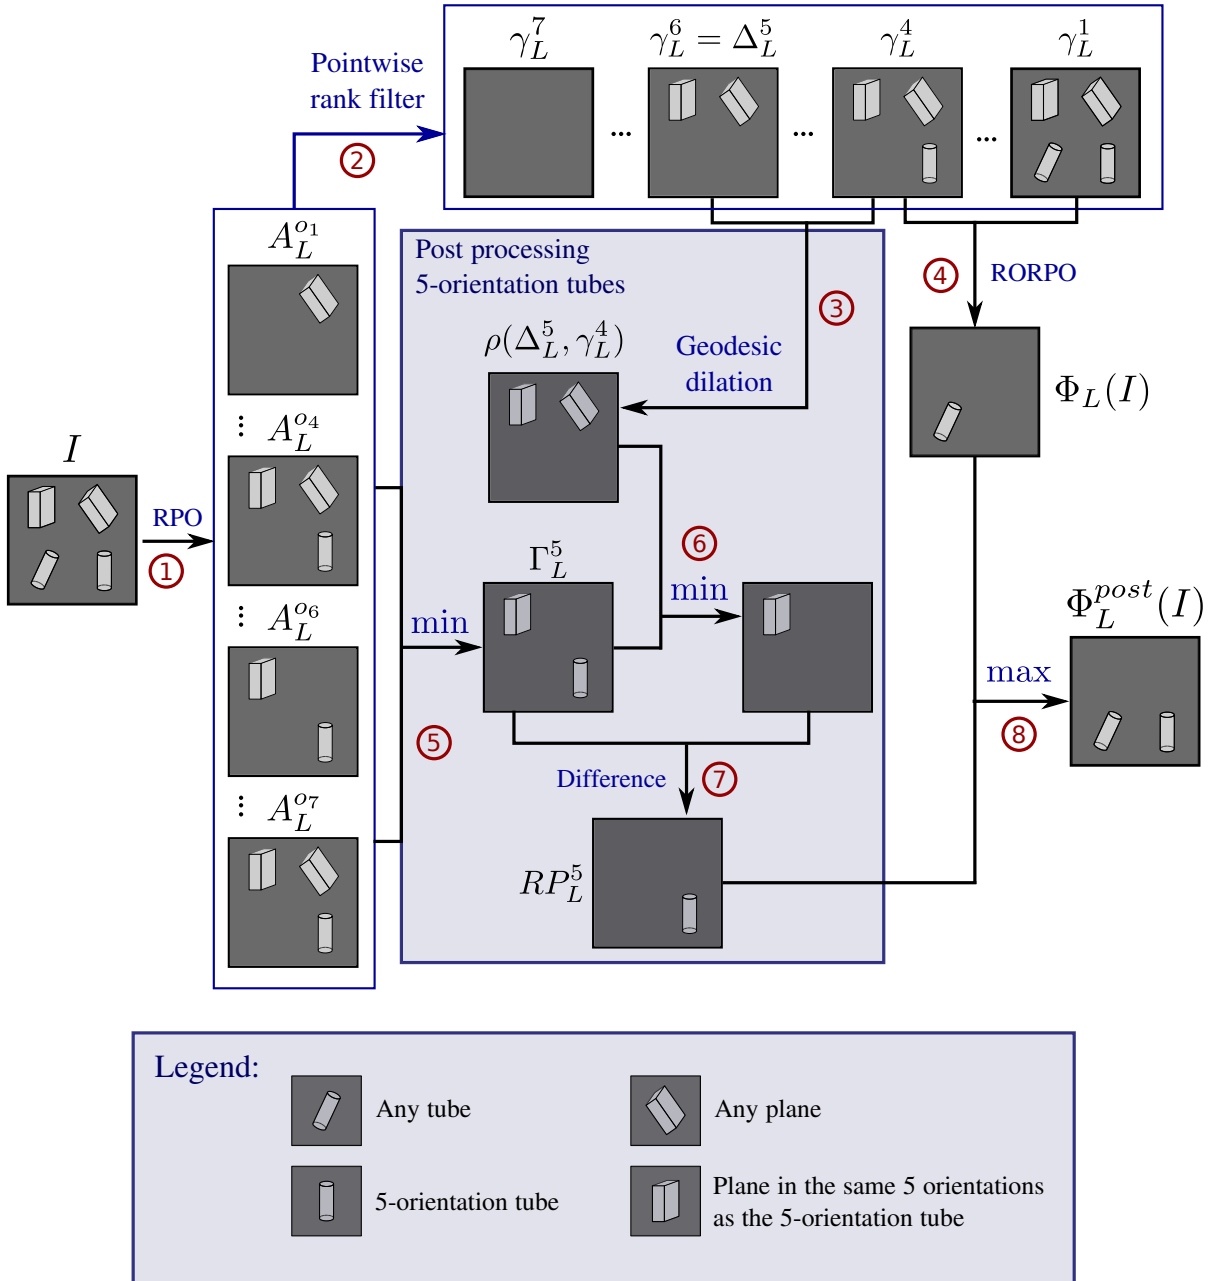


Figure 38: Post-processing pipeline for the 5-orientation tubes.

5.3 DIRECTIONAL FEATURE

Since RORPO is based on oriented path operators, it intrinsically carries information on the direction of the detected curvilinear structures. In order to retrieve this directional information, we have to identify which path opening orientation locally best fits the curvilinear structure. Typically, there is more than one path opening orientation detecting a curvilinear structure (generally 2 or 3, as seen in Section 5.2.2). In the following, we will refer to these orientations as *orientation(s) of interest*. By combining the orientation(s) of interest, we may obtain a reasonably accurate evaluation of the curvilinear structure orientation (see Figure 39).

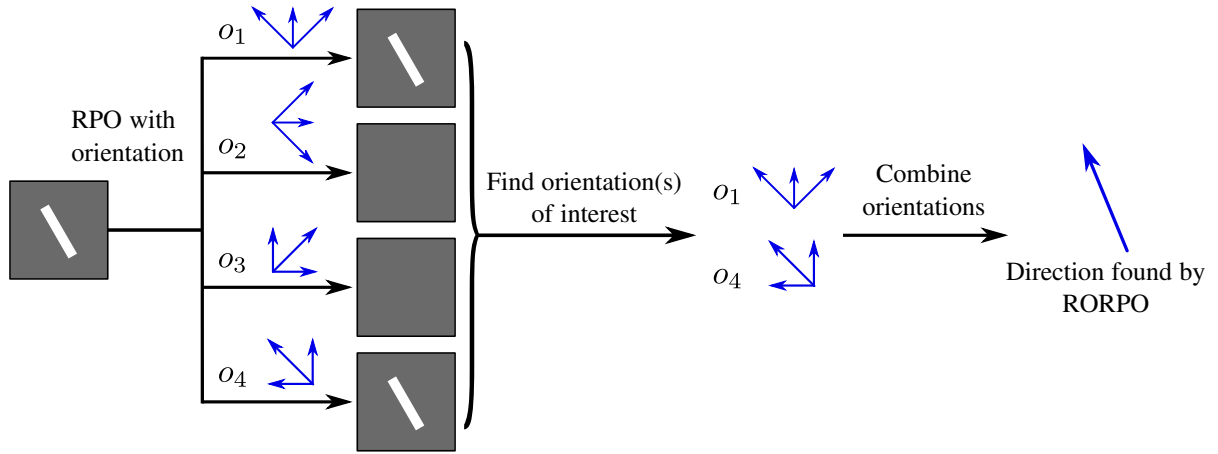


Figure 39: Principle of the directional feature extraction.

5.3.1 Finding the orientations of interest

Finding the orientations of interest requires to take a binary decision for each path opening orientation: either the path opening orientation preserves a curvilinear structure or it does not. However, the output of each path opening is not binary, but a grey-level response. A non-trivial classifying rule is then required to find which path opening orientation is an orientation of interest.

An orientation of interest, o , is characterized by a high path opening value $A_L^o(I)$ compared to the other path opening orientations, which do not detect the curvilinear structure. A first solution to detect these orientations of interest would be to set a threshold t ($t \in G$), such that if $A_L^o(I) > t$, the orientation o detects the curvilinear structure. Nonetheless, this solution is not robust. Indeed, the intensity of the path opening response at each pixel depends on its initial grey-level, and on those of its neighbors, which means that such a unique threshold does not exist.

Consequently, we propose, for each pixel, to classify the n path opening orientations in two classes: the orientations of interest and the remainder, which do not detect a curvilinear structure, based on an homogeneity criterion. Indeed, we expect that all orientations of interest have similar high responses, $A_L^o(I)$, while all the other path opening orientations have similar low responses. For simplicity and efficiency, we chose the standard deviation as the homogeneity criterion and we reformulate the problem as a minimization of the intra-class standard deviation. This is a similar approach to the Otsu thresholding method [67].

More formally, we recall that \mathcal{O} is the set of the n path opening orientations. Let \mathcal{C} be the set of all the combinations of 1, 2 and 3 path opening orientations. One of the combinations of \mathcal{C} is the set of the orientations of interest we seek. Let $P = \{p_1, p_2, \dots, p_k\}$, $1 \leq k \leq 3$, be one of these combinations ($P \in \mathcal{C}$); we note $\sigma_P(x)$ the standard deviation of the path opening responses with the orientations of P at pixel x such as:

$$\sigma_P = \sqrt{\frac{1}{k} \sum_{j=1}^k (A_L^{p_j} - \mu)^2}, \quad \mu = \frac{1}{k} \sum_{j=1}^k A_L^{p_j} \quad (53)$$

Then, we find the orientations of interest by solving the following problem:

$$\underset{P \in \mathcal{C}}{\text{minimize}} \sigma_P + \sigma_{\mathcal{O} \setminus P} \quad (54)$$

To solve this problem, we could compute the intra-class standard deviation of each combination of \mathcal{C} (14 combinations in 2D and 63 in 3D), or we can use the information provided by the Pointwise Rank Filter previously computed for the RORPO intensity feature, to reduce the possible number of combinations.

We remember that $(\gamma_L^i)_{i \in [1, n]}$ are the n ranked path opening orientations previously computed. We call O_i the orientation along which the ranked path opening γ_L^i was computed ($\gamma_L^i = A_L^{O_i}$). We have seen that curvilinear structures are detected by at most 3 path opening orientations (both in 2D and 3D), which means that only γ_L^1 , γ_L^2 or γ_L^3 may contain curvilinear structures. In this way, only O_1 , O_2 or O_3 may be orientations of interest and if O_i is an orientation of interest, so is O_{i-1} , if $i > 1$ (due to the ranking).

Consequently, by using the ranked path opening orientations from RORPO, we can reduce the number of useful combinations of \mathcal{C} to three: $\{O_1\}$, $\{O_1, O_2\}$, $\{O_1, O_2, O_3\}$. An illustration of the choice of orientation of interest is shown in Figure 40.

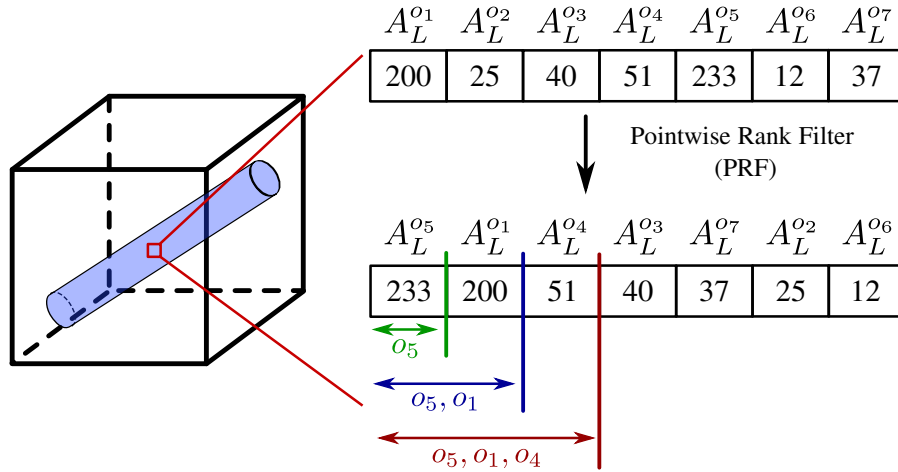


Figure 40: At each pixel of an image (red square), 7 path opening responses $A_L^{o_i}$ are computed. A high value $A_L^{o_i}$ means that orientation o_i is likely to be an orientation of interest. After applying the PRF, only 3 possible sets of path opening orientations may be the orientations of interest (here $\{o_5\}$, $\{o_5, o_1\}$ or $\{o_5, o_1, o_4\}$) depending on which one minimizes the intra-class standard deviation. In this case the orientations of interest are o_5 and o_1 .

5.3.2 Combining the orientations

Once the orientations of interest are found, we have to combine them in order to compute the final direction of the curvilinear structure. Each path opening orientation (both in 2D and 3D) are cones, generated from a main vector (represented in red in Figure 35 and 32). These main vectors constitute the global orientation of any path lying in their path opening orientation. Consequently, the combination of orientations of interest is performed by combining the main vectors of each orientation of interest.

The combination of the main vectors can not be performed by a simple vector sum. We encode the orientation of a curvilinear structure by a vector. However, a vector has a *direction*, which is more specific than an *orientation*. Indeed, each orientation may be encoded by 2 different vectors. For example, if a 3D curvilinear structure has an horizontal orientation, both vectors (*i.e.* direction) $[0,0,1]$ and $[0,0,-1]$ encode its orientation. Nonetheless, a vector sum with one vector or the other would not be the same. An illustration is shown in Figure 41.




Figure 41: Without correction of the principal vectors, the final curvilinear structure direction is incorrect ($[1,0]$). We first need to switch one of the principal vectors from $[1,-1]$ to $[-1,1]$ (both have the same orientation), and sum them to obtain the correct curvilinear structure direction ($[-1,2]$).

To deal with this issue, when performing the combination of the main vectors of the orientations of interest, we first need to ensure to choose the correct main vector for each orientation of interest. Then we can sum these corrected main vectors to obtain the final direction of the curvilinear structure. In practice, we chose the main vectors for which the sum of their pairwise angle is minimal. An example of the RORPO directional feature is presented in Figure 42.

5.4 MULTISCALE APPROACH

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 gen-




Figure 42: A 2D synthetic image (a) and its directional feature (b) (only represented inside the curvilinear structures).

erally highly correlated. To tackle this issue, various multiscale approaches have been developed. A common multiscale approach consists of applying the same filter several times, while changing a scale-related parameter and merging the results. In this section, we propose a multiscale version of RORPO, for both directional and intensity features, based on the length of the structures.

5.4.1 *Length based multiscale approach*

In the literature on multiscale curvilinear structure analysis, the scale parameter is usually related to the diameter of the curvilinear structure [33]. Indeed, whether considering blood vessels or roads, a small curvilinear structure is generally a curvilinear structure with a small diameter as opposed with a large curvilinear structure endowed with a large diameter.

However, in the case of RORPO, the length of the underlying paths, is the only tunable parameter. Had we wanted to vary the diameter, we would have had to combine RORPO with another filter. As it turns out, in many applications, the diameter, curvature and length of curvilinear structures are highly correlated, for physical reasons. For example, small blood vessels are generally more tortuous and shorter than large vessels like the aorta. The same argument can be made for insulation glass fibers or country roads vs. highways.

In the RORPO framework, the path length should not only be set depending on the length of the sought curvilinear structures, but also on its curvature. Indeed, to be preserved, a structure of a given length must lie entirely within a single path opening orientation. This structure may present small-scale curvatures and still be preserved, as the path opening orientations are cones. However, if the structure presents a curvature larger than the cone aperture (here called large-scale curvature), then the structure will not be preserved. In order to still preserve it, more than one path openings with different path opening orientations are required, both with a path length smaller than the real curvilinear structure length (see Figure 43).

To summarize, let us consider two curvilinear structures with the same length L , but one presents a large-scale curvature. The maximal path length preserving the latter is smaller than the maximal path length preserving the straight curvilinear structure. Moreover, we assume that in real applications, the diameter is highly correlated to the curvature of curvilinear structures: the smaller the diameter, the higher the curvature and so the smaller the allowable path length.

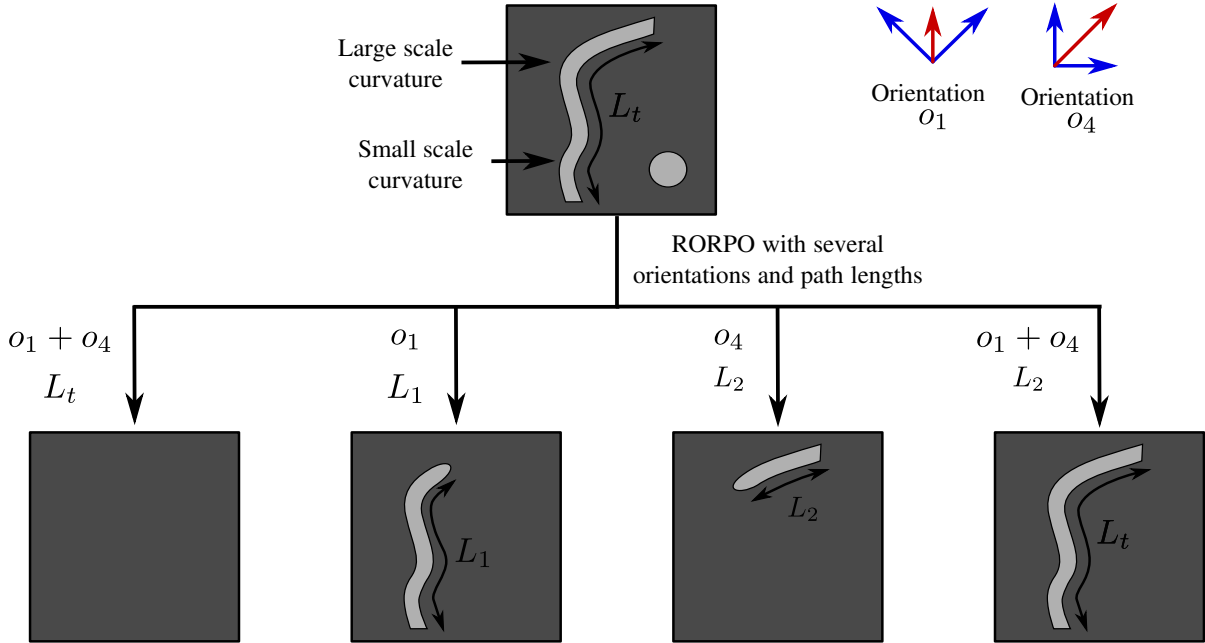


Figure 43: A curvilinear structure with a large-scale curvature cannot be preserved by only one path opening. Path openings with several orientations and a path length smaller than the real curvilinear structure length are required.

5.4.2 Methodology

With our assumption, using the path length as scale parameter is similar to using the diameter for the multiscale filtering of curvilinear structures in real applications. In order to merge the RORPO intensity results with varying path lengths, we propose to use the same approach as for merging the different path opening orientations results: the maximum operator. More formally, let $S = \{L_1, L_2, \dots, L_n\}$ be a set of path length. The multiscale RORPO intensity filter $\Phi(I)$ is defined as follows:

$$\Phi(I) = \bigvee_{L \in S} \Phi_L(I) \quad (55)$$

This multiscale paradigm is also applied to the RORPO directional feature. For each RORPO scale (*i.e.* path length), a direction can be computed; the final direction is the one associated with the highest RORPO intensity response and with the lowest scale.

In the next chapter we present a few algorithmic considerations regarding RORPO.

ALGORITHMIC CONSIDERATIONS

In this chapter, we present a few algorithmic considerations. The first section provides a summary of the RORPO algorithm. Section 6.2 presents our simpler, faster version of the path openings with noise robustness that we developed. Section 6.3 presents the different parameters of the RORPO framework and discusses our choices for the non-tunable parameters. Finally, we discuss in Section 6.4 the computational cost of our algorithm.

Contents

6.1	RORPO algorithm	70
6.2	Simplified robust path opening	70
6.3	Parameters	72
6.3.1	Path length	73
6.3.2	Noise robustness parameter	74
6.3.3	Orientations	76
6.4	Computational cost	77

6.1 RORPO ALGORITHM

In this section we provide a summarize algorithm for the computation of both RORPO features. The algorithm is presented in Algorithms 1 and 2. We recall that we use the algorithm of Luengo Hendriks [54] to compute the path openings (which is called *compute_PO* in Algorithm 1).

We released our C++ code to compute the RORPO intensity feature on github <https://github.com/path-openings/RORPO>, and we will soon release a new version that also compute the RORPO directional feature, on the same repository.

Algorithm 1 : RORPO algorithm without limit orientation correction

Data : Image I , set of scales $S = \{L_1, L_2, \dots, L_n\}$,
set of the 7 orientations $O = \{o_1, o_2, \dots, o_7\}$
Result : Intensity feature Φ and directional feature d

```

begin
  for each scale  $L$  in  $S$  do
     $I_d = \delta(I)$  Dilation of I
    begin
      for  $o$  in  $O$  do
         $A_L^o = \text{compute\_PO}(I_d, o, L)$  Compute the 7 RPO
      end
    end
     $(\gamma_L^i)_{i \in [1,7]}, (O_L^i)_{i \in [1,7]} = \text{ranking}((A_L^o)_{o \in [1,7]})$  Pointwise rank filter
     $\Phi_L = \min(I, \gamma_L^1 - \gamma_L^4)$  Intensity feature at scale L
     $\Phi = \max(\Phi_L, \Phi)$  Multiscale intensity feature
     $d_L = \text{compute\_directions}((\gamma_L^i)_{i \in [1,7]}, (O_L^i)_{i \in [1,7]})$  Directional feature at scale L
    if  $\Phi_L(x) > \Phi(x)$  then
       $d(x) = d_L(x)$  Multiscale directional feature
    end
  end
end

```

6.2 SIMPLIFIED ROBUST PATH OPENING

RORPO is based on path openings which can be sensitive to noise. We have seen in Chapter 4 that two solutions have been proposed to improve the noise robustness of path openings: the Incomplete Path Opening (IPO) and the Robust Path Opening (RPO). Only the latter has been implemented in 3D. In addition, the algorithmic layer handling the robustness is complex, and time and memory consuming. To cope with this problem, we proposed a simplified robust path opening, which yields similar results with only a small extra computation and memory cost.

Algorithm 2 : compute_directions**Data** : $(\gamma_L^i)_{i \in [1,7]}$, and $(O_L^i)_{i \in [1,7]}$ **Result** : Directional feature d **begin** **for each pixel** x **of** I **do**

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

Compute the

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

standard deviation

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

of the 3 useful classes

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

Keep the class with the lowest std

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

Compute the final direction **end****end**

Our method relies on a mask-based second-generation connectivity strategy [68] in order to reconnect the curvilinear structure parts disconnected by noise. A dilation by a cubical structuring element of size $N \times N \times N$ is first performed on the initial image I . This dilated image is later used to compute the regular path openings. This is equivalent to the RPO strategy which uses paths where some pixels can be darker due to noise. Roughly speaking, instead of allowing for noisy pixels inside paths, we compute paths on the dilation, which has already turns the noisy pixels into structure pixels (see Figure 44).




Figure 44: Illustration of the simplified robust path opening. A curvilinear structure is disconnected by a noisy pixel. After dilation this pixel is no longer considered as noise, but belongs to the curvilinear structure, which allows the path opening to preserve the whole curvilinear structure. A minimum operator is finally applied to ensure anti-extensivity.

To preserve the anti-extensivity of the path opening, a minimum operator is applied between the path opening result and the initial image. More formally, let δ_{C_N} be a dilation by a cube of size N , the simplified robust path opening, $A_{L,N}^o$, with path length L and orientation o is then defined as follows:

$$A_{L,N}^o(I) = \min\{I, \delta_{C_N}(A_L^o(I))\} \quad (56)$$

This simplified robust path opening only requires to compute an extra dilation by a cube, which may be computed with a constant complexity. It requires no additional complex algorithmic layer, and yields similar results to RPO. In term of paths, our method is equivalent to allowing noisy pixels inside paths, however, the noisy pixels may be located at the path extremities which is not the case for RPO. For large values of N , using the simplified robust path opening may preserve more noise and false positive structures than RPO. Nonetheless, in real applications, as for RPO, we rarely use values larger than $N = 5$ (equivalent to $Q = \frac{N-1}{2} = 2$ for RPO). A comparison between these two versions of robust path openings is presented in Figure 45.




Figure 45: Comparison of a RPO with $Q = 5$ (b) and the simplified RPO with $N = 5$ (c) both with path length $L = 50$ on a synthetic image (a).

6.3 PARAMETERS

The behaviour of the RORPO framework is controlled by several parameters: The path length, the robustness to noise parameter, and the set of chosen path opening orientations. In practice, the set of path opening orientations is fixed according the dimension of the image (2D or 3D). The robustness to noise parameter is usually set to 0, 3 or 5 depending on the degree of noise in the image; and the path length should be set according to the size of the curvilinear structures one wants to preserve. In this section, we present the behavior of RORPO according to these parameters and how to set them.

6.3.1 Path length

The path length is de facto the only real tunable parameter. As the paths can be tortuous, the path length is not exactly the length of the structure one wants to preserve. Nonetheless an obvious correlation exists between the length and diameter of the structure of interest and the path length. This correlation makes the path length a meaningful parameter, which can be set according to the dimensions of the curvilinear structures to be detected.

In the following, we present a few experiments, which shed some light on how to set the path length according to the length, diameter and curvature of the curvilinear structures of interest.

Path length vs. diameter and length of the curvilinear structures

We computed the RORPO intensity on a straight curvilinear structure in 2D, with either a varying diameter or length. The results are shown in Figure 46. This experiment shows that the minimal path length preserving a straight curvilinear structure only depends on its diameter. Indeed, the minimal path length should be higher than the diameter of the curvilinear structure. Otherwise, all path opening orientations would preserve the structure, and would be considered as a blob-like structure.

The maximal path length detecting a straight curvilinear structure obviously depends on the curvilinear structure length but also on its diameter. Indeed, the path can be tortuous inside the curvilinear structure ; the degree of tortuosity depending on its diameter. This is why the maximal path length is usually larger than the true curvilinear structure length; the larger the diameter, the larger the maximal path length.

Path length vs. curvature of the curvilinear structures

All these experiments assumed straight curvilinear structures , which is rarely the case in real applications. Consequently, we also endeavored to assess the RORPO intensity feature behavior with respect to curvilinear structures presenting curvatures. We computed the RORPO intensity feature on a curvilinear structure with a fixed length and diameter, but with a varying curvature. The degree of curvature is measured by the angle between the two curvilinear structure extremities (see Figure 47). The results of this experiment are shown in Figure 48.

This experiment shows that the range of suitable path lengths decreases as the degree of curvature of the curvilinear structure increases.

A curvilinear structure with a large-scale curvature is usually preserved by more than one path opening orientation each of them preserving a part of the curved curvilinear structure (*cf* Figure 43). The maximal path length is then relative to the length of the parts detected, which is obviously smaller than the whole curvilinear structure. In this case, when the curvature is sufficiently high (more than 80°), each orientation only preserves half of the 120 pixel long curvilinear structure . Consequently, the maximal path length is around 60 pixels.

The minimal path length increases with the curvature of the curvilinear structure. The reason is less intuitive because this is an artifact of the limit orientations in 2D. We did not handle them in 2D as their only impact is on this minimal path length. In 3D the minimal path length would not vary with the curvature.

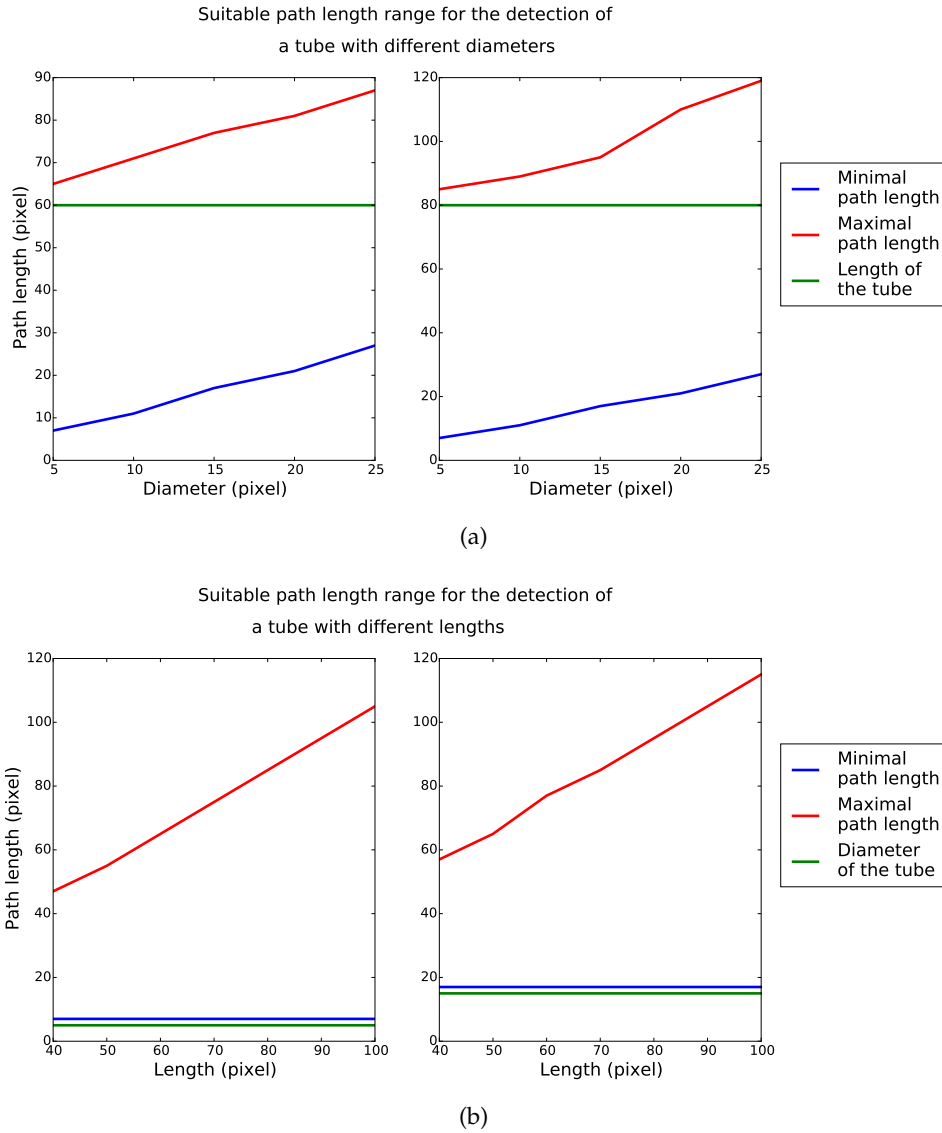


Figure 46: Suitable path length range for the detection of a binary tube with varying lengths (a) or diameters (b). The blue (resp. red) curve shows the minimal (resp. maximal) path length which detects the curvilinear structure for different lengths or diameters. All path length values between the blue and red curves may be used to detect the given curvilinear structure. The green curve indicates the fixed value of the other curvilinear structure dimension. (a) A curvilinear structure with varying diameter but a fixed length of 60 pixels (left) or 80 pixels (right). (b) A curvilinear structure with varying length but a fixed diameter of 5 pixels (right) or 15 pixels (left).

6.3.2 Noise robustness parameter

The noise robustness parameter, N , corresponds to the size of the structuring element of the dilation (see Section 6.2). It is a tunable parameter but only with a short range, as we recommend to set it to 0, 3 or 5 depending on the degree of noise of the initial image. Setting




Figure 47: Curvilinear structures with varying curvatures.

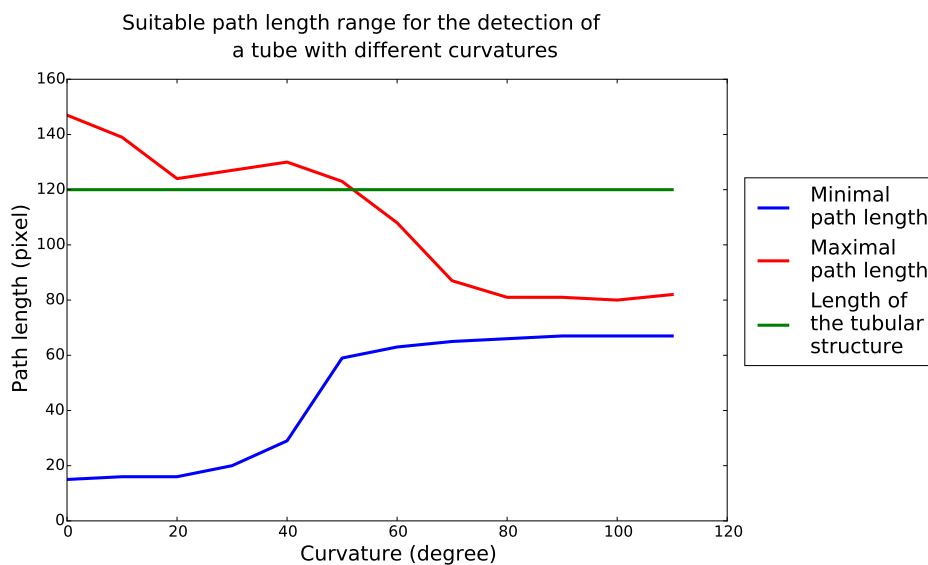


Figure 48: Suitable path length range for the detection of a binary tube with varying degree of curvature. The blue (resp. red) curve shows the minimal (resp. maximal) path length which detects the curvilinear structure for different lengths or diameters. All path length values between the blue and red curves may be used to detect the given curvilinear structure. The curvilinear structure has a fixed length of 120 pixels and a fixed diameter of 10 pixels.

the robust to noise parameter to a larger value tends to increase the detection of false positives more than it tends to preserve disconnected curvilinear structures. Setting $N = 3$ preserves curvilinear structures with disconnection that are at most 2 pixel long, while setting $N = 5$ preserves curvilinear structures with disconnections of at most 4 pixels. Nonetheless, it can also incorrectly preserve noise structures close (2 or 4 pixels) to a real curvilinear structure (see Figure 49).




Figure 49: An image (a) and its simplified robust path opening with path length L and with $N = 0$ (b), $N = 3$ (c) or $N = 5$ (d). The higher N , the more disconnected tubular parts are preserved but the more false positives are also detected.

6.3.3 Orientations

The path opening orientation set is a fixed parameter. The RORPO framework is proposed based on a given orientation set (one in 2D and one in 3D). In this section we motivate the choice of these specific sets of orientations.

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. A path opening orientation is defined based on an adjacency relation. We have seen that this adjacency relation is a set of irreflexive, non-symmetric binary relations defining the local connexity in its 8 (in 2D) or 26 (in 3D) neighborhood. Each adjacency relation is composed of a principal vector, defining the global orientation, and other vectors surrounding the principal vector, in a cone shape, to allow local deviation from the global orientation. Based on these hypothesis, 4 unique global orientations can be defined in 2D and 13 in 3D.

In 2D we chose to use all 4 of global orientations. Using fewer orientations (*i.e.* 2) would obviously not have been sufficient either because of the low overlap or the lack of precision in the directional feature.

In 3D, several combinations are possible among the 13 unique global orientations while keeping a complete cover of the 3D space and a significant overlap of homogeneous orientations: 3, 7 or 13 orientations (see Figure 50). As in the 2-orientation case in 2D, considering only 3 orientations in 3D would not be enough. We showed in the previous chapter that 7 orientations are sufficient to compute both intensity and directional features. However, we may assume that 13 orientations also works.

We have performed the same experiments as in Section ?? but with 13 path opening orientations to set the threshold value, i_t , leading to a correct RORPO intensity feature. Based solely on these synthetic experiments it appears that setting $i_t = 8$ is a good choice since it allows for the preservation of more than 90% of curvilinear structures while removing all plane-like structures. In Chapter , we quantitatively compare the RORPO results with 7 and 13 orientations. Nonetheless, using 13 orientations implies several drawbacks:

- More path opening orientations means a higher computation cost. Indeed, 13 orientations instead of 7 would increase the computation time of RORPO by almost 2.




Figure 50: The 3 possible set of path opening global orientations in 3D. We only represent the main vector of each orientation.

- In 3D, the memory consumption can be significant. Indeed, images may easily contain several millions of pixels and to compute the PRF, each path opening orientation result must be stored in memory. 13 orientations means nearly doubling the memory usage of the algorithm.
- Increasing the number of orientations means reducing the aperture of each cone and so reducing the flexibility of paths inside each orientation. To preserve a curvilinear structure with high curvature, a smaller path length should then be used with 13 path opening orientations than with 7. Consequently, with 13 path opening orientations the path length become less meaningful.
- Based on the synthetic experiment to set the threshold i_t , more limit cases would appear with 13 path opening orientations. To preserve these limit cases, a post processing step would also be required and should be more complex than with 7 path opening orientations.

To conclude, we chose a 3D space discretization in 7 path opening orientations because it constitute the best trade-off between computational efficiency and accuracy.

In a previous work [?], we tried to reduce the aperture of the 7 path opening orientations to reduce the overlap between orientations in order to suppress the limit cases presented in the previous Chapter. However, this is not practical since the reduction of the overlap between orientations leads to the non detection of some curvilinear structures orientations. Indeed, parts of the curvilinear structures was not detected with such path opening set.

6.4 COMPUTATIONAL COST

The computational cost of RORPO is dominated by the path opening. Indeed, the robustness step only requires infimum/supremum operations which have a linear cost $O(|\Omega|)$ with respect to the size of the image Ω , and the pointwise rank filter is also in $O(|\Omega|)$. We recall that we use the Luengo algorithm to compute the path opening, which has a complexity in

$O(|\Omega| \log(L))$ with L the path length. To experimentally assess that our RORPO algorithm really obeys these complexity, we computed RORPO on 3D images with increasing sizes or path lengths. As the density of curvilinear structures may influence the computational time of the algorithm, we use Vascusynth [38] to generate synthetic images. Vascusynth generates synthetic images of vascular trees, where the density of curvilinear structures may be controlled. We generated several images with increasing size but with the same density of curvilinear structures. The computation time of both experiments is shown in Figure 51. Results show that the RORPO algorithm has a linear complexity with respect to the image size and logarithmic complexity with respect to the path chosen path length.




Figure 51: The computation time of RORPO according to the image size with a fixed density of curvilinear structures and path length (a) or according to the path length with a fixed image size and density of curvilinear structures (b).

RESULTS AND COMPARISONS

In this chapter, we validate experimentally the RORPO approach and we present a comparison study with three state of the art methods. We compare our results both on synthetic (Section 7.2) and real images (Section 7.3). The three compared methods along with the comparison framework are presented in Section 7.1. Even though RORPO works both in 2D and 3D, we only present results and comparisons on 3D images.

Contents

7.1	Comparison framework	80
7.1.1	Compared methods	80
7.1.2	Evaluation criteria	84
7.1.3	Parameters Optimization	86
7.2	Synthetic images	86
7.2.1	Intensity feature comparison	88
7.2.2	Directional feature comparison	89
7.3	Real images	91
7.3.1	Intensity feature comparison	91
7.3.2	Directional feature comparison	91

7.1 COMPARISON FRAMEWORK

7.1.1 Compared methods

We compared our results with different methods designed to detect or characterize curvilinear structures. We chose a vesselness measure, namely the Frangi vesselness, which is one of the most common. We also compare our results with a filter based on the structure tensor, called Hybrid Diffusion Filter with Continuous Switch (HDCS), and a filter based on image gradient flux called OOF. Finally, we also decided to compare our results with the RORPO variant using 13 orientations instead of 7 which we call the RORPO₁₃. This last comparison allows us to verify if the choice of 7 orientations instead of 13 really is the best trade-off between efficiency and accuracy.

In the following, we first describe each method.

7.1.1.1 RORPO₁₃

In Section 6.3.3, we proposed a variant to RORPO using 13 orientations instead of 7. This variant, called RORPO₁₃, is based on the very same principle as RORPO, but uses a different threshold value i_t as the number of path opening orientations is different. Let $\gamma_L^i(I)$ be the i^{th} ranked path opening with path length L , then RORPO₁₃ intensity feature, $\Phi_L^{13}(I)$, is defined by:

$$\Phi_L^{13}(I) = \gamma_L^1(I) - \gamma_L^7(I) \quad (57)$$

7.1.1.2 Frangi Vesselness (FV)

The Frangi Vesselness (FV) was proposed by Frangi *et al.* [33]. This filter provides a local, linear and multiscale measure of the "curvilinearity" of a structure, relying on the second derivatives of the image. A directional feature can also be derived from this filter. Let $I(x, y, z)$ be a 3D image, and I_σ be the convolution of I with a Gaussian of variance σ^2 . The Hessian matrix of I_σ , $\mathcal{H}(I_\sigma)$, is the square matrix of its second-order derivatives:

$$\mathcal{H}(I_\sigma) = \begin{bmatrix} \frac{\partial^2 I_\sigma}{\partial x^2} & \frac{\partial^2 I_\sigma}{\partial x \partial y} & \frac{\partial^2 I_\sigma}{\partial x \partial z} \\ \frac{\partial^2 I_\sigma}{\partial y \partial x} & \frac{\partial^2 I_\sigma}{\partial y^2} & \frac{\partial^2 I_\sigma}{\partial y \partial z} \\ \frac{\partial^2 I_\sigma}{\partial z \partial x} & \frac{\partial^2 I_\sigma}{\partial z \partial y} & \frac{\partial^2 I_\sigma}{\partial z^2} \end{bmatrix} \quad (58)$$

We have seen in Chapter 2, that the analysis of the eigenvalues and eigenvectors of the Hessian matrix allows to distinguish between curvilinear structures, plane-like and blob-like structures.

Based on this interpretation, Frangi *et al.* proposed a vesselness measure. Let λ_1 , λ_2 and λ_3 be the three eigenvalues of \mathcal{H} such that $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, and e_1 , e_2 and e_3 be their

associated eigenvectors. Then, for bright curvilinear structures on a dark background, the Frangi vesselness, $\mathcal{V}(p)$, at point p is defined by:

$$\mathcal{V}(p) = \begin{cases} 0 & \text{if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ (1 - \exp(-\frac{\mathcal{R}_A^2}{2a^2})) \exp(-\frac{\mathcal{R}_B^2}{2\beta^2})(1 - \exp(-\frac{S^2}{2c^2})) & \text{otherwise} \end{cases} \quad (59)$$

$$\text{with } \mathcal{R}_A = \frac{|\lambda_2|}{|\lambda_3|} \quad \mathcal{R}_B = \frac{|\lambda_1|}{\sqrt{|\lambda_2\lambda_3|}} \quad S = \sqrt{\lambda_1 + \lambda_2 + \lambda_3} \quad (60)$$

The two terms \mathcal{R}_A and \mathcal{R}_B of this measure have different behaviours depending on the local structure:

- For any p belonging to a curvilinear structure, \mathcal{R}_A is close to 1 and \mathcal{R}_B is close to 0, which results in $\mathcal{V}(p) \simeq 1$
- For any p belonging to a plane-like structure, \mathcal{R}_A is close to 0 and \mathcal{R}_B is close to 1, which results in $\mathcal{V}(p) \simeq 0$
- For any p belonging to a blob-like structure, \mathcal{R}_A is close to 1 and \mathcal{R}_B is close to 1 which results in $\mathcal{V}(p) \simeq 0$

The term S measures the presence of second-order structure in the image. Indeed, S is close to 0 if the local neighborhood does not contain any structure. A directional feature can easily be extracted from the Frangi filter. The direction of a possible curvilinear structure is that of the eigenvector associated to the smallest eigenvalue magnitude.

7.1.1.3 Hybrid Diffusion Filter with Continuous Switch (HDCS)

The Hybrid Diffusion Filter with Continuous Switch (HDCS) is a noise reduction filter proposed by Mendrik *et al.* [58]. HDCS aims at reducing the noise in 3D images while preserving edges of curvilinear and small blob-like structures. This filter introduces a continuous switch, which acts as a balance between two anisotropic diffusion filters, called Edge-Enhancing Diffusion (EED) and Coherence-Enhancing Diffusion (CED) [99, 100]. Both filters are based on the diffusion equation:

$$\frac{\partial u}{\partial t} = \text{div}(D\nabla u) \quad (61)$$

where div is the divergence operator, ∇u is the gradient of the image u and D is the diffusion tensor. Iteratively solving a finite difference version of this equation for a finite time results in a new image where diffusion takes place along the three axes of the 3D space. If D is constant, the diffusion is isotropic and identical whichever the point in the image. This results in a Gaussian blur. In the case of CED and EED, the diffusion is anisotropic and depends on the structure present in the image. The analysis of the image structures relies on the structure tensor.

Let K_s be a Gaussian kernel with standard deviation s and $I_\sigma = I * K_\sigma$ be a smooth image. The structure tensor of an image I , $S(I)$, is the matrix derived from the gradient as follows:

$$S(I) = K_\rho * (\nabla I_\sigma \nabla I_\sigma^T) \quad (62)$$

A spectral analysis of the structure tensor can be carried out. It is similar to that of the Hessian matrix, and can serve to locally characterize a structure. For both CED and EED, the diffusion tensor D depends on the three eigenvalues of the structure tensor, μ_1 , μ_2 and μ_3 , ($\mu_1 \leq \mu_2 \leq \mu_3$). The 3D EED filter enhances the plane-like structures while reducing the noise. The three eigenvalues of its diffusion tensor are defined by:

$$\begin{aligned} \lambda_{e_1} &= \begin{cases} 1 & \text{if } \mu_1 \leq 0 \\ 1 - \exp\left(\frac{-C}{(\frac{\mu_1}{l_e})^4}\right) & \text{if } \mu_1 > 0 \end{cases} \\ \lambda_{e_2} &= 1 \\ \lambda_{e_3} &= 1 \end{aligned} \quad (63)$$

where $C = 3.31488$, and $l_e \in \mathbb{R}$, the EED contrast parameter.

The 3D CED filter enhances curvilinear structures and small blob-like structures. The three eigenvalues of its diffusion tensor are defined by:

$$\begin{aligned} \lambda_{c_1} &= \alpha \\ \lambda_{c_2} &= \alpha \\ \lambda_{c_3} &= \begin{cases} 1 & \text{if } \mu_2 = 0 \text{ or } \mu_3 = 0 \\ \alpha + (1 - \alpha) \exp\left(\frac{-l_c^2 \ln 2}{\kappa}\right) & \text{otherwise} \end{cases} \end{aligned} \quad (64)$$

where $\kappa = \left(\frac{\mu_2}{\alpha + \mu_3}\right)^4$, $\alpha = 0.001$ and $l_c \in \mathbb{R}$ is the CED contrast parameter.

Several methods have been proposed that use either EED or CED, depending on the structure of the image [32, 31]. However, HDCS aims at merging both methods by computing a linear combination of the eigenvalues of the filters. The eigenvalues λ_{h_i} ($1 \leq i \leq 3$), of the HDCS diffusion tensor are defined by:

$$\lambda_{h_i} = (1 - \epsilon)\lambda_{c_i} + \epsilon\lambda_{e_i} \quad (65)$$

$$\text{with } \epsilon = \exp\left(\frac{\mu_2(l_h^2(\zeta - |\zeta|) - 2\mu_3)}{2l_h^4}\right) \quad ; \quad \zeta = \left(\frac{\mu_1}{\alpha + \mu_2} - \frac{\mu_2}{\alpha + \mu_3}\right)$$

where $\lambda_h \in \mathbb{R}$ is the HDCS contrast parameter. The value ζ is the ratio distinguishing between plane-like structure ($\zeta \gg 0$), blob-like structures ($\zeta \simeq 0$) and curvilinear structures ($\zeta \ll 0$).

7.1.1.4 OOF

Optimally Oriented Flux (OOF) [49] is a curvilinear structure detector based on the image gradient. The oriented flux is the amount of the gradient, projected along a direction, flowing out from a sphere. The direction minimizing the oriented flux is the direction of a possible curvilinear structure. More formally, let $f(x, r, \rho)$ be the outward oriented flux along direction ρ , flowing from a sphere S_r of radius r centered at point x :

$$f(x, r, \rho) = \int_{\partial S_r} \left((v(y) \cdot \rho) \rho \right) \cdot ndA \quad (66)$$

where ∂S_r is the surface of the sphere, v is the image gradient, n is the outward unit normal of ∂S_r at $y = x + rn$ and dA is an infinitesimal area on ∂S_r .

The oriented flux can be rewritten as $f(x, r, \rho) = \rho^T Q \rho^T$, where Q is a symmetric 3×3 matrix which each term results from the convolution of the image with a set of specifically designed filter. Law *et al.* [49] proposed to compute the optimal direction, ρ^* by minimizing the inward oriented flux subject to the constraint $\|\rho\| = \rho^T \rho = 1$ (see Figure 52):

$$\rho^* = \underset{\rho}{\text{minimize}} \left(\rho^T Q \rho + \lambda(1 - \rho^T \rho) \right) \quad (67)$$




Figure 52: The direction ρ maximizing the total outward flux is the curvilinear structure direction. Each image represents a curvilinear structure and a sphere centered at x , on which the flux can be computed. (a-c) represent the case of a flux computed along a direction $\vec{\rho}$, which is not the direction of the curvilinear structure; (d-e) represent the case of a flux computed along the curvilinear structure direction. $\vec{v}(y)$ is the image gradient at point y and $\|v\|$ its norm. For y on the edge of the curvilinear structure, $\|v\| \gg 0$ while everywhere else, $\|v\| \simeq 0$. If the flux is computed from these 3 cases, $f(x, r, \rho) = \sum_{i=1}^3 (g(y_i)) = \sum_{i=1}^3 (\vec{v}(y_i) \cdot \vec{\rho}) \cdot \vec{n}$. The direction ρ maximizing the outward flux is the curvilinear structure direction.

The solution of the previous equation turns out to be a generalized eigenvalue problem where ρ^* is the eigenvector of matrix Q :

$$Q \rho^* = \lambda \rho^* \quad (68)$$

The matrix Q may play the same role as the Hessian matrix in the detection of curvilinear structures, as vesselness measures can be derived from its eigenvalues. For example, Law *et al.* embedded the OOF eigenvalues into the Sato vesselness [49] to detect blood vessels.

		Value of the binary result	
		0	1
Value of the ground-truth	0	TN	FP
	1	FN	TP

Table 5: Classification of the errors according to the value of the result and the ground-truth.

The multiscale version of OOF consists of computing the eigenvalues of Q using different radii of spheres. To ensure a coherence between scales, each eigenvalue is normalized by $4\pi r^2$.

7.1.2 Evaluation criteria

To compare and analyze the results of both intensity and directional RORPO features, quantitative measures as long as ground-truth are required. Several classical measures have been developed in the literature, depending on the type of image, ground-truth and problem to deal with. This section presents several quantitative criteria and motivates their choice.

7.1.2.1 The ground-truth

Computing a quantitative criterion requires a ground-truth, *i.e.* a certified image of the expected result. The RORPO intensity feature aims at providing a grey-level image which has a high intensity for curvilinear structures and a low intensity for any other structures. The natural ground-truth for such a feature is a binary image composed of white pixels for curvilinear structures and black pixels, for the rest.

The RORPO directional feature provides for each pixel the direction (*i.e.* a vector) of a possible curvilinear structure going through this pixel. Consequently, the ground-truth is composed of a vector for each pixel belonging to a curvilinear structure. We note that the directional ground-truth is thus related to the intensity ground-truth. Moreover, in all directional ground-truths, we removed the junctions between several curvilinear structures, like the vessels bifurcations. Indeed, these junctions are locally not curvilinear structure but blob-like structure.

7.1.2.2 Intensity feature quantitative criteria

Quantifying the similarity of a grey-level RORPO intensity feature with a binary ground-truth requires a thresholding procedure. The RORPO intensity feature is thresholded at all its grey-level values, resulting in several binary images. Each of these binary images are compared pixel-wise to the ground-truth and the total number of true positives (TP), true negatives (TN), false positives (FP) and false negatives (FN) are computed (see Table 5).

The sum of the TP and FN is the total number of positives (P_{GT}) in the ground-truth whereas the sum of the TN and FP is the total number of negatives (N_{GT}) in the ground-truth. Images containing curvilinear structures are usually sparse, which means that $P_{GT} \ll N_{GT}$. Consequently, the number of FP and TN are potentially much higher than the maximum number of TP and FN. To present meaningful results, we define the false positive rate (FPR), true positive rate (TPR), false negative rate (FNR) and true negative rate (TNR) as follows:

$$\begin{aligned} \text{FPR} &= \frac{FP}{P_{GT}} & \text{TPR} &= \frac{TP}{P_{GT}} \\ \text{FNR} &= \frac{FN}{P_{GT}} & \text{TNR} &= \frac{TN}{P_{GT}} \end{aligned}$$

The TPR is also called the sensitivity while the FPR is also called the fall-out. The closer the TPR to 1 and FPR to 0, the best result. It is important to note that with these definitions, the FPR can exceed 1. For example, if $FPR = 2$, the evaluated method detected twice as many false positives as possible true positives.

Based on these error measures, several quantitative similarity criteria have been proposed. The two widely used criteria to compare filtering and segmentation results are the Accuracy (Eq. 69) and the Dice coefficient (Eq. 70). The closer these coefficients are to 1, the more similar the result to the ground-truth.

$$\text{Acc} = \frac{TP + TN}{TP + FP + FN + TN} \quad (69)$$

$$\text{Dice} = \frac{TP}{2TP + FN + FP} \quad (70)$$

However, these coefficients assume that $P_{GT} \simeq N_{GT}$ which is clearly untrue for sparse images of curvilinear structures. To cope with this problem, we chose a similarity criterion specially designed to deal with unbalance classes: the Matthews Correlation Coefficient (MCC) [57]. The MCC is defined as follows:

$$\text{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (71)$$

Similarly to the Accuracy and Dice, the closer to 1 the MCC is, the more similar the result is to the ground-truth. To illustrate the better behaviour of the MCC compared with the Accuracy and the Dice coefficient, we computed the three of them on synthetic images (see Figure 53). Whereas the Accuracy and the Dice coefficients evaluate the images to be equivalently similar to the ground-truth, the MCC highlights the differences.

The TPR, FPR and MCC are representative of one threshold of the evaluated result. To obtain a global vision of a grey-level result, a Receiver Operating Characteristic (ROC) curve is usually computed. A ROC curve is the curve of the TPR against the FPR at every grey-level value of the evaluated result. The closer the curve to the point $[0, 1]$ ($FPR = 0$ and $TPR = 1$), the more similar the result is to the ground-truth.

7.1.2.3 Directional Feature quantitative criteria

The directional feature is a vector field. The evaluation of this feature requires to compare two vector fields. At each pixel, we compare the vector of the directional feature with the vector of the ground-truth. The criterion we use is the error angle in degree. It is important to note that we look for the orientation of curvilinear structures, which means that the maximum error angle is 90° .




Figure 53: Illustration of the better behavior of the MCC for sparse images. (a) the ground truth, (b,c) two different segmentation results. Although (c) is qualitatively the best result, the Accuracy and Dice coefficients provide the same quantitative measure. The MCC is better able to compare these results as the MCC of (c) is higher than the one of (b). The difference is small in this case, but the larger (and sparser) the image, the larger the difference is.

The pointwise vector comparison is only possible where both the ground-truth and the directional feature are defined. Nonetheless, at some point, the directional feature can be defined whereas the ground-truth is not (*i.e.*, the point does not belong to a curvilinear structure). Conversely, the directional feature may not be defined whereas the ground-truth is. The first case is a FP and the second case a FN. Both cases are not taken into account by the error angle, but they are counted by the intensity feature criterion. Consequently, the FPR and FNR are also criteria for the evaluation of the directional feature.

7.1.3 Parameters Optimization

In the following experiments, we always present the results of each method with optimized parameters. The optimization was performed by an exhaustive search on the range of reasonable values for each parameter according to the MCC: the best set of parameters is the one maximizing the MCC. A few methods use several parameters, some of them are only scalar weights, with no physical meaning or are not relevant for our applications. Table 6 shows, for each method, which parameter we optimized and the value we took for the others.

7.2 SYNTHETIC IMAGES

Publicly available databases of 3D images of curvilinear structures and their associated ground-truth are rare. To cope with this problem in the case of tree-like curvilinear structures, a software package called VasuSynth [38] was proposed by Hamarneh *et al.*. VasuSynth

	Optimized parameters	Fixed parameters and their values
RORPO	$L_{\min}, L_{\max}, \text{nbScales}$	
Frangi	$\sigma_{\min}, \sigma_{\max}, \text{nbScales}$	$\alpha = 0.5, \beta = 0.5, c = 0.5$
OOF	$r_{\min}, r_{\max}, \text{nbScales}$	
HDCS	$\rho, \lambda_h, \lambda_c, \lambda_e$	$\eta = 1, \tau = 0.11, \alpha = 0.001$

Table 6: Optimization of the parameters. \star_{\min} (resp. \star_{\max}) is the minimum (resp. maximum) value of the scale parameter \star . nbScales is the number of scales chosen for the multiscale framework. η is the number of iterations and τ is the time step. Each fixed parameter is set according to the default value used by their authors.




Figure 54: First row: 2D slices of one of the synthetic grey-level image generated with Vas-cuSynth, with background inhomogeneity, and various levels of Gaussian white noise ($\sigma = 0, 10$ and 18). Second row: the associated "vascular model" ground-truth.

simulates 3D images of vessels in the presence of various amounts of noise, and provides their corresponding ground-truth.

Using this software, we produced 10 images of size $100 \times 100 \times 100$. Each image was degraded by 7 levels of additive white Gaussian noise, resulting in a database of 70 images. We also added to each image the same Gaussian Random Field background, in order to simulate undesired, non-homogeneous, smooth, blob-like features. An illustration of these images is presented in Figure 54.

In the following, we compare both our intensity and directional features with the results from the different methods presented above on this database.

7.2.1 Intensity feature comparison

We compare the intensity feature of RORPO₇ and RORPO₁₃ with the results from 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 (see Section 7.1.3). These parameters were then used to process the remaining nine images.

As the results are gray-level images, we thresholded each of them at every gray-level values and kept, for each result, the binary image maximizing the MCC. Table 7 shows the mean best MCC, over the 10 images, for each noise level, and Table 8 shows their associated mean TPR and mean FPR. The ROC curve of one of the 10 images is shown in Fig. 55, and the corresponding results for the best MCC are shown in Figure 56.

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 [33]	0.197	0.192	0.193	0.197	0.199	0.187	0.185
OOF [49]	0.825	0.829	0.818	0.813	0.811	0.772	0.755
HDCS [58]	0.798	0.807	0.792	0.777	0.780	0.725	0.700

Table 7: 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	0.784/ 0.026	0.795/0.031	0.753/0.038	0.743/0.043	0.732/ 0.040	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 [33]	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 [49]	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
HDCS [58]	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

Table 8: Filtering performances on synthetic images, for various levels of noise (Gaussian) – TPR/FPR scores.

We can see that RORPO generally performs better at detecting curvilinear structures, as it detects much less false positives than the other methods, while preserving a high TPR. It




Figure 55: The ROC curves comparing the filtering of one of the ten VascuSynth images with a noise level of $\sigma = 10$.

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 those of RORPO₇, but at the cost of a computation time almost double.

7.2.2 Directional feature comparison

The comparison of the directional feature was performed according to Section 7.1. Since the VascuSynth ground-truth consists of a union of cylinders of various orientations and diameters, its local directions can be computed easily and constitute the directional ground-truth.

Because OOF and HDCS do not provide directional information, we compare the results from RORPO₇ and RORPO₁₃ only with FV. We kept the same parameters as for the comparison of the intensity feature, and computed the directional feature from the three methods for each point belonging to a curvilinear structure according to its best threshold intensity feature result. For a fair comparison, we only compared directions at pixels which were detected as belonging to a curvilinear structure both by FV and RORPO. The FP and FN were already taken into account in the evaluation of the intensity feature.

The results are shown in Table 9. The best directions are provided by RORPO 7 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,




Figure 56: Intensity feature computed on one of the VascuSynth image with noise level $\sigma = 10$ (a), by RORPO (b), RORPO₁₃ (c), FV (d), OOF (e) and HDCS (f).

a more meaningful neighborhood for the analysis of curvilinear structures than the isotropic neighborhood used by Frangi.

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 to curvilinear structures.

Table 9: Directional feature performances on synthetic images, for various levels of noise – Mean values (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 [33]	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)

7.3 REAL IMAGES

Real images of curvilinear structures, with their associated ground-truth are difficult to find or produce. In one of our article [59], we used CTA images of coronary arteries of the Rotterdam Repository [81]. However, this database only provides the centerline of the vessels, which does not allow for a comparison as accurate as with a full segmentation ground-truth. To cope with this issue, the HeartFlow company kindly provided us a coronary CT scan image with its associated high-quality, manually defined full segmentation ground-truth. The directional ground-truth was computed by reference to that of the centerline, also provided by HeartFlow.

The evaluation of both intensity and directional features were performed in an area of 70 pixels around the coronaries ground-truth, to avoid the other curvilinear structure like the ribs. We also optimized the parameters of both methods on this image according to the MCC.

7.3.1 Intensity feature comparison

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

The quantitative results of the intensity feature are shown in Table 10 and the volume rendering of each result is presented in Fig. 57.

The FV yields the worst results as it detects much more FP. However it performs comparatively better than on the synthetic data. OOF and RORPO both present good results, but RORPO again performs the best as it detects less FP than OOF for a similar TPR, which is confirmed by the better MCC value.

7.3.2 Directional feature comparison

With optimal parameters, we computed both RORPO and FV directional features on this image and evaluated the direction error. Once again, we only compare directions for pixels detected as curvilinear structures both by RORPO and FV. An example of results is shown in Fig. 58 and quantitative results of the comparison are shown in Table. 10.

	<i>MCC</i>	<i>TPR</i>	<i>FPR</i>	<i>mean direction error (std)</i>
<i>RORPO</i> 7	0.541	0.557	0.475	16.9 (13.3)
<i>FV</i> [33]	0.405	0.492	0.911	14.8 (13.4)
<i>OOF</i> [49]	0.529	0.582	0.594	-

Table 10: Comparison of RORPO, OOF and FV results on a coronary CT scan.

We see that FV performs slightly better on the accuracy of directions. However it is important to note that this error is computed only for the voxels inside the segmentation ground-






Figure 57: The coronary ground-truth (a) and the filtering results with RORPO (b), OOF (c), FV (d) (volume rendering).

truth. This means that all the false positive directions detected are not taken into account in the error. Indeed, FV 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 conversely computes far fewer false positives directions.

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 FP, and so resulting in generally more accurate directions. In the second case, when the segmentation is already available, FV is preferable, but it was not the problem we sought to address.



(a)



(b)

Figure 58: Illustration of the RORPO (a) and FV (b) direction feature on the coronary CT scan provided by the HeartFlow company [89]

Part II

VARIATIONAL
REGULARIZATION FOR
CURVILINEAR STRUCTURE
FILTERING AND
SEGMENTATION

INTRODUCTION

In the previous part, we presented a new framework called RORPO, which extracts from an image two features characterizing curvilinear structures. In this second part, we propose a strategy for embedding these two features into a segmentation method to better handle curvilinear structures in images. Even though the two RORPO features are generic and can be embedded into any general frameworks, like machine learning or mathematical morphology, we chose to illustrate the possibilities of RORPO in a variational framework.

Many image processing problems can be expressed in the context of a variational formulation. Such approach consists of designing an energy functional, so that its minimization over a space of images provides the desired resulting image. More formally, let u be a 2D image of size $N \times N$ in the Euclidean space $X = \mathbb{R}^{N \times N}$ ($u \in X$). The general variational problem is defined as follows:

$$u^* = \operatorname{argmin}_{u \in X} E(u) \quad (72)$$

With $E(u)$ ($E : X \rightarrow \mathbb{R}$) an energy formulation of image u .

We usually distinguish two distinct components inside the energy: a data fidelity $\mathcal{D}(u)$ and a regularization term $\mathcal{R}(u)$:

$$E(u) = \mathcal{D}(u) + \mathcal{R}(u) \quad (73)$$

The data fidelity term drives the solution to the expected result (*e.g.* a noise free image or a binary segmentation), depending on the input image. Conversely the regularization term represents other constraints, which are not a function of the input image, but depend on intrinsic properties of the sought solution, such as the smoothness of the solution, geometric or domain priors.

Our contribution focuses on the regularization term. A classical regularization term generally used in denoising and segmentation applications is the norm of the image gradient. When the ℓ_1 norm is involved, the regularization term is called the total variation. We propose to embed both RORPO features into the gradient of the total variation, in order to obtain a directional regularization term. We show that this regularization term, by embedding the RORPO features, is better suited than classical regularization terms for the detection and segmentation of curvilinear structures in images.

In the following, we come back on the general variational problem and we explain why a regularization term is required; then we present previous work on total variation and the segmentation model we use (see Chapter 9). Then, we develop our directional total variation regularization term and present some experimental results in Chapter 10.

PREVIOUS WORK

In this chapter, we begin by introducing the historical inverse problem behind numerous variational problems in image processing. Then, we introduce the regularization and why it is an essential term in such problems (Section 9.1). In Section 9.2, we present both the Tikhonov and the Total variation regularizations, and we explain the different approaches which were developed to handle the non-smooth total variation term. Finally, Section 3 presents the Chan *et al.* segmentation model and a few algorithms developed to solve it.

Contents

9.1	Inverse problem and the need of a regularization term	100
9.2	Previous work on total variation in image analysis	101
9.2.1	Tikhonov regularization	101
9.2.2	Total variation	102
9.2.3	TV <i>vs.</i> Tikhonov regularization	103
9.2.4	Solving the total variation	104
9.3	Segmentation model	106
9.3.1	The Chan <i>et al.</i> model	106
9.3.2	Proximal splitting algorithms	106

9.1 INVERSE PROBLEM AND THE NEED OF A REGULARIZATION TERM

The formation of an image can be modeled as a linear inverse problem:

$$Hu + b = o \quad (74)$$

where $o \in X$ is the formed image (the observation), $u \in X$ is the unknown true image, $b \in X$ is a random noise, and H ($H : X \rightarrow X$) is a transformation depending on the application. For example, in image denoising H is the identity matrix, in image deblurring H is a convolution with a blurring kernel and in image reconstruction H can be the Radon transform for CT-scan images or the Fourier transform in the case of MRI images.

However, this problem is only solvable if it is well-posed, *i.e.* if Eq. 74 satisfies the three conditions of Hadamard:

- the solution exists;
- the solution is unique;
- the solution depends continuously on the data.

In reconstruction applications H is generally not invertible, as it is not a squared matrix, making the problem ill-posed. A classical approach to still solve this type of problem is the least square method. The solution, u^* , of Eq. 74 is found by minimizing a functional $F : X \rightarrow \mathbb{R}$:

$$u^* = \operatorname{argmin}_u F(u) \quad (75)$$

such that:

$$F(u) = \|Hu - o\|_2^2 \quad (76)$$

Admitting Eq. 75 is solvable, its solution is obtained using the pseudo inverse $H^+ = (H^T H)^{-1} H^T$:

$$u^* = H^+ o \quad (77)$$

It is interesting to note that if H is invertible, the pseudo inverse H^+ of H is its inverse ($H^+ = H^{-1}$). Consequently, solving Eq. 75 is exactly equivalent to solving Eq. 74.

Even though solving Eq. 75 instead of Eq. 74 ensures the existence of the solution, the problem can still be ill-posed (*e.g.* $H^T H$ not invertible). Moreover, in order to find a numerical solution, these problems are discretized, leading to possible discretization errors. Depending on the problem, the discretization errors may be so important that the numerical solution strongly deviates from the true (analytical) solution. Such problems are called ill-conditioned.

To deal with both ill-posed and ill-conditioned problems, a regularization term, $R(u)$, is classically added to the functional $F(u)$ resulting in a new minimization problem (see Eq. 78). The regularization term selects a unique and continuous solution while ensuring good numerical properties.

$$u^* = \operatorname{argmin}_u F(u) + R(u) \quad (78)$$

9.2 PREVIOUS WORK ON TOTAL VARIATION IN IMAGE ANALYSIS

The historic regularization term in denoising and segmentation applications is the Tikhonov regularization [91]. Later, Rudin, Osher and Fatemi [77], proposed to use the total variation to better preserve the image contours. Section 9.2.1 introduces the Tikhonov regularization term. Then, we present the total variation in Section 9.2.2 and the various approaches in the literature to solve it (Section 9.2.4). The reader is referred to [15] for an exhaustive presentation of the total variation in image analysis.

9.2.1 Tikhonov regularization

The Tikhonov regularization $T_{\Gamma}(u)$ [91] is a quadratic prior defined as follows:

$$T_{\Gamma}(u) = \|\Gamma u\|_2^2 \quad (79)$$

where $\|\cdot\|_2$ is the ℓ_2 norm of an image, also called the Frobenius norm (see Appendix A) and Γ is the Tikhonov matrix. In practice, three different Γ are often used, each used for a specific regularization purpose:

- $\Gamma = I$, the identity operator, aiming at minimizing the norm of u ;
- $\Gamma = \nabla$, the gradient, aiming at minimizing the variation of u ;
- $\Gamma = \Delta$, the Laplacian, which minimizes the curvature of u .

The Tikhonov regularization with the gradient operator is widely used in denoising and segmentation applications as it is easy to optimize. An example of denoising with the Tikhonov gradient regularization is shown in Figure 59.b.




Figure 59: Comparison of the denoising of image (a) with a Tikhonov regularization (b) and a TV regularization (c); image from [93].

9.2.2 Total variation

The Tikhonov regularization tends to induce a blur in the resulting image due to its quadratic nature (see Section 9.2.3). This is not ideal in a segmentation purpose where the position of the contours is crucial. To cope with this problem, Rudin, Osher and Fatemi (ROF) [77] proposed to replace the Tikhonov regularization by the total variation (TV). We first define the general continuous total variation, then we present its discrete expression.

9.2.2.1 Continuous definition

Let $f : \Omega \rightarrow \mathbb{R}$, $\Omega \in \mathbb{R}^n$, be a functional belonging to $L^1_{loc}(\Omega)$. The total variation of f , $TV(f)$, is defined as:

$$TV(f) = \sup \left\{ - \int_{\Omega} f(x) \operatorname{div} \phi(x) \, dx, \phi \in C^1_c(\Omega, \mathbb{R}^n), \|\phi(x)\|_{2,2} \leq 1, \forall x \in \Omega \right\} \quad (80)$$

where $C^1_c(\Omega, \mathbb{R}^n)$ is the set of continuously differentiable vector functions of compact support contained in Ω .

The TV may have other expressions depending on the nature of f :

TV OF A DIFFERENTIABLE FUNCTION Let $f \in C^1(\Omega, \mathbb{R})$, the total variation of such differentiable functional reduces to the following (see the proof in Appendix C.1):

$$TV(f) = \int_{\Omega} \sqrt{(f^x)^2 + (f^y)^2} \, dx \quad (81)$$

where $f^x = \frac{df}{dx}$ and $f^y = \frac{df}{dy}$.

TV OF THE INDICATOR FUNCTION Let E be a set with smooth boundary ($E \in C^{1,1}$) and ι_E the indicator function of E such that:

$$\iota_E(x) = \begin{cases} 1 & \text{if } x \in E \\ 0 & \text{otherwise} \end{cases}$$

The total variation of ι_E , $TV(\iota_E)$ is the perimeter of E (see the proof in Appendix C.2):

$$TV(\iota_E) = \int_{\partial\Omega} \mathbb{1} \, dx \quad (82)$$

where $\partial\Omega$ is the boundary of Ω .

This result may be interpreted in terms of image segmentation. Indeed, a binary segmentation image can be seen as an indicator function. In this case, minimizing the total variation of the segmentation is equivalent to the minimization of the perimeter of the segmentation.

9.2.2.2 Discrete definition for images

As we work on images, we are interested in the discrete expression of the TV. The discrete spatial derivatives (*i.e.* the gradient, see Appendix A.4) of an image are always defined. Then, the TV definition generally used is the discrete equivalent of Eq.(81).

Let u be a 2D image of size $N \times N$ on the Euclidean space $X = \mathbb{R}^{N \times N}$, $u \in X$. The total variation of u , $TV(u)$ is defined as follows:

$$TV(u) = \|\nabla u\|_{2,1} \quad (83)$$

where $\nabla u = (u_x, u_y)$ is the discrete gradient of u and $\|\cdot\|_{2,1}$ is the ℓ_1 norm of the ℓ_2 norm of the gradient (see Appendix A.6).

9.2.3 TV vs. Tikhonov regularization

ROF proposed a denoising model with TV instead of the Tikhonov regularization, to remove its blurring effect.

Let us recall both regularizations, and focus on the norm which differs:

$$TV(u) = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \left| \underbrace{\sqrt{(u_{ij}^x)^2 + (u_{ij}^y)^2}}_{\ell_2 \text{ gradient norm}} \right| \quad T_{\nabla}(u) = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \left(\underbrace{\sqrt{(u_{ij}^x)^2 + (u_{ij}^y)^2}}_{\ell_2 \text{ gradient norm}} \right)^2$$

The TV regularization aims at minimizing the absolute value, $|\cdot|$, of the ℓ_2 gradient norm, whereas the Tikhonov regularization minimizes its square. Figure 60 illustrates the behavior of both regularization terms.

TV penalizes small gradient variations more than Tikhonov. This is why TV results present larger "flat" zones (*i.e.* with constant intensity). Nonetheless, high gradient variations are more penalized by Tikhonov than TV. Large gradient variations generally correspond to the contours of objects in the image. Highly penalizing the contours tends to smooth them, which explains the blurring effect of the Tikhonov regularization. Conversely, the small penalization of high gradient variations of TV better preserves the contours of objects.




Figure 60: Illustration of the different behaviors of the TV and Tikhonov regularizations.

9.2.4 Solving the total variation

We recall the general variational problem we are studying:

$$\underset{u \in X}{\text{minimize}} \mathcal{D}(u) + \lambda \mathcal{R}(u) \quad (84)$$

where $\mathcal{D}(u)$ is a data fidelity prior depending on the application (*e.g.* denoising or segmentation), $\mathcal{R}(u)$ is the regularization term, and $\lambda \in \mathbb{R}$ is the regularization parameter which acts like a balance between both terms.

The classical approach to solve Eq. 84 is to solve its associated Euler-Lagrange equation.

$$\partial \mathcal{D}(u) + \lambda \partial \mathcal{R}(u) = 0 \quad (85)$$

where $\partial \mathcal{D}$ and $\partial \mathcal{R}$ are respectively the subgradients of \mathcal{D} and \mathcal{R} (see Appendix B.1).

Generally, $\mathcal{D}(u)$ is a quadratic function which makes its subgradient equal to its derivative. Nonetheless, in our case, $\mathcal{R}(u) = TV(u)$, which is not differentiable in zero because of the ℓ_1 norm. In the literature, two approaches have been proposed to deal with this non-differentiability issue: the TV approximation and the Chambolle approach [14].

9.2.4.1 TV approximation approach

As the TV is not differentiable in zero, a classical approach is to replace TV by an approximation which is differentiable. The differentiability problem lies in the absolute value function that composes the ℓ_1 norm. A good differentiable approximation of the absolute value is: $|x| \simeq \sqrt{x^2 + \epsilon}$, where $\epsilon \in \mathbb{R}$, is a fixed constant. The differentiable TV approximation, called TV_{diff} , is then defined as follows:

$$TV_{\text{diff}}(u) \simeq \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \sqrt{\|\nabla u\|_2^2 + \epsilon} \quad (86)$$

One can note that the smaller ϵ , the closer the approximation to the real TV (see Figure 61). The functional derivative of this approximation is given by:

$$\nabla TV_{\text{diff}}(u) = -\text{div} \left(\frac{\nabla u}{\sqrt{\|\nabla u\|_2^2 + \epsilon}} \right) \quad (87)$$

With this differentiable regularization term, the solution of Eq. 84 can be obtained by a simple gradient descent:

$$u_{n+1} = u_n - \tau \left(\nabla \mathcal{D}(u_n) - \text{div} \left(\frac{\nabla u_n}{\sqrt{\|\nabla u_n\|_2^2 + \epsilon}} \right) \right) \quad (88)$$

where $\tau \in \mathbb{R}$ is the gradient descent step.

This approach avoids the differentiability problem by approximating the TV. This leads to good results, but the computed solution remains an approximation. To get closer to the initial solution, ϵ should tend to zero. Nonetheless, the smaller ϵ the longer the gradient descent algorithm takes to converge. Indeed the smaller ϵ the higher the curvature of TV_{diff} which requires smaller gradient descent steps.




Figure 61: Various approximations of the total variation (red) for various values of ϵ (blue).

9.2.4.2 Chambolle approach

In 2004, Chambolle proposed an algorithm to minimize the total variation exactly. The TV minimization is performed by computing its proximity operator, prox_{TV} .

$$\text{prox}_{\lambda TV}(g) = \min_{u \in X} \frac{\|u - g\|_2^2}{2} + \lambda TV(u) \quad (89)$$

where $\lambda \in \mathbb{R}$. Chambolle proposed to solve this problem by a duality approach. We only present the outlines of this approach, but the reader is referred to [14] for the complete description.

Let $TV^*(u)$ be the convex conjugate (or Legendre-Fenchel transformation) of $TV(u)$. An interesting property of TV is that its conjugate is the characteristic function of a closed convex set K such that:

$$K = \left\{ \text{div } p, p \in X^2, \|p_{ij}\|_2 \leq 1, 0 \leq i, j \leq N - 1 \right\} \quad (90)$$

with $p = (p^1, p^2)$ and $\|p_{ij}\|_2 = \sqrt{(p_{ij}^1)^2 + (p_{ij}^2)^2}$.

More formally, TV^* may be denoted as:

$$TV^*(u) = \begin{cases} 0 & \text{if } u \in K \\ +\infty & \text{else} \end{cases} \quad (91)$$

In his article, Chambolle introduces the dual problem of Eq. 89:

$$\text{prox}_{\frac{TV^*}{\lambda}}(g/\lambda) = \min_{w \in X} \frac{\|w - g/\lambda\|_2^2}{2} + \frac{1}{\lambda} TV^*(w) \quad (92)$$

Moreover, he shows that the solutions of both problems are linked by the relation:

$$\text{prox}_{\lambda TV}(g) = g - \lambda \text{prox}_{\frac{TV^*}{\lambda}}(g/\lambda) \quad (93)$$

In other words, Chambolle proposes to solve Eq. 92, then compute the solution of the primal problem (Eq. 89) with the above relation. Solving Eq. 92 is much simpler than Eq. 89. Indeed, one can identify Eq. 92 as the projection of g/λ onto the convex K , π_K :

$$\pi_K(g/\lambda) = \min_{w \in K} \frac{\|w - g/\lambda\|_2^2}{2} + \frac{1}{\lambda} TV^*(w) \quad (94)$$

which can be reformulated using Eq. 90 and 91 as:

$$\pi_{\lambda K}(g/\lambda) = \min_{p \in X^2} \left\{ \|\operatorname{div} p - g/\lambda\|_2^2, \quad \|p_{ij}\|_2 \leq 1, \quad 0 \leq i, j \leq N-1 \right\} \quad (95)$$

The latter equation is a convex optimization problem under constraint, which Chambolle propose to solve by a semi-implicit gradient descent as there is no longer a differentiability problem.

9.3 SEGMENTATION MODEL

Our main contribution lies in the regularization term we developed. This term can be used in several applications like denoising, segmentation or reconstruction. In this work, we chose to present the interest of our regularization term on segmentations applications. We present in this section the classical Chan segmentation model we used, and the different approaches in the literature to solve it.

9.3.1 The Chan et al. model

In this work, we use the segmentation model proposed by Chan, Esedoglu and Nikolova [17]:

$$\underset{u \in [0,1]^{N \times N}}{\text{minimize}} \langle c_f, u \rangle_F + \lambda \|\nabla u\|_{2,1} \quad (96)$$

where $(c_f)_{ij} = (c_1 - f_{ij})^2 - (c_2 - f_{ij})^2$, $c_f \in X$ and $\langle u, v \rangle_F = \sum_{i,j} u_{ij} v_{ij}$ is the Frobenius product. The scalar c_1 and c_2 are respectively the foreground and the background constant of f , the initial image ($f \in X$).

The data fidelity term, $\langle c_f, u \rangle_F$, was originally proposed by Chan and Vese in 2001 [18]. However, they proposed to solve the problem with active contours instead of convex optimization.

Eq. 96 is a convex optimization problem under constraint. Chan et al. proposed to handle the constraint with a Lagrangian relaxation. However, this approach introduces a new parameter, namely the Lagrange multiplier, which requires to be tuned additionally to the regularization parameter λ . Moreover, to cope with the differentiability problem of TV, they used the approximation of Eq. 86, which is also not ideal.

9.3.2 Proximal splitting algorithms

Proximal splitting algorithms were introduced to numerically solve convex optimization problems (see [23] for a presentation of such algorithms). A particularity of these algorithms

is that the differentiability condition is relaxed. Moreover, a few of them can handle constraints explicitly, which makes these algorithms good candidates to improve on the solution of Chan *et al.*

Several classes of algorithms have been proposed in the literature. Among the most commonly used are the Forward-Backward algorithm [24], the Douglas Rachford algorithm [29] and the Primal Dual algorithm [16]. In the following, for conciseness, we only present the Forward-Backward algorithm and two recent variants, namely FISTA and FGP.

9.3.2.1 Forward-Backward

The Forward-Backward algorithm considers the following problem:

$$\underset{x \in X}{\text{minimize}} \quad f_1(x) + f_2(x) \quad (97)$$

where $f_1 : X \rightarrow \mathbb{R}$ is a convex differentiable functional with a β -Lipschitz continuous gradient ∇f_1 , $f_2 : X \rightarrow \mathbb{R}$ is a convex non smooth functional ($f_2 \in \Gamma_0(X)$) and $f_1(x) + f_2(x) \rightarrow +\infty$ as $\|x\|_2 \rightarrow +\infty$.

Combettes and Wajs showed in [24] that Eq. 97 admits a solution given by the fixed point equation, suggesting the following iterative scheme:

$$x_{n+1} = \text{prox}_{\gamma_n f_1} \left(x_n - \gamma_n \nabla f_2(x_n) \right) \quad (98)$$

This algorithm is called Forward-Backward, as it can be seen as an explicit (forward) gradient descent step $\left(x_n - \gamma_n \nabla f_2(x_n) \right)$ of the differentiable function f_2 , followed by an implicit (backward) step using the proximity operator of the non differentiable function f_1 .

The full algorithm with a relaxation parameter (λ_n) , which speeds up its convergence, is presented in Algorithm 3.

Algorithm 3 : Forward-Backward algorithm

Data : $\epsilon \in]0, \min\{1, 1/\beta\}[$
 $x_0 \in \mathbb{R}^n$
for $n \geq 0$ **do**
 $\gamma_n \in [\epsilon, 2/\beta - \epsilon]$
 $y_n = x_n - \gamma_n \nabla f_2(x_n)$
 $\lambda_n \in [\epsilon, 1]$
 $x_{n+1} = x_n + \lambda_n \left(\text{prox}_{\gamma_n f_1}(y_n) - x_n \right)$
end

9.3.2.2 FISTA

The convergence of the Forward-Backward algorithm is guaranteed; however, its convergence rate can be rather slow. Several approaches have been investigated in order to increase its convergence rate, and amongst them, the Fast Iterative Shrinkage/Thresholding Algorithm (FISTA) [7] is often used. The full algorithm is presented in Algorithm 4. FISTA is based on the Nesterov acceleration [66] and achieves a $O(\frac{1}{k^2})$ convergence rate.

Algorithm 4 : FISTA**Data :** $x_0 \in \mathbb{R}^n, z_1 = x_0$ and $t_0 = 1$ **for** $n \geq 1$ **do**

$$y_n = z_n - \beta^{-1} \nabla f_2(z_n)$$

$$x_n = \text{prox}_{\beta^{-1} f_1}(y_n)$$

$$t_{n+1} = \frac{1 + \sqrt{4t_n^2 + 1}}{2}$$

$$\lambda_n = \frac{t_n - 1}{t_{n+1}}$$

$$z_{n+1} = x_n + \lambda_n(x_n - x_{n-1})$$

end

9.3.2.3 FGP

These algorithm are well suited to solve the Chan *et al.* segmentation problem, *i.e.* the minimization of a two term energy where one of them is the non differentiable TV. However, the expression of prox_{TV} is required. Moreover, the constraint $u \in [0, 1]^{N \times N}$ should also be dealt with. Both problems can be solved by the FGP algorithm proposed by Beck and Teboulle [6]. We will come back in detail on our methodology in the following chapter (see Section 10.3).

In their work, Beck and Teboulle proposed to solve the constrained TV problem defined as follows:

$$\underset{x \in \mathcal{C}}{\text{minimize}} \|x - o\|_2^2 + 2\lambda \text{TV}(x) \quad (99)$$

where $o \in X$ is the initial image and $\lambda \in \mathbb{R}$ is the regularization parameter. The constraint is expressed as a convex closed set \mathcal{C} , and is handled by an orthogonal projection, $P_{\mathcal{C}}$, onto this set. Even though any convex set would work, Beck and Teboulle focused on the case $\mathcal{C} = B_{a,b}$ of an n -dimensional cube given by:

$$B_{a,b} = \left\{ x, a \leq x_{ij} \leq b \right\}$$

This results in the following orthogonal projection:

$$\left(P_{B_{a,b}}(x) \right)_{ij} = \begin{cases} a & \text{if } x_{ij} < a \\ x_{ij} & \text{if } a \leq x_{ij} \leq b \\ b & \text{else} \end{cases} \quad (100)$$

Beck and Teboulle followed a similar approach than Chambolle by constructing the dual problem of Eq. 99 and solve it with a monotone version of their FISTA algorithm. This results in an algorithm called the Fast Gradient-Based Projection (FGP) which is presented in Algorithm 5.

where $P_{\mathcal{P}}$ is the projection onto the set \mathcal{P} . Beck and Teboulle defined \mathcal{P} as the set of matrix-pairs (p, q) satisfying:

Algorithme 5 : FGP

Data : $o \in X, x_0 \in X^2, z_1 = x_0$ and $t_0 = 1$ **for** $n \geq 1$ **do**

$$y_n = z_n + \frac{1}{8\lambda} \nabla \left(P_{B_{a,b}} [o - \lambda \operatorname{div}(z_n)] \right)$$

$$x_n = P_{\mathcal{P}}[y_n]$$

$$t_{n+1} = \frac{1 + \sqrt{4t_n^2 + 1}}{2}$$

$$\lambda_n = \frac{t_n - 1}{t_{n+1}}$$

$$z_{n+1} = x_n + \lambda_n(x_n - x_{n-1})$$

end

$$p_{ij}^2 + q_{ij}^2 \leq 1, \quad 0 \leq i, j < N - 1$$

$$|p_{i,n}| \leq 1, \quad 0 \leq i, j < N - 1$$

$$|q_{n,j}| \leq 1, \quad 0 \leq i, j < N - 1$$

In the next chapter, we solve the segmentation problem with our directional regularization, using a modified version of the FGP algorithm. We will come back in detail on the modifications in Section 10.3.

DIRECTIONAL TOTAL VARIATION

We present in this chapter the regularization term we developed to better regularize curvilinear structures. Section 10.1 introduces the motivation and our general strategy. The directional total variation is defined in Section 10.2 and the algorithm to solve the segmentation model with the directional total variation is presented in Section 10.3. The segmentation results on retinal images are finally presented in Section 10.4.

Contents

10.1	Motivation	112
10.2	The directional total variation	112
10.2.1	Directional gradient	113
10.2.2	Embedding the curvilinear structure features	115
10.3	Solving the segmentation model	117
10.4	Results and Comparisons	119
10.4.1	Compared methods	119
10.4.2	Synthetic image segmentation	119
10.4.3	Segmentation of the Drive database	122

10.1 MOTIVATION

We consider the classical Chan *et al.* segmentation model defined by:

$$\underset{u \in [0,1]^{M \times N}}{\text{minimize}} \langle c_f, u \rangle_F + \lambda TV(u) \quad (101)$$

where f is the initial image ($f \in X$), $(c_f)_{ij} = (c_1 - f_{ij})^2 - (c_2 - f_{ij})^2$, $c_f \in X$ and $\langle u, v \rangle_F = \sum_{i,j} u_{ij}v_{ij}$ is the Frobenius product. The scalar c_1 and c_2 are respectively the foreground and the background values, assumed constant, of f .

The regularization term (*i.e.* TV) is required to achieve a smooth segmentation. Indeed, we aim at obtaining a result with few low gradient variations, specially the small ones, that are generally due to noise and not structures of the image. The gradient norm minimization is performed to remove these "false" contours. However, the main drawback is that "real" contours may also disappear if the regularization is too strong (*i.e.* for high λ) or if the structures are not contrasted enough.

To sum up, TV is a regularization term which penalizes the image contours. Nonetheless, curvilinear structures are essentially composed of edges (*i.e.* contours) which makes them more susceptible to penalization than other structures and may result in their removal. This is especially the case for contours of low-contrast curvilinear structures, which are the most difficult to detect, both by human and computer methods. This makes the classical TV an unsuitable regularization term for images with curvilinear structures .

To cope with this problem, Miraucourt *et al.* [60] proposed a weighted TV regularization that includes the Frangi vesselness as tubular prior. They used this vesselness to determine where a curvilinear structure is likely to be, and proposed to regularize more strongly outside the putative curvilinear structures than inside. In general, this strategy effectively prevents curvilinear structures from disappearing, however noise is not dealt with inside the curvilinear structures, leading to many disconnections, and the removal of the smallest ones.

We propose a more suitable regularization term for curvilinear structures, by considering not only an intensity feature, as Miraucourt *et al.*, but also a directional feature. Instead of a weaker regularization inside curvilinear structures, we apply a strong regularization but solely along the local curvilinear structure axis. In other words, we propose an intensity and directional spatially variant regularization term (see Figure 62).

10.2 THE DIRECTIONAL TOTAL VARIATION

Our directional regularization term is based on the total variation. Its definition is recalled in Eq. 102.

$$\begin{aligned} TV(u) &= \|\nabla u\|_{2,1} \\ &= \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \sqrt{(u_{ij}^x)^2 + (u_{ij}^y)^2} \end{aligned} \quad (102)$$

To achieve a regularization with a spatially variant direction and intensity, we propose to replace the gradient, ∇ , by a new operator, ∇_d . This operator includes more spatial derivatives than the gradient, and embeds both intensity and directional information. By extension,

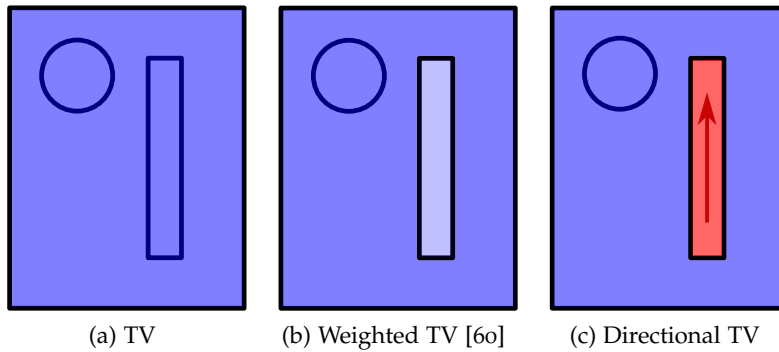


Figure 62: Regularization principle. The dark (resp. bright) blue areas represent a high (resp. low) isotropic (*i.e.* with no privileged direction) regularization, whereas the red color represents a directional regularization. The classical TV regularizes in all directions with the same intensity everywhere. The weighted TV regularizes with a low intensity inside curvilinear structures and a high intensity elsewhere, but always in all directions. Our directional TV regularizes with the same intensity inside and outside curvilinear structures; however, inside curvilinear structures, we only regularize along the curvilinear structure direction.

we call this new operator *directional gradient*, even though it is not strictly a gradient operator. More precisely, we define a general framework in which intensity and directional features may be embedded. To the best of our knowledge, the only alternative to our RORPO features are the Frangi vesselness features. However, as we showed in Part I that the RORPO features significantly outperformed the Frangi ones, we focus on the RORPO features in the following.

A few works have already proposed to adapt the definition of the gradient to make it directional. Peyré *et al.* [72] proposed a gradient computed on a weighted graph, generalized in [25]. Both the graph and its weights are computed, based on the initial image at each step of the optimization algorithm, resulting in an algorithm that may be time-consuming. In this work, we propose a directional TV in which both graph and weights are computed from the directional and intensity features only once, at the initial stage of the process. This results in a regularization that is well adapted to curvilinear structures but more efficient in terms of complexity.

10.2.1 Directional gradient

Let $V = (v_1, v_2, \dots, v_p)$ be a span of p unitary vectors. This span contains all the discrete undirected orientations in a $k \times k$ neighborhood (see Figure 63) resulting in $p = 2(n^2 - n + 1)$ with $n = \frac{k-1}{2}$.

The classical spatial gradient ∇ is defined along two directions, x and y , such that $\nabla u = (u^x, u^y)$, $\nabla u \in X^2$. We generalize this notion by introducing the directional gradient ∇_d ($\nabla_d \in X^p$) defined on the span V :

$$\nabla_d u = \left(D^1 \circ u^{v_1}, D^2 \circ u^{v_2}, \dots, D^p \circ u^{v_p} \right) \quad (103)$$

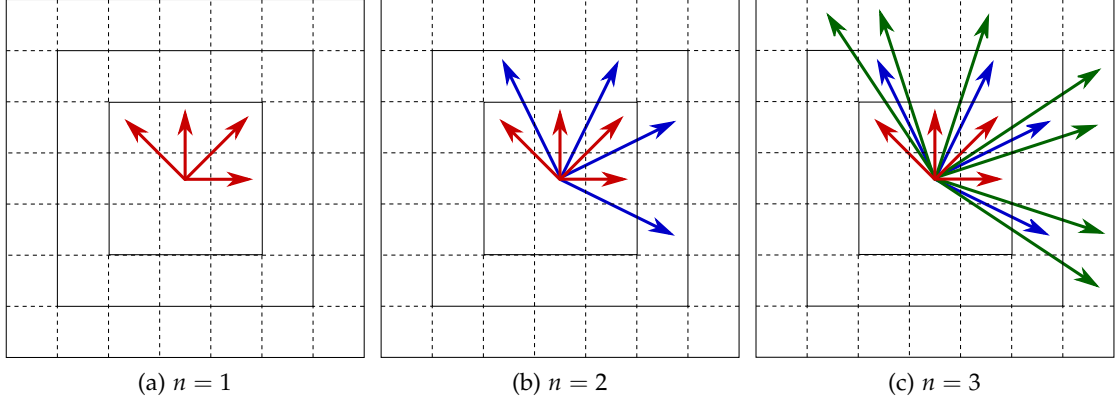


Figure 63: Example of spans for various n . Each square represents a pixel. The span with $n = k$, includes the directions of the previous span $n = k - 1$, plus new ones. The vector colors depict the directions specific of one n .

where u^{v_q} is the discrete derivative along direction v_q (by analogy with u^x and u^y which are the discrete derivatives along the image axes x and y), \circ is the Hadamard product of matrices such that $(A \circ B)_{ij} = A_{ij}B_{ij}$, and $D \in X^p$ is a weight matrix embedded the RORPO features. In other words, the directional gradient is a vector of weighted spatial derivatives. For example, in the case $n = 1$ (i.e. $p = 4$), the directional gradient of u at pixel (i, j) is defined as follows:

$$(\nabla_d u)_{ij} = \left(D_{ij}^1 u_{ij}^{v_1}, D_{ij}^2 u_{ij}^{v_2}, D_{ij}^3 u_{ij}^{v_3}, D_{ij}^4 u_{ij}^{v_4} \right) \quad (104)$$

The definition of this weight matrix requires to introduce the two features. Let $\Phi \in X$ be the RORPO intensity feature defined in Section 5.2, which has been normalized, and $d^q \in X$ be the directional coefficients along direction v_q extracted from the RORPO directional feature. We will define d more precisely in the following section. Both features are embedded into D so that:

$$\begin{aligned} D_{ij}^q &= d_{ij}^q \Phi_{ij} + (1 - \Phi_{ij}), \quad 1 \leq q \leq p \\ \sum_{q=1}^p D_{ij}^q &= 1, \quad 0 \leq i < M \text{ and } 0 \leq j < N \end{aligned} \quad (105)$$

INTUITIVE INTERPRETATION

- If a pixel (i, j) does not belong to a curvilinear structure:
The intensity feature of pixel (i, j) is null, $\Phi_{ij} = 0$; then $D_{ij}^q = 1 \forall q \in [1, p]$. It results that $(\nabla_d u)_{ij} = (u^{v_1}, u^{v_2}, \dots, u^{v_p})$, which is an isotropic gradient. Applying this gradient is equivalent to the classical spatially gradient ∇u , i.e. a regularization with no privileged direction.
- If a pixel (i, j) does belong to a curvilinear structure:
The intensity feature of pixel (i, j) is $\Phi_{ij} = 1$, then $D_{ij}^q = d_{ij}^q \forall q \in [1, p]$, which encodes the direction of the curvilinear structure. For example, if the curvilinear structure is oriented along v_1 , we would have $d_{ij}^1 = 1$ and $d_{ij}^q = 0 \forall q \in [2, p]$. It results that

$(\nabla_d u)_{ij} = (d_{ij}^1 u^{v_1}, d_{ij}^2 u^{v_2}, \dots, d_{ij}^p u^{v_p})$, which is an anisotropic gradient (since all d_{ij}^q are not equal).

The contours perpendicular to the curvilinear structure direction (*i.e.* the major part of its contours) do not appear in the directional gradient, while all the other structure contours still appear (see Figure 64). Consequently, minimizing this new gradient balances the over penalization of curvilinear structures, while still applying a regularization.




Figure 64: Illustration of the classical gradient (b) and the directional gradient (c) of an image (a). Unlike in the classical gradient, most of the curvilinear structure contours do not appear in the directional gradient.

The final directional total variation $TV_d(u)$ is defined by:

$$TV_d(u) = \|\nabla_d(u)\|_{2,1} \quad (106)$$

where ∇_d is the previously defined directional gradient.

10.2.2 Embedding the curvilinear structure features

The previous section presented our directional TV. In the following, we explain how to embed the curvilinear structure features in the case $n = 1$ (*i.e.* with the span $V = (v_1, v_2, v_3, v_4)$, see Figure 63.a).

We have seen that the directional gradient requires two curvilinear structure priors: an intensity feature Φ and a directional feature. We only require to normalize the intensity feature Φ between 0 and 1. However, the directional feature should be transformed to obtain the directions coefficients d^q .

Intensity feature

The RORPO intensity feature can not be interpreted as a probability of curvilinearity. As discussed in the previous chapters, RORPO has an anti-extensive behavior, which means that if a curvilinear structure has a low intensity, I_{low} in the initial image, its resulting intensity will be lower, or in the best case equal to I_{low} , but can never exceed it. Consequently, a classical normalization is not adequate as the low-contrast curvilinear structures would be considered to be less curvilinear than the high-contrast curvilinear structures.

We resolved this by thresholding the RORPO intensity feature to obtain a binary image. The threshold is easily set, and robust, as the intensity feature is generally very contrasted thanks to the top-hat like framework. In practice, other kinds of policies were experimented, but the results were similar given the quasi-binary shape of the RORPO intensity feature. The reader should note that the normalization policy should be adapted to the intensity feature. For example, in the case of the Frangi vesselness, a classical normalization may be a better choice.

Directional feature

Let $\Delta \in X^2$ be the normalized directional feature such that $\Delta_{ij} = (\Delta_{ij}^x, \Delta_{ij}^y)$ is the vector that gives the local direction of the curvilinear structure at pixel (i, j) . We propose to compute the four $(d^q)_{ij}$ as the decomposition coefficients of Δ_{ij} on the span V . In order to obtain a real directional gradient, we impose that the decomposition results in null coefficients except for those of the two span vectors bracketing Δ_{ij} (see Figure 65).




Figure 65: Decomposition of the direction vector at pixel (i, j) , δ_{ij} , on each vector of the span V .

More formally, let v_m and v_n be the two span vectors bracketing Δ_{ij} , then:

$$\Delta_{ij} = \sum_{q=1}^p d_{ij}^q v_q = d_{ij}^m v_m + d_{ij}^n v_n \quad (107)$$

From Eq. (107), we derive the linear system:

$$\begin{aligned} d_{ij}^m v_m \cdot v_m + d_{ij}^n v_n \cdot v_m &= \Delta_{ij} \cdot v_m \\ d_{ij}^m v_m \cdot v_n + d_{ij}^n v_n \cdot v_n &= \Delta_{ij} \cdot v_n \end{aligned}$$

As we consider the span $V = (v_1, v_2, v_3, v_4)$, we have $v_m \cdot v_n = \frac{\sqrt{2}}{2}$ because the angle between adjacent couple v_m and v_n is $\frac{\pi}{4}$. Then, we simply deduce:

$$\begin{aligned} d_{ij}^m &= 2\Delta_{ij} \cdot v_m - \sqrt{2}\Delta_{ij} \cdot v_n \\ d_{ij}^n &= 2\Delta_{ij} \cdot v_n - \sqrt{2}\Delta_{ij} \cdot v_m \end{aligned}$$

10.3 SOLVING THE SEGMENTATION MODEL

Our directional TV can be embedded in various variational image processing frameworks. To illustrate the advantages of this new regularization, we replace the classical TV of the Chan segmentation model (Eq. 101) by our directional TV (see Eq. 108).

$$\underset{u \in [0,1]^{N \times N}}{\text{minimize}} \langle c_f, u \rangle_F + \lambda \text{TV}_d(u) \quad (108)$$

This problem is equivalent to the following:

$$\underset{u}{\text{minimize}} \langle c_f, u \rangle_F + \lambda \text{TV}_d(u) + \iota_{[0,1]^{N \times N}}(u) \quad (109)$$

where $\iota_{[0,1]^{N \times N}}$ is the indicator function defined as follows:

$$\iota_{[0,1]^{N \times N}}(u) = \begin{cases} 0 & \text{if } u \in [0,1]^{N \times N} \\ +\infty & \text{otherwise} \end{cases}$$

It can be shown from [23], that the problem (108) admits a solution given by the *proximal point* splitting algorithm:

$$u_{n+1} = \text{prox}_{\gamma h}(u_n - \gamma c_f) \quad (110)$$

where $h = \lambda \|\nabla_d u\|_{2,1} + \iota_{[0,1]^{N \times N}}(u)$, with γ a step-size parameter.

To run this algorithm, the proximity operator of TV, constrained by $u \in [0,1]^{N \times N}$, $\text{prox}_{\gamma h}$, should be computed. We have seen in Section 9.3.2.3 that Beck and Teboulle [6] proposed an algorithm to solve the proximity operator of the constrained TV when the constraint is a closed convex set, as in our problem, namely FGP. Following, the Beck Teboulle notations, we note our constraint $u \in B_{0,1}$ where $B_{0,1}$ is the n -dimensional cube given by:

$$B_{0,1} = \left\{ x, 0 \leq x_{ij} \leq 1 \right\}$$

The complete FGP algorithm adapted to our directional TV is presented in Algorithm 6.

Algorithm 6 : FGP directional total variation

Data : $o \in X, x_0 \in X^4, z_1 = x_0$ and $t_0 = 1$

for $n \geq 1$ **do**

$$y_n = z_n + \frac{1}{8\lambda} \nabla_d \left(P_{B_{0,1}} [o - \lambda \text{div}_d(z_n)] \right)$$

$$x_n = P_p[y_n]$$

$$t_{n+1} = \frac{1 + \sqrt{4t_n^2 + 1}}{2}$$

$$\lambda_n = \frac{t_n - 1}{t_{n+1}}$$

$$z_{n+1} = x_n + \lambda_n(x_n - x_{n-1})$$

end

We recall the discrete gradient expression, $\nabla_d u = (u^{v_1}, u^{v_2}, u^{v_3}, u^{v_4})$ on the span $V = (v_1, v_2, v_3, v_4)$:

$$\begin{aligned}
 u_{ij}^{v_1} &= \begin{cases} u(i-1, j-1) - u(i, j) & \text{if } 0 < i, j < N \\ 0 & \text{otherwise} \end{cases} \\
 u_{ij}^{v_2} &= \begin{cases} u(i-1, j) - u(i, j) & \text{if } 0 < i < N ; 0 \leq j < N \\ 0 & \text{otherwise} \end{cases} \\
 u_{ij}^{v_3} &= \begin{cases} u(i-1, j+1) - u(i, j) & \text{if } 0 < i < N ; 0 \leq j < N-1 \\ 0 & \text{otherwise} \end{cases} \\
 u_{ij}^{v_4} &= \begin{cases} u(i, j+1) - u(i, j) & \text{if } 0 \leq i < N ; 0 \leq j < N-1 \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

And its adjoint discrete divergence, $\text{div}_d(p)$, $p = (p^1, p^2, p^3, p^4) \in X^4$:

$$\text{div}_d(p) = \begin{cases} p^1(i, j) - p^1(i+1, j+1) + p^2(i, j) - p^2(i+1, j) \\ \quad + p^3(i, j) - p^3(i+1, j-1) + p^4(i, j) - p^4(i, j-1) & \text{if } 0 < i, j < N-1 \\ -p_{i+1, j+1}^1 - p_{i+1, j}^2 + p_{i, j}^4 & \text{if } i = 0 ; j = 0 \\ -p_{i+1, j+1}^1 - p_{i+1, j}^2 - p_{i+1, j-1}^3 - p_{i, 1-1}^4 + p_{i, j}^4 & \text{if } i = 0 ; 0 < j < N-1 \\ -p_{i+1, j}^2 - p_{i+1, j-1}^3 - p_{i, j-1}^4 & \text{if } i = 0 ; j = N-1 \\ -p_{i+1, j+1}^1 + p_{i, j}^2 - p_{i+1, j}^2 + p_{i, j}^3 + p_{i, j}^4 & \text{if } 0 < i < N-1 ; j = 0 \\ p_{i, j}^1 + p_{i, j}^2 - p_{i+1, j}^2 - p_{i+1, j-1}^3 - p_{i, j-1}^4 & \text{if } 0 < i < N-1 ; j = N-1 \\ p_{i, j}^2 + p_{i, j}^3 + p_{i, j}^4 & \text{if } i = N-1 ; j = 0 \\ p_{i, j}^1 + p_{i, j}^2 + p_{i, j}^3 - p_{i, j-1}^4 + p_{i, j}^4 & \text{if } i = N-1 ; 0 < j < N-1 \\ p_{i, j}^1 + p_{i, j}^2 - p_{i, j-1}^4 & \text{if } i = N-1 ; j = N-1 \end{cases}$$

Moreover, the projection $P_p(k)$, $k = (k^1, k^2, k^3, k^4) \in X$, is defined by:

$$P_p(k) = l \tag{111}$$

such that $l = (l^1, l^2, l^3, l^4)$ with:

$$l_{ij}^q = \frac{k_{ij}^q}{\max(1, \mathcal{N}(k))} \quad \forall q \in [1, 4] \tag{112}$$

and:

$$\mathcal{N}(k) = \begin{cases} \sqrt{(k_{ij}^1)^2 + (k_{ij}^2)^2 + (k_{ij}^3)^2 + (k_{ij}^4)^2} & \text{if } 1 < i < N ; 1 < j < N-1 \\ |k_{i, j}^4| & \text{if } i = 0 ; j = 0 \\ 0 & \text{if } i = 0 ; j = N-1 \\ \sqrt{(k_{ij}^2)^2 + (k_{ij}^3)^2 + (k_{ij}^4)^2} & \text{if } 0 < i < N ; j = 0 \\ \sqrt{(k_{ij}^1)^2 + (k_{ij}^2)^2} & \text{if } 1 < i < N ; j = N-1 \end{cases}$$

10.4 RESULTS AND COMPARISONS

In this section, we first present the two methods we compared our results with. Then, we conduct a quantitative comparison both on synthetic image and real images.

10.4.1 Compared methods

We compare our results with two other segmentation models: the classical Chan *et al.* segmentation [17] and the weighted TV of Miraucourt *et al.* [60]. A comparison of the different behaviors of our regularization with these models is given in Section 10.1. We recall the expression of the three models:

10.4.1.1 Chan *et al.* model

$$\underset{u \in [0,1]^{N \times N}}{\text{minimize}} \langle c_f, u \rangle_F + \lambda TV(u) \quad (113)$$

where $f, f \in X$, is the initial image, $(c_f)_{ij} = (c_1 - f_{ij})^2 - (c_2 - f_{ij})^2$, $c_f \in X$ and $\langle u, v \rangle_F = \sum_{i,j} u_{ij}v_{ij}$ is the Frobenius product. The scalar c_1 and c_2 are respectively the foreground and the background values, assumed constant. $\lambda \in \mathbb{R}$ is the regularization parameter.

10.4.1.2 Weighted TV

$$\underset{u \in [0,1]^{N \times N}}{\text{minimize}} \langle c_f, u \rangle + \sum_{0 \leq i,j \leq N} \Lambda \sqrt{\|\nabla u\|_2} \quad (114)$$

where $\Lambda = \frac{1}{\lambda\alpha + (\alpha - 1)\Phi_{ij}}$ is the adaptive regularization parameter, $\lambda \in \mathbb{R}$ is the classical regularization parameter, and $\alpha \in [0,1]$. In the original article, Φ is the Frangi vesselness feature, however, for a fair comparison, we replace it with our RORPO intensity feature.

10.4.1.3 Directional TV

$$\underset{u \in [0,1]^{N \times N}}{\text{minimize}} \langle c_f, u \rangle_F + \lambda TV_d(u) \quad (115)$$

where $f, f \in X$, is the initial image, $(c_f)_{ij} = (c_1 - f_{ij})^2 - (c_2 - f_{ij})^2$, $c_f \in X$ and $\langle u, v \rangle_F = \sum_{i,j} u_{ij}v_{ij}$ is the Frobenius product. The scalar c_1 and c_2 are respectively the foreground and the background values, assumed constant. $\lambda \in \mathbb{R}$ is the regularization parameter.

10.4.2 Synthetic image segmentation

We first compared the three segmentation methods on a synthetic image containing lines and blob-like structures and corrupted by noise (see Figure 66). This experiment constitutes a proof of concept for our regularization term. We applied the three methods on this image with increasing values for the parameter λ . The three models have the same data fidelity parameter, but differ in their regularization terms. We expect that with low values of λ , the results of the three methods should be similar, since the regularization term is not important in the energy, but differences should appear and increase as λ becomes larger.



Figure 66: The synthetic initial image (a) and its RORPO intensity feature (b) and RORPO directional feature (c). The directions are encoded with colors, from 0° (black) to 180° (white).

Figure 67 presents the segmentation results on this synthetic image. We can observe that with low values of λ , the results are indeed similar, and rather noisy, as expected, since little regularization takes place. When λ increases, the lines tend to disappear with the Chan segmentation, as expected, whereas weighted TV succeeds in preserving some of the lines better. Nonetheless, we can see that our model outperformed both methods, as the lines are better preserved at all the λ values. This is particularly clear at high values of λ . Our directional regularization preserves, but also reconnects the line-like structures, as we expected, which results in better segmentations.

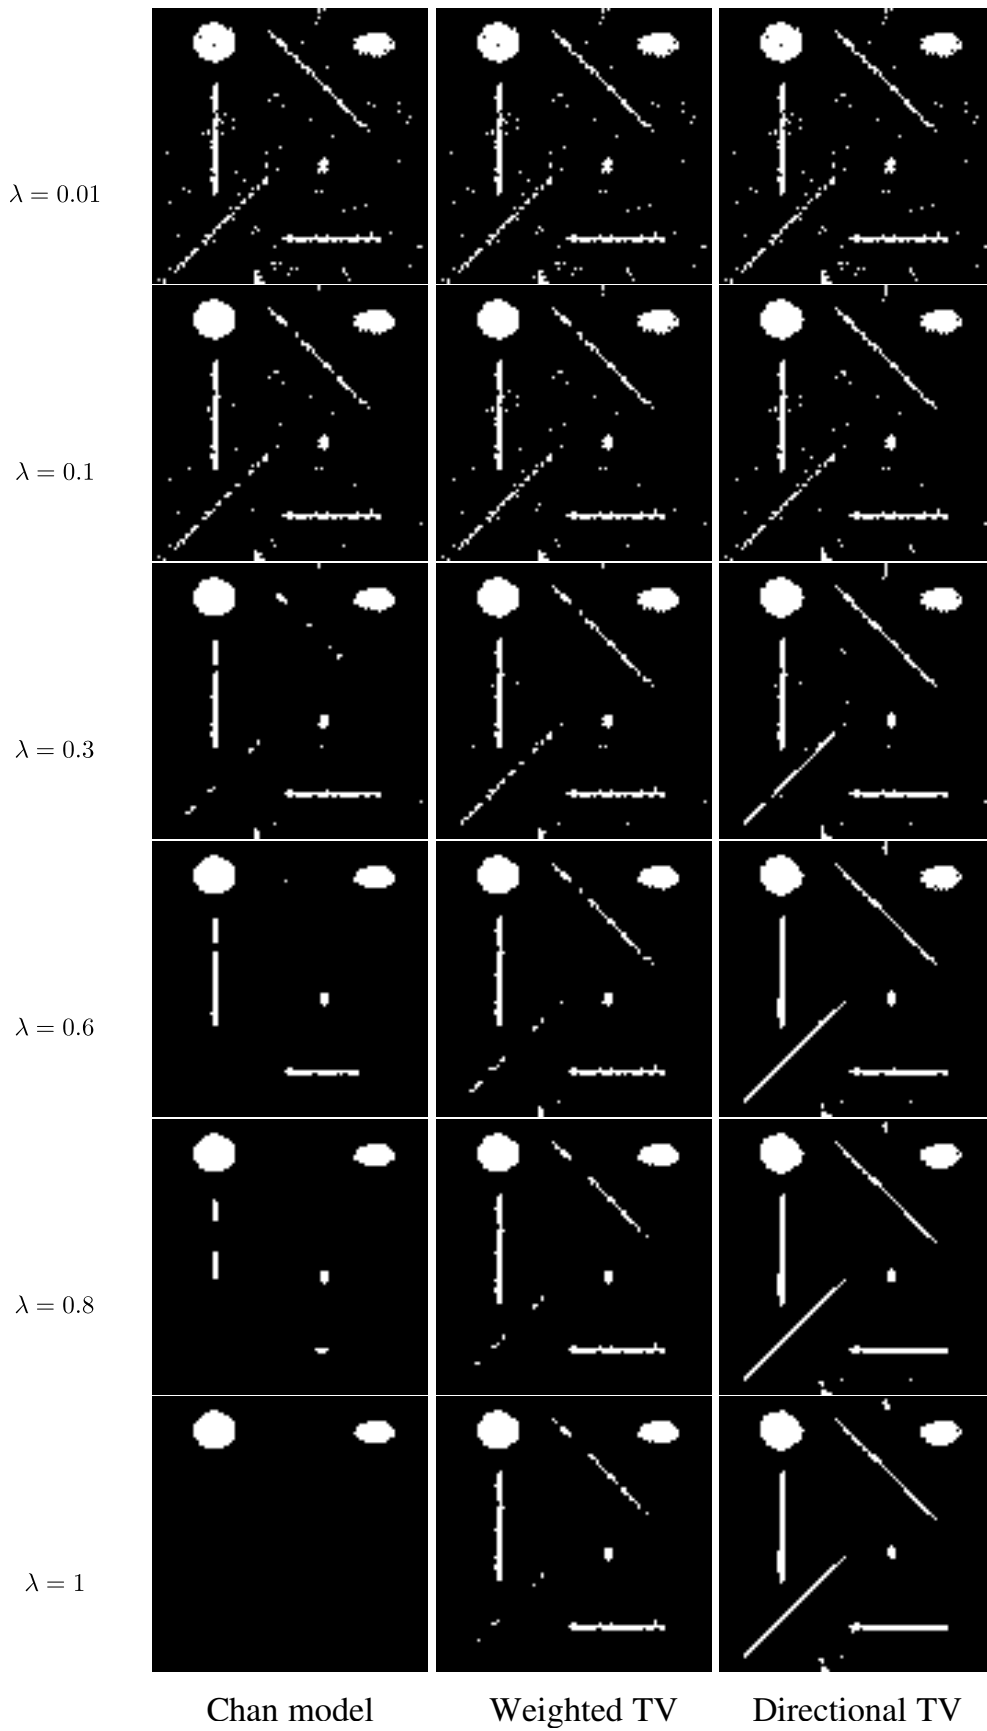


Figure 67: Comparison of three segmentation models on the synthetic image. The rows show the Chan segmentation results (left column), the weighted TV results (middle column) and our directional TV results (right column) for increasing λ .

10.4.3 Segmentation of the Drive database

To quantitatively analyze our directional TV, we segmented the 20 images of the testing set of the Drive database [84]. These are retinal images acquired using a Canon CR5 non-mydratic 3CCD camera, with a 45 degree field of view (FOV). We only analyzed the green channel, where the contrast between the vessels and the background is generally higher. Moreover, we subtracted the median filter result from the green channel, before each segmentation, in order to homogenize the image background, since the Chan *et al.* data fidelity assumes homogeneous background and foreground intensities.

Even though we strongly believe that the MCC is a better similarity criterion than the accuracy to evaluate images with curvilinear structures, in this experiment we are forced to use the accuracy (Acc) as the criterion over which to optimize the results. Indeed, most of the literature on the Drive database segmentation use this criterion. We also computed the true positives (TP) and true negatives (TN) on each best result. Along with the Drive images, Staal *et al.* provided the masks of the FOV. We computed all the quantitative criteria only inside these masks, to reduce the bias of the excessive number of background pixels.

Using the accuracy as criterion, we optimized the RORPO parameters and the regularization parameter, λ , on the first image of the test set, and applied these optimal parameters on the other images. The quantitative results are shown in Table 68. For comparison, we also mentioned some of the state-of-the art results (see [34] for a recent survey). The accuracy of our method is close to state of the art methods that are especially designed to segment these retinal images, such as learning based methods, whereas we propose a generic directional regularization term for curvilinear structures. Indeed, our regularization term is designed to be embedded in more complex variational framework such as [42], in lieu of the classical TV, to improve the segmentation results.

	TP	TN	Acc
Chan	0.656	0.985	0.9421
Proposed segmentation	0.690	0.981	0.9434
Staal [84]	-	-	0.9442
Lupascu [55]	0.6728	0.987	0.9597
Al-Rawi [2]	-	-	0.9535
Human observer	-	-	0.9470

Figure 68: Quantitative segmentation results on the DRIVE database.

The accuracy and TP of our method are higher than that of the classical Chan model. This confirms that our directional TV effectively improves the tubular structure segmentation. Figure 69 shows a visual comparison of the Chan segmentation and ours on the first DRIVE image. A zoom on a few blood vessels extremities shows that we successfully reconnect most small vessels. As these reconnections represent only a few pixels within the image, the accuracy of both methods does not appear quantitatively different, even if the qualitative improvement is real and significant. Indeed, a connected network is a much desired feature in blood vessels segmentation.

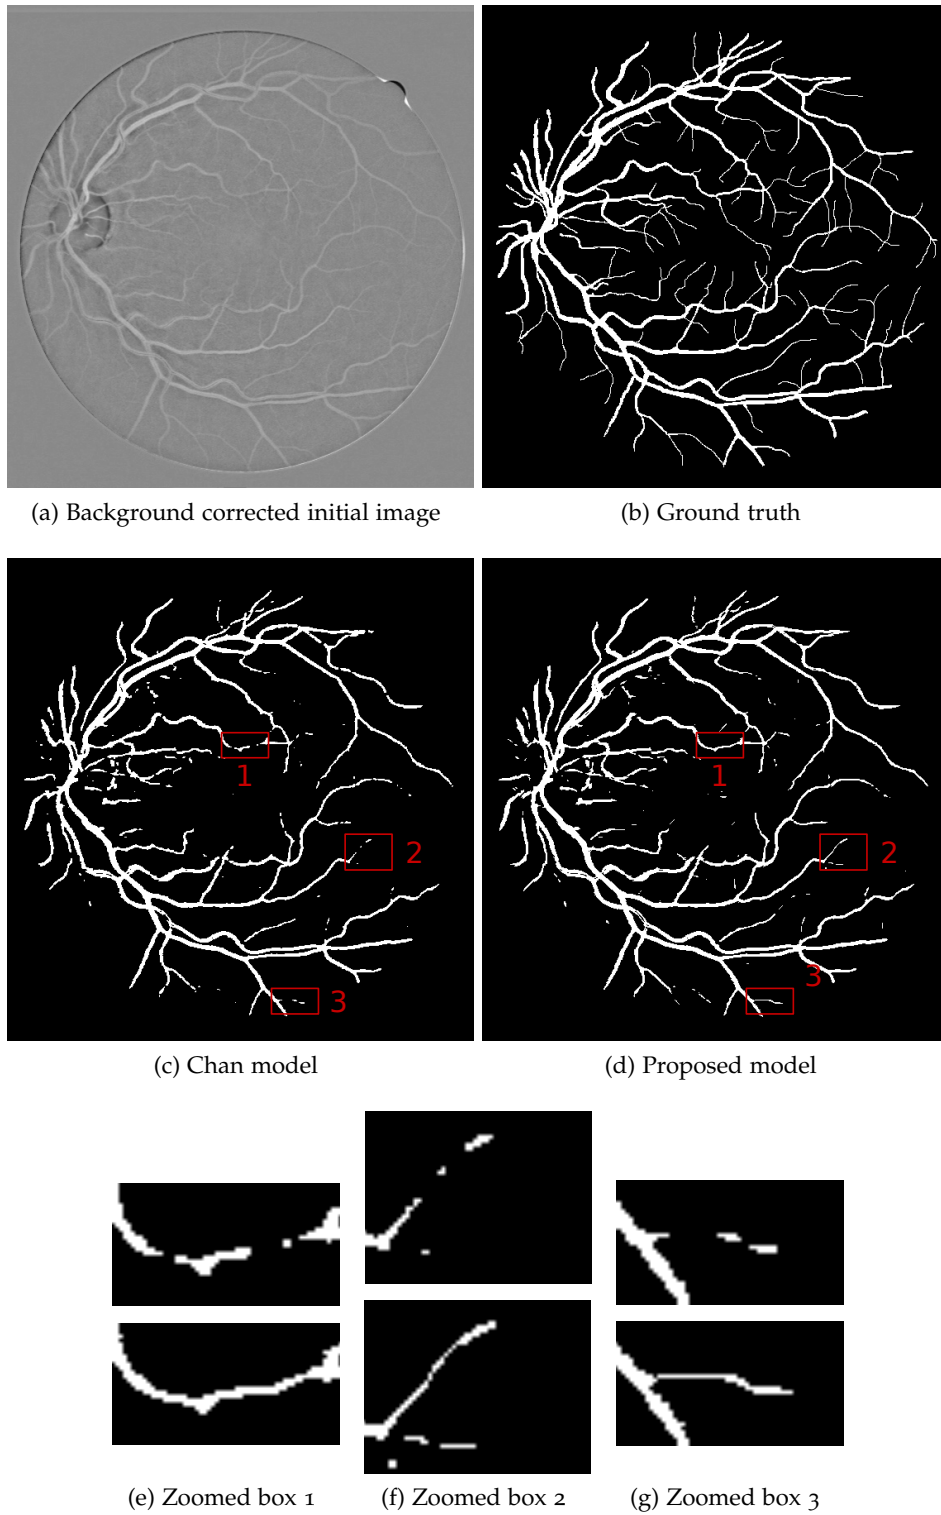


Figure 69: Comparison of the classical Chan *et al.* model (c) with our directional TV model (d) on the first DRIVE image (a) and its ground-truth (b). (e-g) Zoom on a few extremities of blood vessels of the Chan model results (top) and our proposed model results (bottom).

CONCLUSION

In this work, we have proposed a new framework, called RORPO, to characterize curvilinear structures. It consists of an intensity feature which preserves curvilinear structures while removing other structures, and of a directional feature providing at each point the orientation of the curvilinear structures. From a methodological point of view, RORPO is a relevant alternative to the existing methods in the literature. Indeed, most of them rely on local and isotropic eigen-analysis, which is not adapted to the intrinsic anisotropy of curvilinear structures. RORPO, by using path operators, considers non-local, flexible and anisotropic neighborhoods that are better adapted to curvilinear structures.

We have proposed a complete study of RORPO, including experiments on synthetic images to illustrate its behavior. We also have conducted a quantitative comparison study on both synthetic and real images. RORPO was compared with three state of the art methods, namely the Frangi Vesselness, the Optimally Oriented Flux, and the Hybrid Diffusion with Continuous Switch. This study have shown that the RORPO features outperform these other methods in almost every cases. In practice, the strength of RORPO are its noise robustness, its strong specificity to curvilinear structures and its user-friendliness, since it requires only very few, meaningful and intuitive parameters to tune.

As a low-level operator for the characterization of curvilinear structures, RORPO can be used, besides filtering, as a prior in higher-level methods. For instance, both RORPO features can directly be embedded in machine-learning frameworks.

In Part II, we have presented an application of these features for the segmentation of curvilinear structures. Indeed, we have developed a regularization term, based on the total variation, but better suited for curvilinear structures. It replaces the isotropic regularization of TV by an anisotropic regularization along the curvilinear structures axis using both RORPO features. This regularization term, unlike TV, is designed to avoid the loss of low-contrast curvilinear structures and also to reconnect the parts of the curvilinear structures which may have been disconnected by noise. After assessing the good behavior of this regularization term on a synthetic image, we have conducted a quantitative study on retinal images of the Drive database. This confirmed that our regularization term is better suited for curvilinear structures than the classical TV.

FUTURE WORK

During this work, various questions were raised, which could not be addressed in this work but represent interesting avenues for future work.

On RORPO

Recently, Morard *et al.* [61] proposed a parsimonious version of the path opening. It provides comparable results than the path opening, but drastically reduces its computation time. It would be very interesting to adapt RORPO with these parsimonious path openings to take advantage of their reduced computation time.

In its current version, the RORPO framework relies on a combinatorial analysis of the path opening responses for the only purpose of detecting curvilinear structures. A straightforward extension would be to adapt RORPO to specifically detect other structures such as plane-like structures in 3D or blob-like structures.

On the directional regularization term

We proposed a general definition of our regularization term for a span of n vectors. However, we only worked with $n = 4$, which corresponds to the immediate 3×3 neighborhood. There is no theoretical limitation in the adaptation of the directional TV to larger neighborhoods, but only practical decisions to make. For example we have to investigate different projection policies of the curvilinear structure direction on the span vectors.

Another interesting research lead would be to extend the directional TV for 3D images to work on the VIVABRAIN project angiography.

For now, we have focused on a binary segmentation framework, where both RORPO features are used as constant prior information. We plan to extend this work to a multi-label segmentation where the different labels would be the RORPO orientations. In this way we would be able to optimize on the RORPO directions together with the segmentation, which certainly would improve the segmentation result.

On the segmentation of curvilinear structures

To present results of our directional regularization term, we chose to work on the Chan *et al.* segmentation model. During the analysis of these results, we observe that the data fidelity term is not very effective. Indeed, the Chan *et al.* fidelity term can be interpreted as an optimal thresholding criterion, similar to that of Otsu. While this is interesting, it means that the Chan *et al.* formulation is essentially equivalent to apply a denoising followed by a thresholding. Moreover, the homogeneous background hypothesis and the two constants required by this data fidelity term are very restrictive. A possible future work would be to replace this data fidelity term, and propose a full high-level variational segmentation model for curvilinear structures. A possibility would be to replace the two constant of the background and foreground by spatially variant parameters which would relax the homogeneous background hypothesis.

We began, during my stay at HeartFlow, to work on the segmentation of blood vessels in 3D images using machine learning. We used random forest with a few features based on RORPO, in order to segment coronaries in CT-scan images. For now, this work is at an embryonic stage, but showed encouraging results, which should be pursued.

APPENDIX

NOTATIONS

We recall here a few notations used in Part II.

A.1 IMAGE

We define an image as a 2D matrix of size $M \times N$ on the Euclidean space $X = \mathbb{R}^{M \times N}$. For example, let $u \in X$ be a 2D image then:

$$\begin{aligned} u : [0, M-1] \times [0, N-1] &\rightarrow \mathbb{R} \\ (x, y) &\mapsto u(x, y) \end{aligned}$$

A.2 DOT PRODUCT OF IMAGES

The Frobenius dot product of two images u and v ($u, v \in X$) is defined by:

$$\langle u, v \rangle_F = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} u_{ij} v_{ij} \quad (116)$$

A.3 NORMS OF AN IMAGE

The ℓ_1 norm $\|u\|_1$ of an image $u \in X$ is defined by:

$$\|u\|_1 = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} |u_{ij}| \quad (117)$$

The ℓ_2 norm $\|u\|_2$ of an image $u \in X$ is defined by:

$$\|u\|_2 = \sqrt{\sum_{i=0}^{M-1} \sum_{j=0}^{N-1} u_{ij}^2} \quad (118)$$

A.4 GRADIENT OF AN IMAGE

We define the discrete gradient ∇u of an image $u \in X$ as:

$$(\nabla u)_{i,j} = \left(u_{ij}^x, u_{ij}^y \right) \quad \forall i, j \in \mathbb{R}^2 \quad (119)$$

With u^x (resp. u^y) the spatial gradient of u in the x direction (resp. in the y direction):

$$u_{ij}^x = \begin{cases} u(i-1, j) - u(i, j) & \text{if } 0 \leq i < N-1 \\ 0 & \text{if } i = N-1 \end{cases}$$

$$u_{ij}^y = \begin{cases} u(i, j+1) - u(i, j) & \text{if } 0 \leq j < N-1 \\ 0 & \text{if } j = N-1 \end{cases}$$

A.5 DIVERGENCE OF AN IMAGE

The divergence may be defined as the adjoint operator of the gradient. That is, $\forall u \in X$, $\forall p = (p^1, p^2) \in X \times X$:

$$\langle -\operatorname{div} p, u \rangle_X = \langle p, \nabla u \rangle_{X \times X} \quad (120)$$

which results in the following expression:

$$\operatorname{div}_a(p) = \begin{cases} p^1(i, j) - p^1(i+1, j) + p^2(i, j) - p^2(i, j-1) & \text{if } 0 < i, j < N-1 \\ -p_{i+1, j}^1 + p_{i, j}^2 & \text{if } i = 0; j = 0 \\ -p_{i+1, j}^1 - p_{i, j-1}^2 + p_{i, j}^2 & \text{if } i = 0; 0 < j < N-1 \\ -p_{i+1, j}^1 - p_{i, j-1}^2 & \text{if } i = 0; j = N-1 \\ p_{i, j}^1 - p_{i+1, j}^1 + p_{i, j}^2 & \text{if } 0 < i < N-1; j = 0 \\ p_{i, j}^1 - p_{i+1, j}^1 - p_{i, j-1}^2 & \text{if } 0 < i < N-1; j = N-1 \\ p_{i, j}^1 + p_{i, j}^2 & \text{if } i = N-1; j = 0 \\ p_{i, j}^1 - p_{i, j-1}^2 + p_{i, j}^2 & \text{if } i = N-1; 0 < j < N-1 \\ p_{i, j}^1 - p_{i, j-1}^2 & \text{if } i = N-1; j = N-1 \end{cases}$$

A.6 NORMS OF THE GRADIENT OF AN IMAGE

The ℓ_2 norm of the ℓ_2 norm of the gradient $\nabla u = (u^x, u^y)$ ($\nabla u \in X^2$) is defined by:

$$\|\nabla u\|_{2,2} = \sqrt{\sum_{i=0}^{M-1} \sum_{j=0}^{N-1} (u_{ij}^x)^2 + (u_{ij}^y)^2} \quad (121)$$

The ℓ_1 norm of the ℓ_2 norm of the gradient $\nabla u = (u^x, u^y)$ ($\nabla u \in X^2$) is defined by:

$$\|\nabla u\|_{2,1} = \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} \sqrt{(u_{ij}^x)^2 + (u_{ij}^y)^2} \quad (122)$$

SUBGRADIENT AND PROXIMITY OPERATOR

B.1 SUBGRADIENT

The subgradient is a generalization of the derivative to convex functions that are not differentiable everywhere. We recall that convex functions are always differentiable almost everywhere. Let f be a convex, proper, real-valued function on the Euclidean space $X = \mathbb{R}^n$. The subgradient of f , ∂f , is defined as follows:

$$\begin{aligned} \partial f : \mathbb{R}^n &\rightarrow 2^{\mathbb{R}^n} \\ x &\rightarrow \left\{ v \in \mathbb{R}^n, \forall y \in \mathbb{R}^n, f(y) - f(x) \geq \langle v, y - x \rangle \right\} \end{aligned} \quad (123)$$

where $2^{\mathbb{R}^n}$ is the power set of \mathbb{R}^n and $\langle \cdot, \cdot \rangle$ is the dot product.

In the case $n = 1$, the subgradient can be seen as the set of scalars u such that the line $T_{f_u}(x) = f(x) + (y - x)^T u$, is always beneath f . One can see that $T_f(x)$ is a generalization of the tangent equation when there is not a unique line at point x beneath f . The definition of the subgradient correspond to that of the normal gradient for smooth functions. In this case, the set u is reduced to the unique gradient of f , ∇f . Several examples of subgradients are illustrated in Figure 70 in the case $n = 1$. Equations 124 and 125 are respectively the subgradients of Figure 70.b and 70c.

$$\partial f_2(x) : \begin{cases} -1 & \text{if } x < 0 \\ 1 & \text{if } x > 0 \\ [-1, 1] & \text{else} \end{cases} \quad (124)$$

$$\partial f_3(x) : \begin{cases} \frac{1}{2}x^2 & \text{if } x < 0 \\ x & \text{else} \end{cases} \quad (125)$$




Figure 70: Examples of subgradients at x for a smooth function (a) (note that u is a unique scalar, $u = \{\nabla f(x)\}$) and at $x = 0$ for two non smooth functions (b)-(c). The green color show the areas where a line passing through the point $(x, f(x))$ (called T_{f_u}) is permitted while being beneath f , and so correspond to the allowed $u_i \in u$.

B.2 PROXIMITY OPERATOR

The proximity operator of a convex function is a natural extension of the notion of the projection operator onto a convex set C .

Let ι_C be the indicator function of the convex set C , defined as follows :

$$\iota_C(x) = \begin{cases} 0 & \text{if } x \in C \\ +\infty & \text{else} \end{cases} \quad (126)$$

ι_C is a convex but non differentiable function. Finding the projection of x onto C is equivalent to find the closest point to x onto C (see Figure 71). This problem can be formulated as follows:

$$y^* = \operatorname{argmin}_{y \in C} \left\{ \iota_C(y) + \frac{1}{2} \|y - x\|^2 \right\} \quad (127)$$

where y^* is the projection of x onto C .




Figure 71: y is the projection of x onto C

The proximity operator is the generalization of this projection to other functions than indicator functions. Let f ($f \in \mathbb{R}^n$) be a real-valued function, the proximity operator of f , prox_f is defined by:

$$\begin{aligned} \operatorname{prox}_f : \mathbb{R}^n &\rightarrow \mathbb{R}^n \\ x &\rightarrow \operatorname{argmin}_y \left\{ f(y) + \frac{1}{2} \|y - x\|^2 \right\} \end{aligned} \quad (128)$$

In the following, we present two examples of proximity operators.

Example 1 : 1D convex set




Figure 72: Projection of several x_i on C_1

Let C_1 be the convex set defined by its indicator function ι_{C_1} :

$$\iota_{C_1} = \begin{cases} 0 & \text{if } x \in [0,1] \\ +\infty & \text{else} \end{cases} \quad (129)$$

The proximity operator of C_1 is defined by (see Figure 72):

$$\text{prox}_{\iota_{C_1}}(x) = \begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x > 1 \\ x & \text{else} \end{cases} \quad (130)$$

Example 2 : 2D convex set


Let C_2 be a circle with indicator function ι_{C_2} , its proximity operator is defined in Eq. 131, see Figure 73.

$$\text{prox}_{\iota_{C_2}}(x) = \begin{cases} x & \text{if } \|x\|_2 = 1 \\ \begin{bmatrix} a \\ \frac{\|x\|_2}{b} \end{bmatrix} & \text{else} \end{cases} \quad (131)$$

B.3 LINK BETWEEN THE SUBGRADIENT AND THE PROXIMITY OPERATOR

The proximity operator of a convex non-smooth function f may be computed from the sub-gradient:

$$\forall (x, p) \in \mathbb{R}^N, \quad p = \text{prox}_f(x) \iff x - p \in \partial f(p) \quad (132)$$

Figure 73: Projection of x on C_2

Proof. By definition:

$$\text{prox}_f(x) = p = \underset{y}{\operatorname{argmin}} \left\{ f(y) + \frac{1}{2} \|x - y\|_2^2 \right\} \quad (133)$$

The Euler-Lagrange equation of Eq. 133 is given by:

$$\begin{aligned} 0 &\in \partial \left(f(p) + \frac{1}{2} \|x - p\|_2^2 \right) \\ \Leftrightarrow 0 &\in \partial f(p) + p - x \\ \Leftrightarrow x - p &\in \partial f(p) \end{aligned}$$

□

The reader may note that, for a convex smooth function g , Eq. 132 is equivalent to :

$$\text{prox}_f(x) = p = x - \nabla f(p)$$

which may be seen as an implicit gradient descent iteration :

$$\text{prox}_f(x) = x_{n+1} = x_n - \nabla f(x_{n+1})$$

The proximity operator plays the role of an iteration of the gradient descent for non smooth function

PROOF OF SOME STATEMENTS

C.1 TV OF A DIFFERENTIABLE FUNCTION

Let $f \in C^1(\Omega, \mathbb{R})$ be a functional. The total variation of f is defined as follows:

$$TV(f) = \int_{\Omega} \|\nabla f(x)\|_2 \, dx \quad (134)$$

Proof. As ∇ and $-\text{div}$ are adjoint, then:

$$\langle -\text{div}\phi(x), f(x) \rangle = \langle \phi(x), \nabla f(x) \rangle$$

which is equivalent to:

$$-\int_{\Omega} f(x) \text{div}\phi(x) \, dx = \int_{\Omega} \phi(x) \cdot \nabla f(x) \, dx$$

Consequently,

$$TV(f) = \sup \left\{ \int_{\Omega} \phi(x) \cdot \nabla f(x) \, dx, \phi \in C_c^1(\Omega, \mathbb{R}^n), \|\phi(x)\|_2 \leq 1, \forall x \in \Omega \right\}$$

One can choose a good ϕ for which the supremum is attained. Let $\phi(x) = \frac{\nabla f}{\|\nabla f\|_2}$ then $\|\phi(x)\|_2 \leq 1 \forall x \in \Omega$ and:

$$\begin{aligned} \int_{\Omega} \phi(x) \cdot \nabla f(x) \, dx &= \int_{\Omega} \frac{\nabla f}{\|\nabla f\|_2} \cdot \nabla f(x) \, dx \\ &= \int_{\Omega} \|\nabla f\|_2 \, dx \end{aligned}$$

□

C.2 TV OF THE INDICATOR FUNCTION

Let E be a set with smooth boundary ($E \in C^{1,1}$) and ι_E the indicator function of E such that:

$$\iota_E(x) = \begin{cases} 1 & \text{if } x \in E \\ 0 & \text{else} \end{cases}$$

The total variation of ι_E is the perimeter of E :

$$TV(\iota_E) = \int_{\partial E} \mathbf{1} \, dx$$

Where $\partial\Omega$ is the boundary of Ω and \mathbf{n} is the outward pointing unit normal field of $\partial\Omega$.

Proof. Let recall the definition of TV:

$$TV(\iota_E) = \sup \left\{ - \int_{\Omega} \iota_E(x) \operatorname{div} \phi(x) \, dx, \phi \in C_c^1(\Omega, \mathbb{R}^n), \|\phi(x)\|_2 \leq 1, \forall x \in \Omega \right\}$$

Then:

$$\int_{\Omega} \iota_E(\mathbf{x}) \operatorname{div} \phi(\mathbf{x}) \, d\mathbf{x} = \int_E \operatorname{div} \phi(\mathbf{x}) \, d\mathbf{x}$$

Applying the divergence theorem (Gauss's theorem) leads to:

$$\int_E \operatorname{div} \phi(x) \, dx = \int_{\partial E} \phi(x) \cdot n \, dx$$

Consequently,

$$TV(\iota_E) = \sup \left\{ - \int_{\partial E} \phi(x) \cdot n \, dx, \phi \in C_c^1(\Omega, \mathbb{R}^n), \|\phi(x)\|_2 \leq 1, \forall x \in \Omega \right\}$$

One can chose a good ϕ for which the supremum is attained. Let $\phi(x) = -\mathbf{n}$, then $\|\phi(\mathbf{x})\|_2 \leq 1$ and:

$$TV(\iota_E) = \int_{\partial E} \mathbf{1} \, dx$$

□

LIST OF PUBLICATIONS AND COMMUNICATIONS

JOURNAL

- O. Merveille, N. Passat, L. Najman and H. Talbot Tubular Structure Analysis by Ranking the Orientation Responses of Path Operators **In revision**

INTERNATIONAL CONFERENCE

- 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. (ECCV)*, volume 8690 of Lect. Notes Comput. Sci., pp. 203–218, Springer, 2014
- O. Merveille, H. Talbot, L. Najman, and N. Passat, "Ranking orientation responses of path operators: Motivations, choices and algorithmics" *In Proc. Int. Symp. Mathematical Morphology (ISMM)*, ser. Lect. Notes Comput. Sci., vol. 9082. pp. 633– 644, Springer, 2015
- O. Merveille, O. Miraucourt, S. Salmon, N. Passat, and H. Talbot, "A variational model for thin structure segmentation based on a directional regularization. *In IEEE Int. Conf. on Image Process. (ICIP)*, pp. 4324-4328, 2016
- N. Passat, S. Salmon, J.-P. Armspach, B. Naegel, C. Prud'homme, H. Talbot, A. Fortin, S. Garnotel, O. Merveille, O. Miraucourt, R. Tarabay, V. Chabannes, A. Dufour, A. Jezierska, O. Balédent, E. Durand, L. Najman, M. Szopos, A. Ancel, J. Baruthio, M. Delbany, S. Fall, G. Pagé, O. Génevaux, M. Ismail, P. Loureiro de Sousa, M. Thiriet, and J. Jomier, "From real MRA to virtual MRA: towards an open-source framework" *In Proc. Med. Image Computing and Comp.-Assist. Intervention (MICCAI)*, pp. 335–343, Springer, 2016

CONFERENCE WITHOUT PROCEEDINGS

- O. Merveille, H. Talbot and N. Passat, "Nouveaux opérateurs morphologiques pour la détection d'objets tubulaires", *37ème journée ISS France*, 2014 (Oral Presentation)
- O. Merveille, H. Talbot, L. Najman, and N. Passat, "Tubular Structure Filtering by Ranking Orientation Responses of Path Operators" *Reims Image 2015* (Poster Presentation)

ACRONYMS

ANR	Agence Nationale de la Recherche
CED	Coherence-Enhancing Diffusion
CT-scan	Computerized Tomography scanner
EED	Edge-Enhancing Diffusion
FGP	Fast Gradient-Based Projection
FISTA	Fast Iterative Shrinkage/Thresholding Algorithm
FN	false negatives
FNR	false negative rate
FP	false positives
FPR	false positive rate
FV	Frangi Vesselness
HDCS	Hybrid Diffusion Filter with Continuous Switch
IPO	Incomplete Path Opening
MCC	Matthews Correlation Coefficient
MIP	Maximum Intensity Projection
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging
OOF	Optimally Oriented Flux
PRF	Pointwise Rank Filter
ROC	Receiver Operating Characteristic
RORPO	Ranking Orientation Responses of Path Operators
RPO	Robust Path Opening
SE	Structuring Element
TN	true negatives
TNR	true negative rate
TP	true positives

TPR true positive rate

TV total variation

BIBLIOGRAPHY

- [1] G. Agam and C. Wu. Probabilistic modeling based vessel enhancement in thoracic ct scans. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, volume 2, pages 649–654 vol. 2, 2005.
- [2] Mohammed Al-Rawi, Munib Qutaishat, and Mohammed Arrar. An improved matched filter for blood vessel detection of digital retinal images. *Comput. Biol. Med.*, 37(2):262–267, 2007.
- [3] N. Armande, P. Montesinos, and O. Monga. A 3d thin nets extraction method for medical imaging. In *Proc. Int. Conf. Patt. Rec.*, volume 1, pages 642–646, 1996.
- [4] R. H. Bamberger and M. J. T. Smith. A multirate filter bank based approach to the detection and enhancement of linear features in images. In *Proc. Int. Conf. Acoust. Speech Signal Process.*, pages 2557–2560 vol.4, 1991.
- [5] Christian Bauer and Horst Bischof. A novel approach for detection of tubular objects and its application to medical image analysis. In *Proc. DAGM Symp. Patt. Recogn.*, pages 163–172, 2008.
- [6] Amir Beck and Marc Teboulle. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Trans. Image Process.*, 18(11):2419–2434, 2009.
- [7] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.
- [8] Fethallah Benmansour and Laurent D. Cohen. Tubular structure segmentation based on minimal path method and anisotropic enhancement. *Int. J. Comp. Vis.*, 92(2):192–210, 2011.
- [9] Vincent Bismuth, Régis Vaillant, Hugues Talbot, and Laurent 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)*, pages 9–16, 2012.
- [10] M. Buckley and H. Talbot. Flexible linear openings and closings. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, volume 18, pages 109–118, 2000.
- [11] Benoît Caldaïrou, Benoît Naegel, and Nicolas Passat. Segmentation of complex images based on component-trees: Methodological tools. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, pages 171–180, 2009.
- [12] John Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, (6):679–698, 1986.
- [13] C. Cañero and P. Radeva. Vesselness enhancement diffusion. *Pattern Recogn. Lett.*, 24(16):3141 – 3151, 2003.

- [14] A. Chambolle. An algorithm for total variation minimization and applications. *J. Math. Imaging Vision*, 20(1):89–97, 2004.
- [15] Antonin Chambolle, Vicent Caselles, Daniel Cremers, Matteo Novaga, and Thomas Pock. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery*, 9(263-340):227, 2010.
- [16] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [17] T.F. Chan, S. Esedoglu, and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM J. Appl. Math.*, 66:1632–1648, 2006.
- [18] T.F. Chan and Vese L.A. Active active contours without edges. *IEEE Trans. Image Process.*, 10(2):266–277, 2001.
- [19] Chris A. Cocosco, Vasken Kollokian, Remi K.-S. Kwan, G. Bruce Pike, and Alan C. Evans. Brainweb: Online interface to a 3d mri simulated brain database. *Neuromimage*, 5:425, 1997.
- [20] Laurent D. Cohen and Ron Kimmel. Global minimum for active contour models: A minimal path approach. *Int. J. Comp. Vis.*, 24(1):57–78, 1997.
- [21] F. Cokelaer, H. Talbot, and J. Chanussot. Efficient robust d-dimensional path operators. *IEEE J. Sel. Topics Signal Process.*, 6(7):830–839, 2012.
- [22] François Cokelaer. *Image processing and analysis for non destructive real time quality control in 3D RX tomography*. Theses, Université de Grenoble, February 2013.
- [23] Patrick L Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- [24] P.L. Combettes and V.R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.*, 4:1168–1200, 2005.
- [25] Camille Couprie, Leo Grady, Laurent Najman, Jean-Christophe Pesquet, and Hugues Talbot. Dual constrained tv-based regularization on graphs. *SIAM J. Imaging Sci.*, 6(3):1246–1273, 2013.
- [26] Thomas Deschamps and Laurent D. Cohen. Fast extraction of minimal paths in 3d images and applications to virtual endoscopy¹. *Med. Image Anal.*, 5(4):281 – 299, 2001.
- [27] Maxime Descoteaux, D. Louis Collins, and Kaleem Siddiqi. A geometric flow for segmenting vasculature in proton-density weighted {MRI}. *Med. Image Anal.*, 12(4):497 – 513, 2008.
- [28] David L Donoho and Ofer Levi. Fast x-ray and beamlet transforms for three-dimensional data. *Modern signal processing*, 46:79–116, 2002.
- [29] Jim Douglas and Henry H Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Trans. Amer. Math. Soc.*, 82(2):421–439, 1956.

- [30] 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.*, 17(2):147–164, 2013.
- [31] José-Jesús Fernández and Sam Li. An improved algorithm for anisotropic nonlinear diffusion for denoising cryo-tomograms. *J. Struct. Biol.*, 144(1–2):152 – 161, 2003.
- [32] Achilleas S. Frangakis and Reiner Hegerl. Noise reduction in electron tomographic reconstructions using nonlinear anisotropic diffusion. *J. Struct. Biol.*, 135(3):239 – 250, 2001.
- [33] A.F. Frangi, W.J. Niessen, K.L. Vincken, and M.A. Viergever. Multiscale vessel enhancement filtering. In *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, volume 1496 of *Lect. Notes Comput. Sci.*, pages 130–137. Springer, 1998.
- [34] Muhammad Moazam Fraz, Paolo Remagnino, Andreas Hoppe, Bunyarit Uyyanonvara, Alicja R Rudnicka, Christopher G Owen, and Sarah A Barman. Blood vessel segmentation methodologies in retinal images—a survey. *Comput. Meth. Programs Biomed.*, 108(1):407–433, 2012.
- [35] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(9):891–906, 1991.
- [36] Germán González, François Aguet, François Fleuret, Michael Unser, and Pascal Fua. *Steerable Features for Statistical 3D Dendrite Detection*, pages 625–632. Springer Berlin Heidelberg, 2009.
- [37] Jasper J. Gronde, Herman R. Schubert, and Jos B. T. M. Roerdink. Fast computation of greyscale path openings. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, pages 621–632. Springer, 2015.
- [38] Ghassan Hamarneh and Preet Jassi. Vascusynth: Simulating vascular trees for generating volumetric image data with ground truth segmentation and tree analysis. *Comput. Med. Imag. Graph.*, 34(8):605–616, 2010.
- [39] H. J. A. M. Heijmans, M. Buckley, and H. Talbot. Path openings and closings. *J. Math. Imaging Vision*, 22(2–3):107–119, 2005.
- [40] Georg Heygster. Rank filters in digital image processing. *Computer Graphics and Image Processing*, 19(2):148 – 164, 1982.
- [41] M. Jacob and M. Unser. Design of steerable filters for feature detection using canny-like criteria. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(8):1007–1019, 2004.
- [42] A. Jezierska et al. A non-local Chan-Vese model for sparse, tubular object segmentation. In *Proc. Int. Conf. Image Process.*, pages 907–911, 2014.
- [43] Thomas Markus Koller, Guido Gerig, Gabor Szekely, and Daniel Dettwiler. Multiscale detection of curvilinear structures in 2-d and 3-d image data. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 864–869. IEEE, 1995.
- [44] K. Krissian. Flux-based anisotropic diffusion applied to enhancement of 3-d angiogram. *IEEE Trans. Med. Imag.*, 21(11):1440–1442, 2002.

- [45] Karl Krissian, Grégoire Malandain, and Nicholas Ayache. Directional anisotropic diffusion applied to segmentation of vessels in 3d images. In *Proc. International Conf. Scale-Space Theo. Comput. Vis.*, pages 345–348, 1997.
- [46] Karl Krissian, Grégoire Malandain, Nicholas Ayache, Régis Vaillant, and Yves Troussel. Model-based detection of tubular structures in 3d images. *Comput. Vis. Image Understand.*, 80(2):130 – 171, 2000.
- [47] R. Kutka and S. Stier. Extraction of line properties based on direction fields. *IEEE Trans. Med. Imag.*, 15(1):51–58, 1996.
- [48] C. Lantuejoul and F. Maisonneuve. Geodesic methods in quantitative image analysis. *Pattern Recogn.*, 17(2):177 – 187, 1984.
- [49] Max WK Law and Albert CS Chung. Three dimensional curvilinear structure detection using optimally oriented flux. In *Proc. Eur. Conf. Comput. Vis.*, pages 368–382, 2008.
- [50] Max WK Law and Albert CS Chung. An oriented flux symmetry based active contour model for three dimensional vessel segmentation. In *Proc. Eur. Conf. Comput. Vis.*, pages 720–734, 2010.
- [51] 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.*, 13(6):819–845, 2009.
- [52] Qiang Li, Shusuke Sone, and Kunio Doi. Selective enhancement filters for nodules, vessels, and airway walls in two-and three-dimensional ct scans. *Med. Physics*, 30(8):2040–2051, 2003.
- [53] Cristian Lorenz, I-C Carlsen, Thorsten M Buzug, Carola Fassnacht, and Jürgen 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*, volume 1205 of *Lect. Notes Comput. Sci.*, pages 233–242, 1997.
- [54] C. L. Luengo Hendriks. Constrained and dimensionality-independent path openings. *IEEE Trans. Image Process.*, 19(6):1587–1595, 2010.
- [55] Carmen Alina Lupascu, Domenico Tegolo, and Emanuele Trucco. Fabc: retinal vessel segmentation using adaboost. *IEEE Trans. Inform. Tech. Biomed.*, 14(5):1267–1274, 2010.
- [56] Rashindra Manniesing, Max A. Viergever, and Wiro J. Niessen. Vessel enhancing diffusion: A scale space representation of vessel structures. *Med. Image Anal.*, 10(6):815 – 825, 2006.
- [57] B.W. Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2):442 – 451, 1975.
- [58] Adriënne M Mendrik, Evert-Jan Vonken, Annemarieke Rutten, Max A Viergever, and Bram van Ginneken. Noise reduction in computed tomography scans using 3-d anisotropic hybrid diffusion with continuous switch. *IEEE Trans. Med. Imag.*, 28(10):1585–1594, 2009.

- [59] 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.*, volume 8690 of *Lect. Notes Comput. Sci.*, pages 203–218. Springer, 2014.
- [60] O. Miraucourt, A. Jezierska, H. Talbot, S. Salmon, and N. Passat. Variational method combined with Frangi vesselness for tubular object segmentation. *Comput. Math. Biomed Eng.*, pages 485–488, 2015.
- [61] V. Morard, P. Dokládal, and E. Decenci re. Parsimonious path openings and closings. *IEEE Trans. Image Process.*, 23(4):1543–1555, 2014.
- [62] Vincent Morard, Etienne Decenci re, and Petr Dokládal. Efficient geodesic attribute thinnings based on the barycentric diameter. *J. Math. Imaging Vision*, 46(1):128–142, 2013.
- [63] Rodrigo Moreno and  rjan Smedby. Gradient-based enhancement of tubular structures in medical images. *Med. Image Anal.*, 26(1):19 – 29, 2015.
- [64] Beno t Naegel, Nicolas Passat, and Christian Ronse. Grey-level hit-or-miss transforms—part ii: Application to angiographic image processing. *Pattern Recognition*, 40(2):648–658, 2007.
- [65] Laurent Najman and Hugues Talbot, editors. *Mathematical morphology: from theory to applications*. John Wiley & Sons, 2010.
- [66] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $o(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.
- [67] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *Trans. Syst. Man Cybern.*, 9:62–66, 1979.
- [68] Georgios K Ouzounis and Michael HF Wilkinson. Mask-based second-generation connectivity and attribute filters. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(6):990–1004, 2007.
- [69] 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.*, 10(2):259 – 274, 2006.
- [70] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12(7):629–639, 1990.
- [71] 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.*, 37(6):1162–1176, 2015.
- [72] G. Peyr , S. Bogleux, and L. Cohen. Non-local regularization of inverse problems. In *Proc. Eur. Conf. Comput. Vis.*, volume 5304 of *Lect. Notes Comput. Sci.*, pages 57–68. Springer, 2008.
- [73] V. Prinet, O. Monga, C. Ge, S. L. Xie, and S. D. Ma. Thin network extraction in 3d images: application to medical angiograms. In *Proc. Int. Conf. Patt. Rec.*, volume 3, pages 386–390 vol.3, 1996.

- [74] Rangaraj M. Rangayyan, Fábio J. Ayres, Faraz Oloumi, Foad Oloumi, and Peyman Eshghzadeh-Zanjani. Detection of blood vessels in the retina with multiscale gabor filters. *J. Electronic Imaging*, 17(2):023018–023018–7, 2008.
- [75] Y. Rouchdy and L. D. Cohen. Image segmentation by geodesic voting. application to the extraction of tree structures from confocal microscope images. In *Proc. Int. Conf. Patt. Rec.*, pages 1–5, 2008.
- [76] Youssef Rouchdy and Laurent D. Cohen. Geodesic voting methods: overview, extensions and application to blood vessel segmentation. *Comput. Meth. Biomech. Biomed. Eng.*, 1(2):79–88, 2013.
- [77] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, 1992.
- [78] P. Salembier, A. Oliveras, and L. Garrido. Antiextensive connected operators for image and sequence processing. *IEEE Trans. Image Process.*, 7(4):555–570, 1998.
- [79] P. Salembier and H. Sanson. Robust motion estimation using connected operators. In *Proc. Int. Conf. Image Process.*, volume 1, pages 77–80 vol.1, 1997.
- [80] Yoshinobu Sato, Shin Nakajima, Hideki Atsumi, Thomas Koller, Guido Gerig, Shigeyuki Yoshida, and Ron Kikinis. 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In *Proc. joint CVRMed-MRCAS*, volume 1205 of *Lect. Notes Comput. Sci.*, pages 213–222, 1997.
- [81] Michiel Schaap, Coert T. Metz, Theo van Walsum, Alina G. van der Giessen, Annick C. Weustink, Nico R. Mollet, Christian Bauer, Hrvoje Bogunović, Carlos Castro, Xiang Deng, Engin Dikici, Thomas O’Donnell, Michel Frenay, Ola Friman, Marcela Hernández Hoyos, Pieter H. Kitslaar, Karl Krissian, Caroline Kühnel, Miguel A. Luengo-Oroz, Maciej Orkisz, Örjan Smedby, Martin Styner, Andrzej Szymczak, Hüseyin Tek, Chunliang Wang, Simon K. Warfield, Sebastian Zambal, Yong Zhang, Gabriel P. Krestin, and Wiro J. Niessen. Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Med. Image Anal.*, 13(5):701 – 714, 2009.
- [82] P. Soille and H. Talbot. Directional morphological filtering. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(11):1313–1329, 2001.
- [83] Pierre Soille, Edmond J. Breen, and Ronald Jones. Recursive implementation of erosions and dilations along discrete lines at arbitrary angles. *IEEE Trans. Pattern Anal. Mach. Intell.*, 18(5):562–567, 1996.
- [84] J.J. Staal, M.D. Abramoff, M. Niemeijer, M.A. Viergever, and B. van Ginneken. Ridge based vessel segmentation in color images of the retina. *IEEE Trans. Med. Imag.*, 23(4):501–509, 2004.
- [85] H. Talbot. *Oriented patterns in image analysis: from thin objects to flow-based methods*. Habilitation Thesis, Université Paris-Est, 2013.
- [86] H. Talbot and B. Appleton. Efficient complete and incomplete path openings and closings. *Image Vis. Comput.*, 25(4):416–425, 2007.

- [87] O. Tankyevych, H. Talbot, and P. Dokladal. Curvilinear morpho-hessian filter. In *Proc. IEEE Int. Symp. Biomed. Imaging*, 2008.
- [88] Olena Tankyevych, Hugues Talbot, Nicolas Passat, Mariano Musacchio, and Michel Lagneau. Angiographic image analysis. In *Medical Image Processing*, pages 115–144. Springer, 2011.
- [89] Charles A Taylor, Timothy A Fonte, and James K Min. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: scientific basis. *J. Am. Coll. Cardiology*, 61(22):2233–2241, 2013.
- [90] 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.*, 12(3):385–392, 1993.
- [91] Andrei Nikolaevich Tikhonov. Regularization of incorrectly posed problems. SOVIET MATHEMATICS DOKLADY, 1963.
- [92] Phan T.H. Truc, Md. A.U. Khan, Young-Koo Lee, Sungyoung Lee, and Tae-Seong Kim. Vessel enhancement filter using directional filter bank. *Comput. Vis. Image Understand.*, 113(1):101 – 112, 2009.
- [93] Markus Unger. *Convex optimization for Image Segmentation*. PhD thesis, Graz University of Technology, 2012.
- [94] Erik R. Urbach, Niek J. Boersma, and Michael H.F. Wilkinson. Vector-attribute filters. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, pages 95–104, 2005.
- [95] M. Van Droogenbroeck and M. J. Buckley. Morphological erosions and openings: Fast algorithms based on anchors. *J. Math. Imaging Vision*, 22(2):121–142, 2005.
- [96] Marcel van Herk. A fast algorithm for local minimum and maximum filters on rectangular and octagonal kernels. *Pattern Recogn. Lett.*, 13(7):517 – 521, 1992.
- [97] A. Vasilevskiy and K. Siddiqi. Flux maximizing geometric flows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(12), 2002.
- [98] L. Vincent. Minimal path algorithms for the robust detection of linear features in gray images. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, volume 12 of *Computational Imaging and Vision*, pages 331–338. Springer, 1998.
- [99] J. Weickert. *Theoretical Foundations of Anisotropic Diffusion in Image Processing*, pages 221–236. Springer Vienna, 1996.
- [100] Joachim Weickert. Coherence-enhancing diffusion filtering. *Int. J. Comp. Vis.*, 31(2):111–127, 1999.
- [101] R. Wiemker, T. Klinder, M. Bergtholdt, K. Meetz, I. C. Carlsen, and T. Bülow. A radial structure tensor and its use for shape-encoding medical visualization of tubular and nodular structures. *IEEE Trans. Visual Comput. Graph.*, 19(3):353–366, 2013.
- [102] Michael H. F. Wilkinson. Hyperconnectivity, attribute-space connectivity and path openings: Theoretical relationships. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, pages 47–58, 2009.

- [103] Michael HF Wilkinson and Michel A Westenberg. Shape preserving filament enhancement filtering. In *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, pages 770–777, 2001.
- [104] O. Wink, W. J. Niessen, B. Verdonck, and M. A. Viergever. Vessel axis determination using wave front propagation analysis. In *Proc. Med. Image Computing & Comp.-Assist. Intervention (MICCAI)*, pages 845–853, 2001.
- [105] A. Woiselle, J.-L. Starck, and J. Fadili. 3d curvelet transforms and astronomical data restoration. *Appl. Comput. Harmon. Analysis*, 28(2):171 – 188, 2010.
- [106] Chenyang Xu and J. L. Prince. Snakes, shapes, and gradient vector flow. *IEEE Trans. Image Process.*, 7(3):359–369, 1998.
- [107] Yongchao Xu, Thierry Géraud, and Laurent Najman. Two applications of shape-based morphology: Blood vessels segmentation and a generalization of constrained connectivity. In *Proc. Int. Symp. Mathematical Morphology (ISMM)*, pages 390–401, 2013.
- [108] F. Zana and J. C. Klein. Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation. *IEEE Trans. Image Process.*, 10(7):1010–1019, 2001.