

Scale and Curvature Invariant Ridge Detector for Tortuous and Fragmented Structures

Roberto Annunziata¹ *, Ahmad Kheirkhah², Pedram Hamrah², and Emanuele Trucco¹

¹ School of Computing, University of Dundee, Dundee, UK

² Massachusetts Eye and Ear Infirmary, Harvard Medical School, Boston, USA

Abstract. Segmenting dendritic trees and corneal nerve fibres is challenging due to their uneven and irregular appearance in the respective image modalities. State-of-the-art approaches use hand-crafted features based on local assumptions that are often violated by tortuous and point-like structures, e.g., straight tubular shape. We propose a novel ridge detector, SCIRD, which is simultaneously rotation, scale and curvature invariant, and relaxes shape assumptions to achieve enhancement of target image structures. Experimental results on three datasets show that our approach outperforms state-of-the-art hand-crafted methods on tortuous and point-like structures, especially when captured at low resolution or limited signal-to-noise ratio and in the presence of other non-target structures.

1 Introduction

Segmenting automatically curvilinear structures such as blood vessels, dendritic arbors, or corneal nerve fibres has been the subject of much research as a fundamental pre-processing step, e.g. for lesion quantification and evaluation of therapy progress [1–9]. Most of the methods estimate a local *tubularity measure* (e.g. *vesselness* in [2,9]) based on hand-crafted features (henceforth HCFs) modelling local geometrical properties of ideal tubular structures; enhancement filters are then built based on such models [2,3,5,9]. Recently, combining hand-crafted and learned filters, or using HCF in the learning process, has been shown to improve detection [7,8]. It is therefore important to identify HCF sets leading to high detection performance, to be used either separately or in combination with learning; this is the purpose of this paper. While HCF methods are typically fast, the result quality can be improved by capturing more general local structures found in blood vessels, dendritic arbors, or corneal nerve fibres. Specifically, highly fragmented and tortuous structures violate two usual assumptions of most HCF models, i.e. continuous and locally straight tubular shapes. While discontinuity

*Corresponding author. E-mail: r.annunziata@dundee.ac.uk. This research was supported by the EU Marie Curie ITN REVAMMAD, n° 316990. The authors are grateful to S. McKenna, J. Zhang (CVIP, Dundee) for valuable comments and to Amos Sironi (CVlab, EPFL) for providing the VC6 and BF2D datasets.

can be addressed by elongated kernels (e.g., Gabor [3]), no *hand-crafted* model for non-straight tubular shapes has been proposed so far. Here we present a novel hand-crafted ridge detector, *SCIRD* (*Scale and Curvature Invariant Ridge Detector*) which is simultaneously rotation, scale and curvature invariant, and removes the assumption of locally straight tubular structures by introducing a *curved-support Gaussian model* (Figure 1). Our tubularity measure is obtained by convolving the image with a filter bank of second-order directional derivatives of such curved-support Gaussians. To cope with irregular, point-like structures, we generate a subsets of directionally elongated kernels, similar to a Gabor filter bank [3]. Experimental results on three challenging datasets [7] show that our approach outperforms state-of-the-art HCF methods on tortuous structures.

2 Methods

Curved-Support Gaussian Models. Our HCF model is inspired by the curved-Gaussian models introduced by Lin and Dayan [10] who analyzed the correlation between currencies in time (exchange rate curvature analysis). Consider a multivariate zero-mean (n -D) Gaussian function with diagonal covariance matrix,

$$G(\boldsymbol{\varphi}; \boldsymbol{\sigma}) = \frac{1}{\sqrt{(2\pi)^n \prod_{i=1}^n \sigma_i^2}} e^{-\sum_{i=1}^n \frac{\varphi_i^2}{2\sigma_i^2}} \quad (1)$$

where $\boldsymbol{\varphi} = (\varphi_1, \varphi_2, \dots, \varphi_n)$ represents a point in the $\{\varphi\}$ coordinate system, and $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_n)$ describes its amount of dispersion around the average (set to 0 without loss of generality). To bend the support of this Gaussian, we consider a volume-preserving non-linear transformation $\mathcal{T} : \mathbb{R}^n \mapsto \mathbb{R}^n$ with $\mathcal{T}(\mathbf{x}) = \boldsymbol{\varphi} = (\varphi_1, \varphi_2, \dots, \varphi_n)$ of the form

$$\varphi_n = x_n + \sum_{i=0}^{n-1} k_{ni} m_{ni}(x_1, x_2, \dots, x_{n-1}) \quad (2)$$

where k_{ni} are weights (see below) and the non-linear functions $m_{ni}(x_1, x_2, \dots, x_{n-1})$ (and their inverses) have continuous partial derivatives [10]. Now, using the change of variable theorem, the non-linear transformation of the normalised $G(\boldsymbol{\varphi}; \boldsymbol{\sigma})$ in the $\{\varphi\}$ coordinate system will be already normalised in the new $\{x\}$ coordinate system since the transformation considered is volume-preserving. In the 2-D case (i.e. $n = 2$), the application of the transformation \mathcal{T} to $G(\boldsymbol{\varphi}; \boldsymbol{\sigma})$ in Eq. (1) leads to the curved-support bivariate Gaussian function:

$$G(x_1, x_2; \boldsymbol{\sigma}, \mathbf{k}) = \frac{1}{\sqrt{4\pi^2 \sigma_1^2 \sigma_2^2}} e^{-\frac{(x_1 + k_{10})^2}{2\sigma_1^2} - \frac{[x_2 + k_{20} m_{20}(x_1) + k_{21} m_{21}(x_1)]^2}{2\sigma_2^2}}. \quad (3)$$

If we now consider *quadratic* non-linear functions $m_{ni} = x_i^2$ for $0 < i < n$ and $m_{n0} = 1$, some of the parameters k_{ni} have the intuitive interpretation of curvatures [10]. In fact, we observe that k_{10} controls the elongation *imbalance*

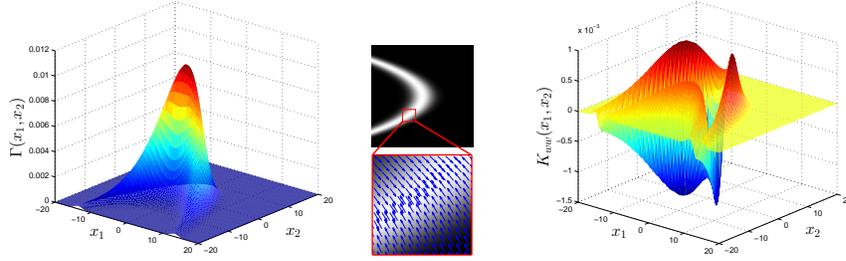


Fig. 1. An example of our “curved” shape model (left), its gradients map (centre) and the derived filter used to detect tortuous and point-like structures (right).

of the shape (i.e. $k_{10} \neq 0$ makes one side tail longer than the other), k_{21} its curvature and k_{20} is simply a translation parameter. We are now in a position to define our shape model:

$$\Gamma(x_1, x_2; \boldsymbol{\sigma}, k) = \underbrace{\frac{1}{\sqrt{2\pi\sigma_1^2}} e^{-\frac{x_1^2}{2\sigma_1^2}}}_{\text{longitudinal}} \underbrace{\frac{1}{\sqrt{2\pi\sigma_2^2}} e^{-\frac{(x_2 + kx_1^2)^2}{2\sigma_2^2}}}_{\text{orthogonal}}. \quad (4)$$

where (x_1, x_2) is a point in the principal component coordinate system of the target structure³, (σ_1, σ_2) control the elongation of the shape (“memory”) and its width, respectively; in fact, the first term of $\Gamma(x_1, x_2; \boldsymbol{\sigma}, k)$ controls the longitudinal Gaussian profile of the model, while the second controls the cross-sectional Gaussian profile. We set $k_{10} = 0$ and $k_{20} = 0$. Importantly, we add a new parameter, k , to control the curvature of the Gaussian support (see Figure 1 - left).

Unsupervised SCIRD. Various detectors of tubular image structures compute the contrast between the regions inside and outside the tube or ridge [2,3,9]. We extend this idea to curved-support Gaussian models by computing the second-order directional derivative in the gradient direction at each pixel.

Let $I(x, y)$ represent the grey-level of a monochrome image at the location (x, y) in image coordinates. The grey-level at this location can be expressed using the first-order gauge coordinate system (\mathbf{v}, \mathbf{w}) , where $\mathbf{w} = \frac{\nabla I(x, y)}{\|\nabla I(x, y)\|}$ and $\mathbf{v} = \mathbf{w}_\perp$. Our “curved tubularity” measure is

$$I_{\mathbf{w}\mathbf{w}}(x, y) = D_{\mathbf{w}}[D_{\mathbf{w}}I(x, y)] = D_{\mathbf{w}}[\mathbf{w}^T \nabla I(x, y)] \triangleq \mathbf{w}^T H_I \mathbf{w}, \quad (5)$$

where $D_{\mathbf{w}}$ is the directional derivative operator along \mathbf{w} .

$$H_I = \begin{bmatrix} I_{xx}(x, y) & I_{xy}(x, y) \\ I_{yx}(x, y) & I_{yy}(x, y) \end{bmatrix} \quad (6)$$

³Notice that this formulation requires a prior knowledge of the structure orientation. Later, we remove this constraint using multiple rotated kernels.

is the Hessian matrix, and $\nabla I(x, y) = [I_x(x, y), I_y(x, y)]^T$ is the gradient of $I(x, y)$. Substituting Eq. (6) in Eq. (5):

$$I_{ww} = \frac{(I_x I_{xx} + I_y I_{yx})I_x + (I_y I_{xy} + I_x I_{yy})I_y}{I_x^2 + I_y^2}. \quad (7)$$

where we have omitted arguments $(x, y; \sigma, k)$ for compactness. However, computing these differential operators separately and then combining them is very expensive; hence we differentiate by convolving with derivatives of curved-support Gaussian, which simultaneously smooths noise [2]. This leads to an efficient tubularity enhancement filter:

$$I_{ww}(x, y; \sigma, k) = I(x, y) * K_{ww}(x, y; \sigma, k), \quad (8)$$

where K_{ww} represents our *tubularity* probe kernel (see example in Figure 1 - right):

$$K_{ww} = \frac{(\tilde{I}_x \Gamma_{xx} + \tilde{I}_y \Gamma_{yx})\tilde{I}_x + (\tilde{I}_y \Gamma_{xy} + \tilde{I}_x \Gamma_{yy})\tilde{I}_y}{\tilde{I}_x^2 + \tilde{I}_y^2}. \quad (9)$$

Here, \tilde{I} is a curved-support Gaussian model (our model) with a constant (i.e. non-Gaussian) longitudinal profile, whose gradient direction is orthogonal to the centreline (Figure 1 - centre). To achieve scale and curvature invariance in the discrete domain, we create a filter bank from multiple kernels generated by making σ_2 and k span scale and curvature range for the specific application at hand. Rotation invariance is obtained augmenting our filter bank with kernel replicas rotated by angles $\theta \in [0, 2\pi)$, i.e., symbolically, $K_{ww}(x, y; \sigma, k, \theta)$. Fragmented (point-like) structures are dealt with by tuning the “memory” parameter σ_1 (see Section 3 for how to tune σ_1). Figure 2(a) shows some of the kernels used in our experiments. Finally, SCIRD selects the maximum response over all kernels. Notice that since we are interested in ridge-like structures, we set $\text{SCIRD}(x, y) = 0$ for all (x, y) locations such that $\text{SCIRD}(x, y) < 0$ [2].

We observed that in some cases, e.g. *in vivo* confocal microscopy, non-linear contrast variations across the image can make interesting structures appear thin and poorly contrasted. In other cases, e.g. bright-field micrographs, such structures may not be interesting and therefore they should not be enhanced. To address this, we introduce a contrast normalization term and a parameter $\alpha \in \mathbb{R}$ that is positive in the former case (enhancement required) and negative in the latter, obtaining

$$\text{SCIRD}(x, y) = \frac{\max_{\sigma, k, \theta} I_{ww}(x, y; \sigma, k, \theta)}{1 + \alpha I^C(x, y)}, \quad (10)$$

where

$$I^C(x, y) = \frac{1}{(N+1)^2} \sum_{i=x-\frac{N}{2}}^{x+\frac{N}{2}} \sum_{j=y-\frac{N}{2}}^{y+\frac{N}{2}} \max_{\sigma_2} \|\nabla I(i, j; \sigma_2)\| \quad (11)$$

is a contrast measure based on multiscale gradient magnitude estimation averaged on a patch $(N+1) \times (N+1)$ around the pixel (x, y) . Notice, N is not a free

parameter, but the width (and height) of the largest filter in the SCIRD filter bank ($N = 8\sigma_{2max}$)⁴. Since SCIRD is designed to respond maximally at the centreline, it can be employed for both curvilinear structure segmentation *and* centreline detection. The former is achieved, e.g., by thresholding directly the *tubularity map* obtained from Eq. (10). The latter is achieved, e.g., by Canny-like non-maxima suppression.

Supervised SCIRD. Segmentation and centreline detection using unsupervised SCIRD can be improved by introducing a supervised classifier. Hence, we combine n feature maps ($I_{ww}^{(i)}$, $i = 1, \dots, n$) obtained using our filter bank with unsupervised SCIRD to form a feature vector \mathbf{f} :

$$\mathbf{f} = \left[\text{SCIRD}, I_{ww}^{(1)}, I_{ww}^{(2)}, \dots, I_{ww}^{(n)} \right]^T. \quad (12)$$

We employ a Random Decision Forest [7, 11] to classify each pixel. Thresholding directly the resulting probability map (which can be seen as a tubularity map) leads indeed to a more accurate and less noisy segmentation (see precision-recall curves in Section 3). Supervised centreline detection is obtained using Canny-like non-maxima suppression on the tubularity map. As local orientation for both supervised and unsupervised centreline detection we choose that of the kernel responding maximally.

3 Experiments and Results

Datasets. We validate SCIRD on 3 datasets including low and high resolution images of corneal nerve fibres and neurons, showing very diverse curvilinear structures as illustrated in Figure 2(b). **IVCM** [1] is a dataset of 100 384×384 confocal microscopy images of the corneal subbasal nerve plexus with different grades of tortuosity. Nerve centrelines were manually traced by a specialist as ground truth. Low resolution, non-uniform and poor contrast, tortuosity, and fibre fragmentation make this dataset particularly challenging. As usually done to evaluate methods extracting onepixel-wide curves [12], we introduce a tolerance factor ρ : a predicted centreline point is considered a true positive if it is at most ρ distant from a ground truth centreline point. In our experiments we set $\rho = 2$. Following the usual benchmarking procedure [7], we average performance measures over 10 random sub-sampling cross-validation runs, using 50 images for training and the rest for testing in each run. The resulting precision-recall curves are reported in Figure 4 (mean and standard deviation of the results from individual runs). The **BF2D** dataset [7] consists of two minimum intensity projections of bright-field micrographs that capture neurons, annotated by an expert. The images have a high resolution (1024×1792 and 768×1792) but a low signal-to-noise ratio because of irregularities in the staining process; the dendrites often appear as point-like (fragmented) structures easily mistaken for

⁴The code is available at <http://staff.computing.dundee.ac.uk/rannunziata/>.

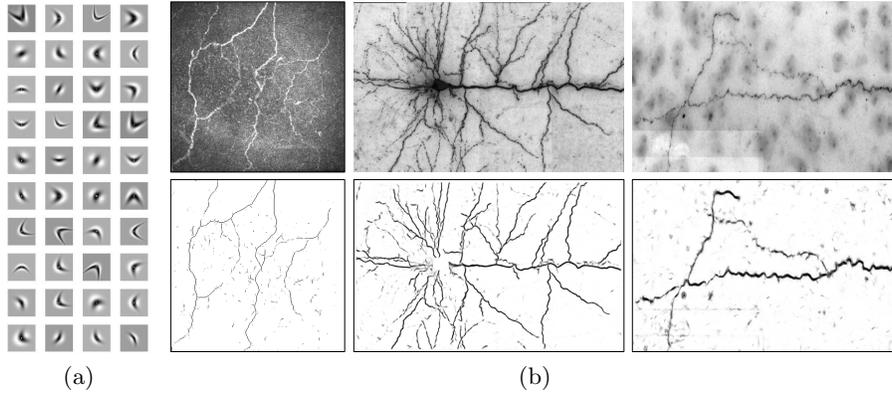


Fig. 2. (a) A subset of SCIRD filters used in our experiments. Notice, the model includes straight support too; (b) original images (top), supervised SCIRD *tubularity* maps (bottom) from IVCm (left), BF2D (centre), VC6 (right) testing sets.

noise. As a consequence, the quality of the annotations themselves is limited [7]. We adopted the same set partition described in [7]. The **VC6** dataset was created by the authors in [7] from a set of publicly available 3D images showing dendritic and axonal subtrees from one neuron in the primary visual cortex. It consists of three images obtained computing minimum intensity projections of a subset of 3-D images, with numerous artifacts and poor contrast, hence challenging for automatic segmentation. We retained two images for training and the third one for testing, adopting the same set partition in [7].

Parameters setting. SCIRD’s key parameters are σ_1 , σ_2 and k controlling the filter memory, width and curvature, respectively. The ranges of σ_1 over datasets were chosen considering the level of fragmentation (i.e. the more fragmented, the higher). We set σ_2 values taking into account the maximum and minimum width of the target structures in each dataset as they depend on resolution; curvature values were set according to the level of tortuosity (or bending) shown by the specific curvilinear structures. The discretisation step of the orientations θ was set to achieve a good trade-off between accuracy and speed. The contrast normalisation parameter α was set as discussed in Section 2. We set the range of θ and k for all the datasets: $\theta = \{\frac{\pi}{12}, \frac{\pi}{6}, \dots, 2\pi\}$, and $k = \{0, 0.025, \dots, 0.1\}$. We set other parameters separately, given the significant difference in resolution between datasets. For the IVCm, $\sigma_1 = \{2, 3, 4\}$, $\sigma_2 = \{2, 3\}$, $\alpha = 1$; for the BF2D dataset, $\sigma_1 = 5$, $\sigma_2 = \{2, 3, 4\}$, $\alpha = -0.075$; for the VC6, $\sigma_1 = 3$, $\sigma_2 = \{2, 3\}$, $\alpha = 0$. For the supervised SCIRD, we used the same range of θ and k , but we doubled the discretisation steps for computational efficiency. We trained Random Forests with the number of decision trees and maximum number of samples in each leaf set using the out-of-bag error. We randomly selected 50,000 pixels from the training set for each dataset. Parameters for SCIRD and baseline methods

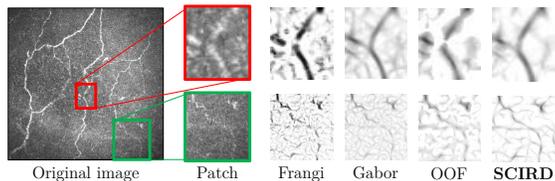


Fig. 3. Qualitative comparison: tubularity estimation results on IVCN. SCIRD (our approach) shows better connectivity and higher signal-to-noise ratio than others.

were tuned separately to achieve their best performance on each dataset to provide a fair comparison.

Results and Discussion. We compare SCIRD against 3 HCF ridge detectors: Frangi [2], based on Hessian matrix eigenvalue analysis, Gabor [3] and the recent Optimally Oriented Flux (OOF) [5], widely acknowledged as excellent tubular structure detectors.

Qualitative (Figure 3) and quantitative (Figure 4) results show that SCIRD outperforms the other methods on all datasets. Specifically, SCIRD shows higher precision from medium to high recall values for the IVCN dataset, suggesting that our filters behave better than others at low resolution and low SNR when dealing with tortuous and fragmented structures. The low number of false positives when false negatives are low, implies that SCIRD selects target structures with higher confidence. Notice that contrast enhancement (“SCIRD $\alpha \neq 0$ ”, $\alpha = 1$) boosts performance on this dataset achieving the level of performance obtained in the supervised setting (“SCIRD, RF”), at a lower computational cost. For the BF2D dataset, unsupervised SCIRD shows a significant improvement from low to high recall values, suggesting that fewer non-target structures are detected and targets are enhanced with higher accuracy (e.g. point-like structures are correctly reconnected, tortuous structure profiles are better preserved). Contrast reduction (“SCIRD $\alpha \neq 0$ ”, $\alpha = -0.075$) proves helpful for these images. Supervised classification improves performance further. For the VC6 dataset, SCIRD shows better performance from low to medium recall values, indicating a better discrimination between curvilinear structures and artifacts, in addition to a more accurate profile segmentation for tortuous structures. Contrast enhancement did not help for this dataset, while supervised classification contributes significantly to improve results. The time to run SCIRD using a single core on Intel i7-4770 CPU @ 3.4 GHz and MATLAB code is 3.84s (IVCN), 2.77s (VC6) and 23s (BF2D). Since SCIRD is highly parallelizable the time to run can be reduced further.

4 Conclusion

We designed a novel HCF ridge detector based on curved-support Gaussians, which is simultaneously invariant to orientation, scale, and, unlike its peers, cur-

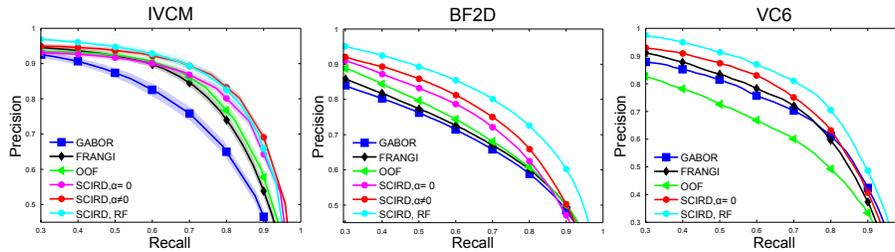


Fig. 4. Precision-recall curves for SCIRD and baselines on IVCN, BF2D and VC6 datasets. Curves are obtained applying different thresholds on the *tubularity* maps.

vature invariant. Experimental results (precision-recall plots) show that SCIRD outperforms current state-of-the-art HCF ridge detectors on 3 challenging datasets, two of which used in the recent literature for similar methods. Our future work will investigate the combination of SCIRD with learned filters in order to capture or discard structures difficult to model (e.g. crossings and artifacts).

References

- Annunziata, R., Kheirkhah, A., Aggarwal, S., Cavalcanti, B.M., Hamrah, P., Trucco, E.: Tortuosity classification of corneal nerves images using a multiple-scale-multiple-window approach. In: OMIA Workshop , MICCAI. (2014)
- Frangi, A., Niessen, W., Vincken, K., Viergever, M.: Multiscale vessel enhancement filtering. In: MICCAI. (1998)
- Soares, J.V., Leandro, J.J., Cesar, R.M., Jelinek, H.F., Cree, M.J.: Retinal vessel segmentation using the 2-d gabor wavelet and supervised classification. IEEE TMI (2006)
- Santamara-Pang, A., Colbert, C., Saggau, P., Kakadiaris, I.: Automatic centerline extraction of irregular tubular structures using probability volumes from multi-photon imaging. In: MICCAI. (2007)
- Law, M.W., Chung, A.C.: Three dimensional curvilinear structure detection using optimally oriented flux. In: ECCV. (2008)
- González, G., Aguet, F., Fleuret, F., Unser, M., Fua, P.: Steerable features for statistical 3d dendrite detection. In: MICCAI. (2009)
- Rigamonti, R., Lepetit, V.: Accurate and efficient linear structure segmentation by leveraging ad hoc features with learned filters. In: MICCAI. (2012)
- Becker, C., Rigamonti, R., Lepetit, V., Fua, P.: Supervised feature learning for curvilinear structure segmentation. In: MICCAI. (2013)
- Hannink, J., Duits, R., Bekkers, E.: Crossing-preserving multi-scale vesselness. In: MICCAI. (2014)
- Lin, J.K., Dayan, P.: Curved gaussian models with application to the modeling of foreign exchange rates. In: Computational Finance, MIT Press (1999)
- Criminisi, A., Shotton, J.: Decision forests for computer vision and medical image analysis. Springer (2013)
- Sironi, A., Lepetit, V., Fua, P.: Multiscale centerline detection by learning a scale-space distance transform. In: CVPR. (2014)