Towards Robust and Efficient Automated Curvilinear Structure Detection in Medical Images

R. Annunziata, E. Trucco
School of Computing,
University of Dundee,
Dundee, UK.

P. Hamrah, A. Kheirkhah
Department of Ophthalmology,
Harvard Medical School,
Boston, USA.

ABSTRACT

Segmenting dendritic trees and corneal nerve fibres is challenging due to their uneven and irregular shape. Our first contribution is a novel ridge detector, SCIRD, which is rotation, scale and curvature invariant [2]. Then, we present a novel feature boosting method combining hand-crafted appearance features with learned context filters. Experimental results on 3 challenging and diverse datasets show that our methods outperform state-of-the-art approaches.

INTRODUCTION

Automated systems analysing curvilinear structures in the medical domain would allow cost-effective large screening programs aimed at early diagnosis. To measure morphometric properties such as tortuosity [1, 6], these structures have to be segmented accurately. Many solutions have been proposed to cope with problems such as low signal to noise ratio at small scales, confounding non-target structures and non-uniform illumination [3, 4, 7, 9]. Most approaches are based on a local *tubularity* measure estimated via hand-crafted features (HCFs) [3, 4, 9], or learned from training data [8]. While HCF methods are fast, they are based on assumptions that are violated by highly fragmented and tortuous structures, i.e. continuous and locally straight tubular shapes. While discontinuity can be addressed by elongated kernels (e.g., Gabor [9]), no *hand-crafted* model for non-straight tubular shapes has been proposed so far. So, our first contribution is a novel HCF ridge detector, SCIRD (Scale and Curvature Invariant Ridge Detector) which is rotation, scale and curvature invariant. Recently, combining HCFs with learned filters has proven successful as it exploits the efficiency of a fast HCF approach to reduce the amount of learned filters [7]. Although this method outperforms state-of-the-art approaches such as [3, 4], our second contribution is motivated by its modest success in segmenting fragmented structures such as corneal nerve fibres and neurites. Integrating *context* information with *appearance* features has been recently found to reduce these shortcomings, by learning multiple discriminative models sequentially [10]. In this paper, we propose a novel approach for combining HCFs (i.e. appearance) with learned context filters which improves results in [7] at the *same* computation cost (i.e. learning a single discriminative model).

METHODS

Curved-Support Gaussian Models

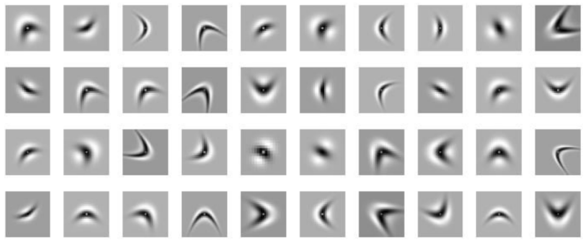

Figure 1: A subset of SCIRD filters.

We model a 2-D curvilinear structure using curved-Gaussian models [5]:

$$\Gamma(x_1, x_2; \boldsymbol{\sigma}, k) = \frac{1}{\sqrt{2\pi\sigma_1^2}} e^{-\frac{x_1^2}{2\sigma_1^2}} \frac{1}{\sqrt{2\pi\sigma_2^2}} e^{-\frac{\left(x_2 + k x_1^2\right)^2}{2\sigma_2^2}}. \quad (1)$$

where (x_1, x_2) is a point in the principal component coordinate system of the target structure, (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. Importantly, we add a new parameter, k, to control the curvature of the Gaussian support.

Scale and Curvature Invariant Ridge Detector

Various detectors of tubular image structures compute the contrast between the regions inside and outside the ridge (e.g., [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. To improve efficiency, we formulate the problem of tubularity estimation (I_{ww}) as a filtering operation:

$$I_{ww}(x, y; \boldsymbol{\sigma}, k) = I(x, y) * K_{ww}(x, y; \boldsymbol{\sigma}, k), \quad (2)$$

where $I(x, y)$ represents the grey-level of a monochrome image at the location (x, y) and

$$K_{ww} = (\tilde{\Gamma}_x \Gamma_{xx} + \tilde{\Gamma}_y \Gamma_{yx})\tilde{\Gamma}_x + (\tilde{\Gamma}_y \Gamma_{xy} + \tilde{\Gamma}_y \Gamma_{yy})\tilde{\Gamma}_y. \quad (3)$$

Here, $\tilde{\Gamma}$ is a curved-support Gaussian model with a constant longitudinal profile. Single subscript indicates first derivative (e.g., $\tilde{\Gamma}_x$), double subscripts is used for second derivatives (e.g., $\tilde{\Gamma}_{xx}$). To achieve scale and curvature invariance, we create a filter bank generated by making σ_2 and k span scale and curvature range for the specific appli-

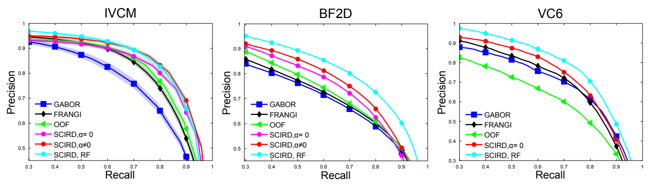

Figure 2: P-R curves for SCIRD and baselines.

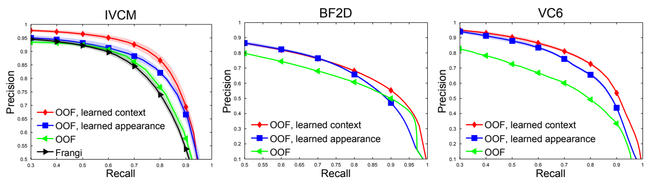

Figure 3: P-R curves for our feature boosting method and baselines.

cation at hand. Rotation invariance is obtained adding rotated filters replicas. Fragmented structures are dealt with by tuning the "memory" parameter σ_1. Figure 1 shows some of the kernels used in our experiments. Finally, SCIRD selects the maximum response over all kernels. Notice, SCIRD can be employed for both segmentation (after thresholding) *and* centreline detection (after non-maxima suppression).

Learning Context Filters

Auto-context [10] is a well-known method to take context information into account. However, it is based on learning multiple discriminative models in a sequential pipeline, making this approach impractical for high resolution images or large datasets. Our goal is to exploit context information without increasing computational cost with respect to the solution in [7]. This is achieved by learning a *single* discriminative model which takes as input both *appearance* (i.e. likelihood computed on the original image [10]) and context information (i.e. relations between objects [10]). We use the Optimally Oriented Flux (OOF) feature [4] as HCF, shown to outperform other HCFs on the datasets we use for our tests in [7]. Then, we learn *context* filters applying the K-means algorithm to the patches of OOF feature maps. Intuitively, we learn a set of probing filters capturing key configurations actually present in the OOF filtered maps. Notice that our context feature learning includes HCF into the filter learning process, while appearance feature learning as proposed in [7] does not. So, while the latter could potentially learn redundant filters, i.e. filters reconstructed by a combination (linear or non-linear) of the HCF already modelling the appearance, our context filters learning method is complementary.

RESULTS AND DISCUSSION

We evaluate our method and baselines performance on 3 very diverse datasets, showing corneal nerve fibres and neurites (IVCM, BF2D and VC6). First, we compare SCIRD against 3 HCF state-of-the-art ridge detectors: Frangi [3], Gabor [9] and Optimally Oriented Flux (OOF) [4]. The results in Figure 2 show that SCIRD outperforms the other methods on all datasets 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.

Second, we compare our method of feature boosting (i.e. context learning) against the method proposed in [7] (i.e. appearance learning). As Figure 3 shows, our method outperforms the baselines on all datasets. In particular, boosting OOF with learned context features outperforms the boosting technique proposed in [7] at the *same* computational cost (i.e. learning a single discriminative model). In terms of efficiency, learning a filter bank of 100 filters using K-means algorithm is significantly faster than sparse coding. In fact, less than 30 seconds are typically required to learn our filter banks which compares favourably with several days reportedly needed to learn 121 filters using sparse coding [7].

CONCLUSION

We addressed the problem of curvilinear structure segmentation by proposing a novel HCF (SCIRD) and a new method combining HCFs with learned context filters. Experimental results show that our methods outperform state-of-the-art approaches. Our future work will investigate the combination of SCIRD with learned context filters. Part of this work (i.e., SCIRD) will be presented at MICCAI 2015 [2].

REFERENCES

[1] R. Annunziata, A. Kheirkhah, S. Aggarwal, B. M. Cavalcanti, P. Hamrah, and E. Trucco. Tortuosity classification of corneal nerves images using a multiple-scale-multiple-window approach. In *OMIA Workshop , MICCAI*, 2014.

[2] R. Annunziata, A. Kheirkhah, P. Hamrah, and E. Trucco. Scale and curvature invariant ridge detector for tortuous and fragmented structures. In *MICCAI*, 2015, (In press).

[3] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. In *MICCAI*. 1998.

[4] M. W. Law and A. C. Chung. Three dimensional curvilinear structure detection using oof. In *ECCV*. 2008.

[5] J. K. Lin and P. Dayan. Curved gaussian models with application to the modeling of foreign exchange rates. In *Computational Finance*. MIT Press, 1999.

[6] A. Lisowska, R. Annunziata, and E. Trucco. An experimental assessment of five indices of retinal vessel tortuosity with the ret-tort public dataset. In *IEEE EMBS*, 2014.

[7] R. Rigamonti and V. Lepetit. Accurate and efficient linear structure segmentation by leveraging ad hoc features with learned filters. In *MICCAI*. 2012.

[8] A. Sironi, V. Lepetit, and P. Fua. Multiscale centerline detection by learning a scale-space distance transform. In *CVPR*, 2014.

[9] J. V. Soares, J. J. Leandro, R. M. Cesar, H. F. Jelinek, and M. J. Cree. Retinal vessel segmentation using the 2-d gabor wavelet and supervised classification. *IEEE TMI*, 2006.

[10] Z. Tu and X. Bai. Auto-context and its application to high-level vision tasks and 3d brain image segmentation. *IEEE TPAMI*, 2010.