

Crossing–Preserving Multi–Scale Vesselness

Julius Hannink, Remco Duits and Erik Bekkers

Eindhoven University of Technology, Department of Biomedical Engineering and
Department of Mathematics and Computer Science, Eindhoven, the Netherlands

J.Hannink@gmx.de, R.Duits@tue.nl and E.J.Bekkers@tue.nl

Abstract. The multi–scale Frangi vesselness filter is an established tool in (retinal) vascular imaging. However, it cannot properly cope with crossings or bifurcations since it only looks for elongated structures. Therefore, we disentangle crossings/bifurcations via (multiple scale) invertible orientation scores and apply vesselness filters in this domain. This new method via scale–orientation scores performs considerably better at enhancing vessels throughout crossings and bifurcations than the Frangi version. Both methods are evaluated on a public dataset. Performance is measured by comparing ground truth data to the segmentation results obtained by basic thresholding and morphological component analysis of the filtered images.

Keywords: Multi–scale vesselness filters, continuous wavelet transforms, line detection, gauge frames, retinal imaging

1 Introduction

The retinal vasculature enables non–invasive observation of the human circulatory system. A variety of eye–related and systematic diseases such as glaucoma, age–related macular degeneration, diabetes, hypertension, arteriosclerosis or Alzheimer’s disease affect the vasculature and may cause functional or geometric changes [1]. Automated quantification of these defects promises massive screenings for systematic and eye–related vascular diseases on the basis of fast and inexpensive imaging modalities, i.e., retinal photography.

To automatically assess the state of the retinal vascular tree, vessel segmentations and/or models have to be created and analyzed. Because retinal images usually suffer from low contrast at small scales, the vasculature needs to be enhanced prior to model creation/segmentation. One well–established approach is the Frangi vesselness filter [2]. It is frequently used in robust retinal vessel segmentation methods [3, 4]. However, the Frangi filter has a known drawback. It cannot properly enhance vessels throughout crossings or bifurcations that make up huge parts of the retinal vascular network. To generically deal with this issue, we apply the principle of image processing via invertible orientation scores (Fig. 1) and derive a multi–scale crossing–preserving vesselness filter.

The construction of invertible orientation scores is inspired by the functional architecture of the human cortical columns in the primary visual cortex [5],

where decomposition of local orientation allows the visual system to separate crossing/bifurcating structures. Similar approaches rely on other Lie-groups, see [6, 7]. Here, we will develop vesselness filters on the extended Lie-group domains of the rotation translation group $SE(2)$ and the rotation, translation and scaling group $SIM(2)$. In the $SIM(2)$ case we arrive at continuous wavelet transforms on $SIM(2)$ [8]. In the $SE(2)$ case our approach is closely related (see [9, ch:2.3]) to the work by Krause et al. [10] who rely on a local Radon transform.

The general idea is as follows. Frangi’s vesselness relies on a Hessian in the image domain and it only copes with one orientation per location. Since complex structures in retinal images exhibit multiple orientations per position, invertible (multiple scale) orientation scores provide a generic disentanglement of all orientations and scales without tampering of data-evidence, see Fig. 2. The orientation score carries a curved (sub)-Riemannian geometry. When moving in θ -direction, spatial derivatives must rotate accordingly. We employ differential frames in the score, providing us Hessians, and subsequent vesselness filters that cope generically with (multiple-scale) crossings/bifurcations without having to classify them.

In the end, we show the performance of this new type of vesselness filters by comparison to the multi-scale Frangi vesselness [2], both qualitatively and quantitatively on the High Resolution Fundus (HRF) image dataset available at <http://www5.cs.fau.de/research/data/fundus-images/>.

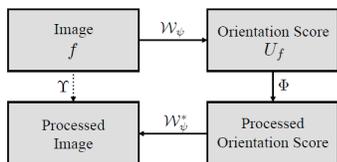


Fig. 1: Image processing via invertible orientation scores.

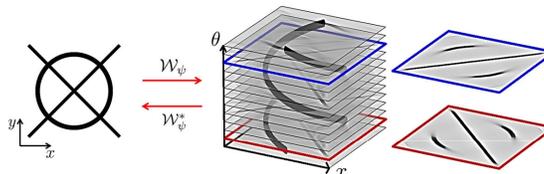


Fig. 2: Exemplary image and corresponding orientation score.

2 Methods

An orientation score $U_f : SE(2) \rightarrow \mathbb{C}$ is obtained by correlating an input image f with a specially designed anisotropic wavelet ψ :

$$U_f(\mathbf{x}, \theta) = (\overline{\psi_\theta} \star f)(\mathbf{x}, \theta) = \int_{\mathbb{R}^2} \overline{\psi(R_\theta^{-1}(\mathbf{y} - \mathbf{x}))} f(\mathbf{y}) d\mathbf{y} \quad (1)$$

where R_θ denotes a 2d counter-clockwise rotation matrix and $\psi_\theta(\mathbf{x}) = \psi(R_\theta^{-1}\mathbf{x})$. Exact image reconstruction is achieved by

$$f = \mathcal{F}^{-1} \left[M_\psi^{-1} \mathcal{F} \left[\mathbf{x} \mapsto \frac{1}{2\pi} \int_0^{2\pi} (\psi_\theta \star U_f)(\mathbf{x}, \theta) d\theta \right] \right] \quad (2)$$

where $*$ denotes convolution, $\mathcal{F}[\cdot]$ represents the unitary Fourier transform on $\mathbb{L}_2(\mathbb{R}^2)$ and M_ψ is given by $\int_0^{2\pi} |\mathcal{F}[\psi_\theta]|^2 d\theta$. Theoretically, reconstruction is well posed for $0 < \delta < M_\psi < \infty$ with arbitrary small δ . One type of wavelets that meet the stability criteria in [11, 5, 12] are cake wavelets described by [13, 9]. They uniformly cover the Fourier domain up to a radius of about the Nyquist frequency ρ_n to satisfy the discrete version of $M_\psi \approx 1$.

2.1 Gaussian Derivatives in Orientation Scores

The orientation score domain is essentially the 2d special Euclidean motion group $SE(2) \equiv \mathbb{R}^2 \rtimes S^1$ [13, ch.2]. Because of this, all operations Φ on this domain (see Fig. 1) have to be left-invariant to produce a Euclidean invariant net operator \mathcal{Y} on the image [13, ch.2]. This is desirable since the result should be independent of rotation and translation of the input. Φ is left-invariant iff $\Phi \circ \mathcal{L}_g = \mathcal{L}_g \circ \Phi$ for all $g = (\mathbf{x}, \theta) \in SE(2)$ with the shift-twist operator on the score given by

$$\mathcal{L}_g U_f(\mathbf{h}) = U_f(g^{-1}\mathbf{h}) = U_f(R_\theta^{-1}(\mathbf{x}' - \mathbf{x}), \theta' - \theta), \quad (3)$$

for all $g = (\mathbf{x}, \theta), \mathbf{h} = (\mathbf{x}', \theta') \in SE(2)$. So we must rely on left-invariant derivatives, attached to each $(x, y, \theta) \in SE(2)$ in the domain of the score, given by

$$\partial_\xi := \cos \theta \partial_x + \sin \theta \partial_y, \partial_\eta := -\sin \theta \partial_x + \cos \theta \partial_y, \text{ and } \partial_\theta, \quad (4)$$

when constructing vesselness filters on $SE(2)$. These derivatives provide a moving frame of reference on the group steered by the orientation of the wavelet. Their non-zero commutators are given by $[\partial_\theta, \partial_\xi] = \partial_\eta$ and $[\partial_\theta, \partial_\eta] = -\partial_\xi$. Later we will adapt this frame locally to the score, following the theory of best exponential curve fits presented in [13, ch.6]. This will compensate for the fact that our wavelet kernel is not always perfectly aligned with all local orientations present in the image (For details, see [13, ch.6]).

Since orientation and spatial direction have different physical units, a conversion factor is needed. This stiffness parameter β has unit 1/length. It determines the shape of both *exponential curves* [13, ch.6] (i.e. the auto-parallel curves, see [14, app.C]) and the *geodesics* (i.e. shortest distance curves) in $SE(2)$. Intuitively, increase of β makes it cheaper to bend curves, whereas decrease of β makes it cheaper to stretch curves. See e.g. [14, Fig.17]. Mathematically, β appears as the only free parameter in the (sub-)Riemannian metric on $SE(2)$ given by

$$d(\mathbf{g}_1, \mathbf{g}_2) = \inf_{\substack{\gamma(0)=\mathbf{g}_1 \\ \gamma(l)=\mathbf{g}_2 \\ \gamma \in \Delta, l \geq 0}} \int_0^l \sqrt{\mathcal{G}_\beta|_{\gamma(s)}(\dot{\gamma}(s), \dot{\gamma}(s))} ds, \quad (5)$$

with $\Delta = \text{span}\{\partial_\xi, \partial_\eta, \partial_\theta\}$, $\gamma(s) = (\mathbf{x}(s), \theta(s))$ and $\mathcal{G}_\beta|_\gamma(\dot{\gamma}, \dot{\gamma}) = \beta^2(\dot{x} \cos \theta + \dot{y} \sin \theta)^2 + \beta^2(-\dot{x} \sin \theta + \dot{y} \cos \theta)^2 + \dot{\theta}^2$ in the Riemannian case. In the sub-Riemannian case, the allowed part of the tangent space is $\Delta = \text{span}\{\partial_\xi, \partial_\theta\}$. The functional in (5) then reduces to $\int_0^l \sqrt{\kappa^2(s) + \beta^2} ds$ with the curvature κ of the spatially projected curve $\mathbf{x}(s) = \mathcal{P}_{\mathbb{R}^2}\gamma(s)$, for details see [14]).

In order to extract local features in $SE(2)$ with well-posed, left-invariant derivative operators, regularization is needed. The only left-invariant diffusion regularization in $SE(2)$ that preserves the non-commutative group structure via the commutators is elliptic diffusion, isotropic w.r.t. the β -metric (5). Then regularization is achieved via a spatially isotropic Gaussian with scale $\frac{1}{2}\sigma_s^2$ and a 1d-Gaussian in θ with scale $\frac{1}{2}(\beta\sigma_s)^2$ [13, ch.5]. The regularized derivative operators are convolutions with differentiated β -isotropic Gaussians and generalize the concept of Gaussian derivatives used in the Frangi vesselness filter [2] to $SE(2)$. In our subsequent extension to $SIM(2)$, where we include scaling $a > 0$, we will choose $\beta = 0.05/a$ as we must take into account both physical dimensions, and typical tortuosities of retinal vessels.

2.2 Scale–Orientation Scores

To make the cake kernels described by [13, 9] scale-selective, the pieces of cake have to be further divided. By cutting out pieces in the log-radial direction, they are made sensitive to a specific frequency range that can be identified with a scale a in the spatial domain. To construct scale-selective cake kernels (anisotropic wavelets), [12] uses a radial envelope function

$$B^{\text{MS}}(\rho) = \sum_{l=0}^{N_\rho-1} B_l^k(\rho) := \sum_{l=0}^{N_\rho-1} B^k\left(\frac{\log(\rho a_l)}{s_\rho}\right), \text{ with } a_l = a_0 e^{ls_\rho}, \quad (6)$$

where $B^k(x)$ is the k -th order B-spline function, N_ρ is the total number of scales to sample in the Fourier domain and $s_\rho > 0$ denotes the stepsize in log-scale. The multiplicative character in $a_l = a_0 e^{ls_\rho}$ reflects the typical scale transitions at bifurcations in bloodvessels. Because of the B-spline approach, the scale-selective envelopes $B_l^k(\rho)$ sum to 1 and the $M_\psi \approx 1$ requirement is met (Fig. 3). Scale layers outside a spatially defined range of interest are merged to reduce computational load (Fig. 3). We propose the following multi-scale cake kernel

$$\psi^{\text{MS}}(\mathbf{x}) = \left(\mathcal{F}^{-1} \left[M^{-1} \mathcal{F} \left[\tilde{\psi}^{\text{MS}} \right] (\boldsymbol{\omega}) \right] \right) (\mathbf{x}) \quad (7)$$

where $\tilde{\psi}^{\text{MS}}(\mathbf{x})$ denotes the wavelet

$$\tilde{\psi}^{\text{MS}}(\mathbf{x}) = \left(\mathcal{F}^{-1} \left[\boldsymbol{\omega} \mapsto A(\varphi) B_0^k(\rho) \right] \right) (\mathbf{x}) G_{s_x, s_y}(\mathbf{x}) \quad (8)$$

at the finest scale a_0 . The function $A(\varphi)$ is given by $B^k(s_\theta^{-1}[(\varphi \bmod 2\pi) - \pi/2])$ for $\rho > 0$ and $1/N_\theta$ for $\rho = 0$ with the angular stepsize $s_\theta = 2\pi/N_\theta$. The anisotropic Gaussian window $G_{s_x, s_y}(\mathbf{x})$ reduces long tails along the orientation of the wavelet and suppresses oscillations perpendicular to it induced by narrow sampling bandwidths in $B^{\text{MS}}(\rho)$. Changes in the Fourier domain are resolved via normalization with $M(\boldsymbol{\omega}) = N_\rho^{-1} N_\theta^{-1} \sum_{l=0}^{N_\rho-1} \sum_{j=1}^{N_\theta} a_l^{-1} |\mathcal{F}[\tilde{\psi}^{\text{MS}}](a_l R_{\theta_j}^{-1} \boldsymbol{\omega})|$. Thereby, approximative reconstruction is done by summation over scales and angles. The continuous wavelet transform $(\mathcal{W}_\psi f)(\mathbf{x}, \theta, a) = (\overline{\psi}_\theta^a \star f)(\mathbf{x})$ with $\psi_\theta^a(\mathbf{x}) = a^{-1} \psi^{\text{MS}}(a^{-1} R_\theta^{-1} \mathbf{x})$ is now set and processing can begin.

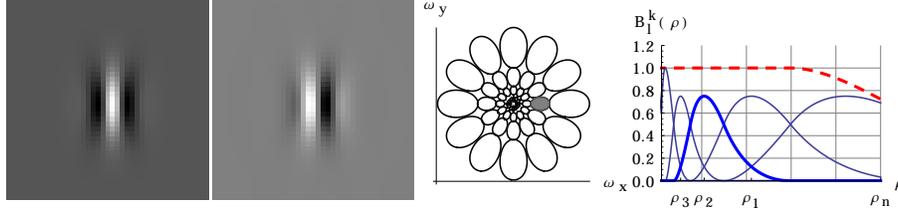


Fig. 3: Real and imaginary part of the multi-scale cake kernel at a_2 (left), Fourier contours of all wavelets (at 70% of the maximum) and $B^{MS}(\rho)$ (red, dashed) with components $B_l^k(\rho)$ (blue).

2.3 Vesselness Filtering on Scale–Orientation Scores

The single-scale Frangi vesselness filter $\mathcal{V}_0^{\text{Fr}}$ makes use of an anisotropy measure \mathcal{R} and a structure measure \mathcal{S} based on second order derivatives in a coordinate system aligned with elongated structures [2]. This approach is now generalized to (scale-)orientation scores for crossing-preserving vesselness filtering. Given a convexity criterion $\mathcal{Q} > 0$ on transversal vessel profiles and the measures \mathcal{R} resp. \mathcal{S} , the setup for the $SE(2)$ vesselness expression $\mathcal{V}_0^{SE(2)}(\mathcal{U}_f^a) : SE(2) \rightarrow \mathbb{R}^+$ is identical to the one proposed by [2]:

$$\mathcal{V}_0^{SE(2)}(\mathcal{U}_f^a) = \begin{cases} 0 & \text{if } \mathcal{Q} \leq 0 \\ \exp\left(-\frac{\mathcal{R}^2}{2\sigma_1^2}\right) \left[1 - \exp\left(-\frac{\mathcal{S}}{2\sigma_2^2}\right)\right] & \text{if } \mathcal{Q} > 0 \end{cases} \quad (9)$$

where $\mathcal{U}_f^a(\mathbf{x}, \theta) = (\mathcal{W}_\psi f)(\mathbf{x}, \theta, a)$, $a > 0$ fixed, is a single scale layer of the multi-scale wavelet transform. We always set $\sigma_1 = 0.5$ and $\sigma_2 = 0.2 \|\mathcal{S}\|_\infty$.

There are two natural generalizations of $\mathcal{V}_0^{\text{Fr}}$ to $SE(2)$ that differ in the choice of coordinate system used to define \mathcal{R}, \mathcal{S} and \mathcal{Q} . One option is to work in the moving frame of reference $\{\partial_\xi, \partial_\eta, \partial_\theta\}$, recall Eq. (4). Second order Gaussian derivatives along the directions ($\xi = x \cos \theta + y \sin \theta, \eta = -x \sin \theta + y \cos \theta$) carry the same information as the eigenvalues of the Hessian in $\mathcal{V}_0^{\text{Fr}}$ [2], and we set

$$\mathcal{R} = \frac{(\partial_\xi^2 \mathcal{U}_f^a)^{s,\beta}}{(\partial_\eta^2 \mathcal{U}_f^a)^{s,\beta}} \quad , \quad \mathcal{S} = [(\partial_\xi^2 \mathcal{U}_f^a)^{s,\beta}]^2 + [(\partial_\eta^2 \mathcal{U}_f^a)^{s,\beta}]^2 \quad , \quad \mathcal{Q} = (\partial_\eta^2 \mathcal{U}_f^a)^{s,\beta} \quad (10)$$

where the superscripts s,β indicate Gaussian derivatives at spatial scale $s = \frac{1}{2}\sigma_s^2$ and angular scale $\frac{1}{2}(\beta\sigma_s)^2$. The generalization of the filter in the $\{\partial_\xi, \partial_\eta, \partial_\theta\}$ frame is referred to as $\mathcal{V}_0^{\xi,\eta}$ at single scales and as $\mathcal{V}^{\xi,\eta}$ in the multiple scale recombination, similar to the notation in [2]. The other possible coordinate system is the Gauge frame $\{\partial_{\mathbf{a}}, \partial_{\mathbf{b}}, \partial_{\mathbf{c}}\}$ determined by the eigendirections of the 3D-Hessian $(\mathcal{H}^{s,\beta} \mathcal{U}_f^a)(\mathbf{g})$, given in [13, ch:5.3], computed with scale s at $\mathbf{g} \in SE(2)$, normalized w.r.t the β -metric (5). As the filter is no longer confined to fixed θ -slices, and the Gauge frame is fully aligned with the 3D-line structures in the score, the analogies to [2] are even stronger in this frame. Given the eigenvalues

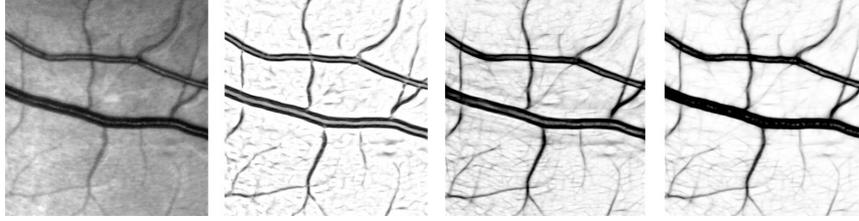


Fig. 4: Retinal image f and multi-scale vesseness filtering results for the Frangi filter $\mathcal{V}^{\text{Fr}}(f)$ and our two methods $\mathcal{V}^{\xi\eta}(f)$ resp. $\mathcal{V}^{a,b,c}(f)$ (left to right).

of the Hessian $(\mathcal{H}^{s,\beta} U_f^a)(\mathbf{g})$ ordered $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$, $\mathcal{R}, \mathcal{S}, \mathcal{Q}$ are computed as

$$\mathcal{R} = \frac{\lambda_1}{c} \quad , \quad \mathcal{S} = \lambda_1^2 + c^2 \quad , \quad \mathcal{Q} = c \quad (11) \quad \text{🗨️}$$

with $c = \frac{1}{2}(\lambda_2 + \lambda_3)$. As such, c is comparable to the orientation confidence defined by [13, ch.5]. The $SE(2)$ -generalization of the vesseness filter in this frame is referred to as $\mathcal{V}_0^{a,b,c}$, whereas the multi-scale $SIM(2)$ -generalization is written as $\mathcal{V}^{a,b,c}$. In the $SIM(2)$ -generalizations of the vesseness filters (regardless the choice of reference frame) the final image reconstruction from vesseness filtered scale-orientation scores are obtained via

$$(\mathcal{V}^{SIM(2)}(f))(\mathbf{x}) = \mu_\infty^{-1} \sum_{l=1}^{N_s} \mu_{l,\infty}^{-1} \sum_{j=1}^{N_\theta} (\mathcal{V}_0^{SE(2)}(U_f^{a_l}))(\mathbf{x}, \theta_j) \quad (12)$$

where μ_∞ and $\mu_{l,\infty}$ are the maximum values, i.e. $\|\cdot\|_\infty$ norms, taken over the subsequent sums. For comparison, the multi-scale Frangi vesseness filter is also computed via summation over single scale results and max-normalized.

Fig. 4 shows multi-scale vesseness filtering results for a retinal image f obtained with the Frangi filter and our two methods for five scales $\{1.5, 2.4, 3.8, 6.0, 9.5\}$ px, $\beta = 0.05/a$ and 12 orientations. Both our methods clearly outperform the Frangi filter at crossings and bifurcations. The Gauge-frame method $\mathcal{V}^{a,b,c}$ gives best results as it aligns with 3D-elongated structures in the score.

3 Experiments

To show the benefit of crossing-preservation in multiple scale vesseness filtering, we devised a simple segmentation algorithm to turn a vesseness filtered image $\mathcal{V}(f)$ into a binary vessel map. First, a local thresholding is applied so that we obtain the binary image $f_B = \Theta([\mathcal{V}(f) - G_\gamma * \mathcal{V}(f)] - t)$ where Θ is the Heaviside step function and G_γ is a Gaussian with standard deviation $\gamma \gg 1$. In a second step, the connected morphological components in f_B counting less than τ pixels or showing elongations below a threshold ν are removed. The parameters γ, τ and ν are fixed at 100 px, 500 px and 0.85 respectively. $\mathcal{V}(f)$ is either obtained with \mathcal{V}^{Fr} or via the $SIM(2)$ method $\mathcal{V}^{a,b,c}$ using the settings mentioned earlier.

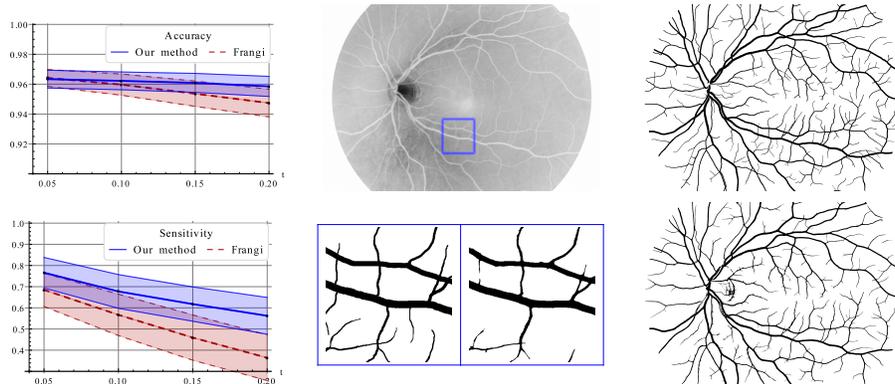


Fig. 5: Mean accuracy and sensitivity on the HRF dataset over threshold values t . Shading shows $\pm 1\sigma$ (left). Retinal image and patch ground truth/segmentation at $t = 0.05$ (center). Full ground truth and segmentation (right).

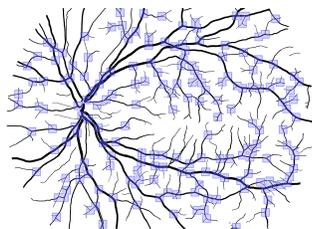


Fig. 6: Shaded regions specify areas containing junction points. Junction points are found by thinning and pruning the binary ground truth.

Method	Junctions		Non-junctions	
	Sensitivity	Accuracy	Sensitivity	Accuracy
Our method				
All	0.84	0.92	0.76	0.97
Healthy	0.82	0.94	0.75	0.97
Glaucoma	0.84	0.93	0.74	0.97
Diabetic	0.85	0.90	0.78	0.96
Frangi [2]				
All	0.71	0.92	0.67	0.97
Healthy	0.71	0.93	0.67	0.97
Glaucoma	0.75	0.93	0.68	0.97
Diabetic	0.69	0.91	0.65	0.97

Table 1: Results for bifurcation and crossing areas (Junctions) compared to the complement of this set (Non-junctions). See Fig. 6 for sample areas.

This segmentation algorithm is evaluated on the HRF dataset consisting of fundus images for a healthy, diabetic retinopathy and glaucoma group (15 images each, ground truths provided). Average sensitivity and accuracy on the whole dataset are shown in Fig. 5 over threshold values t . Our method via invertible scale-orientation scores performs considerably better than the method based on the multi-scale Frangi filter. The segmentation results obtained with $\mathcal{V}^{a,b,c}$ are more stable w.r.t variations in the threshold t and the performance on the small vasculature has improved as measured via the sensitivity. Average sensitivity, specificity and accuracy at a threshold $t = 0.05$ resp. given by 0.786, 0.988, 0.969 (healthy), 0.811, 0.963, 0.953 (diabetic retinopathy) and 0.797, 0.976, 0.964 (glaucoma) compare well with other algorithms evaluated on the HRF dataset (see [3, Tab. 5]). On the diabetic retinopathy group, our method even outperforms existing segmentation methods. Fig. 5 shows a full segmentation computed with

the proposed method and an in-detail patch. In Table 1 and Fig. 6 we see that our method improves sensitivity both at non-crossing line structures (due to line propagation of anisotropic wavelets) and at crossing/bifurcating structures. As expected, we observe a larger improvement at crossings.

4 Discussion

We developed (multi-scale) crossing-preserving vesselness filters as generalizations of [2] to the extended Lie-group domains $SE(2)$ resp. $SIM(2)$. The new filters were evaluated qualitatively and quantitatively on a public dataset and outperformed the Frangi filter and existing segmentation methods. This shows the method’s potential for application in other areas of vascular imaging. Future work includes concatenation with enhancements [13, 12] and tracking [9].

Acknowledgements: The research leading to these results has received funding from the ERC council under the EC’s 7th Framework Programme (FP7/2007–2013) / ERC grant agr. No. 335555.

References

- [1] Ikram, M.K., Ong, Y.T., Cheung, C.Y., Wong, T.Y.: Retinal Vascular Caliber Measurements: Clinical Significance, Current Knowledge and Future Perspectives. *Ophthalmologica* **229**(3) (2013) 125–136
- [2] Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. *LNCS* **1496** (1998) 130–137
- [3] Budai, A., Bock, R., Maier, A., Hornegger, J., Michelson, G.: Robust Vessel Segmentation in Fundus Images. *IJBI* **2013** (2013)
- [4] Lupascu, C.A., Tegolo, D., Trucco, E.: FABC: Retinal Vessel Segmentation Using AdaBoost. *IEEE T-ITB* **14**(5) (2010) 1267–1274
- [5] Duits, R., Felsberg, M., Granlund, G., ter Haar Romeny, B.: Image Analysis and Reconstruction using a Wavelet Transform Constructed from a Reducible Representation of the Euclidean Motion Group. *IJCV* **72**(1) (2007) 79–102
- [6] Duits, R., Janssen, B., Bruurmijn, M., Florack, L., Van Assen, H.: Evolution Equations on Gabor Transforms and their Applications. *ACHA* (to appear 2014)
- [7] Barbieri, D., Citti, G., Cocci, G., Sarti, A.: A cortical-inspired geometry for contour perception and motion integration. *arXiv preprint arXiv:1301.3433* (2013)
- [8] Jacques, L., Antoine, J.P.: Multiselective pyramidal decomposition of images: wavelets with adaptive angular selectivity. *International Journal of Wavelets, Multiresolution and Information Processing* **5** **05** (2007) 785–814
- [9] Bekkers, E., Duits, R., Berendschot, T., ter Haar Romeny, B.: A Multi-Orientation Analysis Approach to Retinal Vessel Tracking. *JMIV* (2014) 1–28
- [10] Krause, M., Alles, R.M., Burgeth, B., Weickert, J.: Fast retinal vessel analysis. *JRTIP* (2013) 1–10
- [11] Fuehr, H.: *Abstract Harmonic Analysis of Continuous Wavelet Transforms*. Springer (2005)
- [12] Sharma, U., Duits, R.: Left-invariant evolutions of wavelet transforms on the Similitude Group. *arXiv preprint arXiv:1306.1800* (2013)
- [13] Franken, E.: *Enhancement of Crossing Elongated Structures in Images*. PhD thesis, Technical University Eindhoven (2008)
- [14] Duits, R., Boscaïn, U., Rossi, F., Sachkov, Y.: Association Fields via Cuspless Sub-Riemannian Geodesics in $SE(2)$. *JMIV* **1** (2013) 32