RETINAL VESSEL DETECTION USING SELF-MATCHED FILTERING

Nai-Xiang Lian, Vitali Zagorodnov, and Yap-Peng Tan

Nanyang Technological University
50 Nanyang Avenue, Singapore 639798, Singapore

ABSTRACT

Automated analysis of retinal images usually requires estimating the positions of blood vessels, which contain important features for image alignment and abnormality detection. Matched filtering can produce the best results but is difficult to implement because the vessel orientations and widths are unknown beforehand. Many researchers use Hessian filtering, which provides an estimate for vessel orientation through the use of three orientation templates. We propose a novel filtering approach, called self-matched filtering, which is based on the 180° rotated version of the noisy vessel signal in the local neighborhood. We show that even though the proposed filter achieves half the signal-to-noise ratio of a matched filter, it does not require the estimation of the vessel scale and orientation, and can outperform Hessian filtering by up to a factor of two in terms of miss detection error.

***Index Terms*—** Medical image processing, Medical signal detection, Hessian matrices, Matched filters,

1. INTRODUCTION

Detecting temporal retinal changes can assist in diagnosis of various eye diseases, such as diabetic retinopathy and age-related macular degeneration [1]. Prior to the change detection, the images need to be aligned to compensate for possible eye motion. The features for such an alignment usually come from the locations of retinal blood vessels [2]. Hence a substantial amount of recent work on both blood vessel detection (vessel tracing) [3,4] and blood vessel segmentation [5].

Detecting vessel locations is similar to detecting a signal in the presence of noise and hence can be optimally solved by using a matched filter [5], yielding responses with the best signal-to-noise ratio (SNR) [6]. However, matched filtering requires precise knowledge of the blood vessel shape parameters (thickness,orientation, shape profile), which substantially vary across the image. Trying all possibilities is computationally too demanding and several approaches have been proposed to reduce the complexity. For example, the search for the vessel orientation can be sped up through the use of steerable filtering [7], which allows computation of filter response in any direction from the output of the filters at several fixed orientations. The most popular steerable filter used in vessel detection is Hessian filter, which approximates the vessel shape via three second derivatives of Gaussian [4, 8, 9].

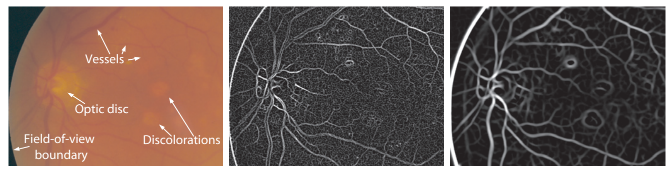

(a) (b) (c)

Fig. 1. A retinal image and the responses of Hessian filtering. (a) Retinal image sample (b) Small scale filtering (c) Large scale filtering

A simple solution to the unknown thickness problem is to choose some effective scale that works sufficiently well for all vessel widths. However, in this case small-scale filters might yield two large responses around wide vessel edges and low response for vessel centerline location, see Figure 1(b). On the other hand, large-scale filters perform poorly on narrow vessels, as shown in Figure 1(c). In addition, the narrow vessels typically have a lower contrast compared to the wide vessels. Hence using large-scale filters to detect them makes their detection even more difficult. Some researchers proposed using multiscale Hessian filtering to improve the performance [4, 9]. However, the downsides of this approach include possible overfitting to noise and the computational overhead due to repeated filtering for several scales.

Another challenge to vessel detection comes from a variety of distracting structures often present in retinal images, such as the field-of-view boundary, optic disc, discolorations, etc., as shown in Figure 1(a). These structures could cause large filtering responses and hence a wrong classification as vessels [see Figures 1(b) and 1(c)]. Removal of such false detections often requires further post-processing, such as quadrature filtering [7] or verification of two opposite-sign step edges in the vicinity of the vessel center [4], which could be computationally expensive.

We propose to improve and simplify the detection of retinal vessels by using a novel filter we call the "self-matched filter". A self-matched filter is a location-dependent template that is obtained by 180° rotation of the noisy vessel signal in the local neighborhood. Since a vessel is locally similar to a straight line, it is rotationally symmetric around any of its centerline points. Hence the so-obtained template will approximately coincide with the original signal regardless of the vessel width, orientation, and even shape profile, making it an excellent approximation to the ideal matched filter. However, due to the presence of noise, the SNR performance of such a filter is reduced by a factor of two compared with the ideal matched filter, as we shall show in Section 2. Nevertheless, the proposed filter yields a more favorable overall performance compared to fixed-scale Hessian filter, as detailed in Section 4. Moreover, the self-matched filter gives a negative response around step edges, preventing the appearance of field-of-view boundary, optic disc and discolorations in the resultant segmentation (Section 2).

2. SELF-MATCHED FILTER

Let $\tilde{s}(\boldsymbol{x}) = s(\boldsymbol{x}) + n(\boldsymbol{x}), \boldsymbol{x} = [x, y]^T \in \mathcal{D}$ be a noisy local image patch, where $s(\boldsymbol{x})$ is the intensity of a feature (such as a vessel) to be detected and $n(\boldsymbol{x})$ is a zero-mean i.i.d. noise. Let $f(\boldsymbol{x})$ be the matched filter defined on $\mathcal{D}$, then the output of such filter at the feature (vessel) center $\boldsymbol{x}_0$ is $r(\boldsymbol{x}_0) = \int_{\mathcal{D}} \tilde{s}(\boldsymbol{x}+\boldsymbol{x}_0) f(\boldsymbol{x}) d\boldsymbol{x}$. Without loss of generality we assume $\int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x} = 0$. (If this is not originally the case, the property can be approximately satisfied by subtracting $\int_{\mathcal{D}} \tilde{s}(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}$ from $\tilde{s}(\boldsymbol{x}+\boldsymbol{x}_0)$, $\forall \boldsymbol{x} \in \mathcal{D}$.) It follows that the optimal matched filter has the same shape as $s(\boldsymbol{x})$:

$$f(\boldsymbol{x}) = s(\boldsymbol{x}+\boldsymbol{x}_0), \qquad \forall \boldsymbol{x} \in \mathcal{D} \tag{1}$$

Let $\xi^2(\boldsymbol{x}_0) = \int_{\mathcal{D}} s^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}$ and $\sigma_n^2 = E\left\{n^2(\boldsymbol{x})\right\}$. The matched filter achieves the best possible SNR given as

$$\mathbf{SNR}_{opt}(\boldsymbol{x}_0) = \frac{\left[\int_{\mathcal{D}} s^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}\right]^2}{E\left\{\left[\int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)n(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}\right]^2\right\}} = \frac{\xi^2(\boldsymbol{x}_0)}{\sigma_n^2} \tag{2}$$

However, the optimal matched filter is often unattainable, as we can only observe $\tilde{s}(\boldsymbol{x})$ and not $s(\boldsymbol{x})$. Using an approximation $f(\boldsymbol{x}) = \hat{s}(\boldsymbol{x}+\boldsymbol{x}_0)$ will lead to an SNR smaller than (2); the exact reduction depends on the approximation and hence is difficult to predict.

Consider a pseudo-matched filter whose shape is the same as the noisy version of the signal, $\tilde{s}(\boldsymbol{x})$:

$$f(\boldsymbol{x}) = \tilde{s}(\boldsymbol{x}+\boldsymbol{x}_0), \qquad \forall \boldsymbol{x} \in \mathcal{D} \tag{3}$$

According to the following lemma, when the noise is not very large (compared to the signal energy), such a 'naive' filter can achieve a quarter of the best possible SNR.

Lemma 1 *If $\sigma_s^2 \gg \sigma_n^2$, the SNR of the pseudo-matched filtering (3) is*

$$\mathbf{SNR}_{pseudo} = \frac{1}{4}\mathbf{SNR}_{opt} \tag{4}$$

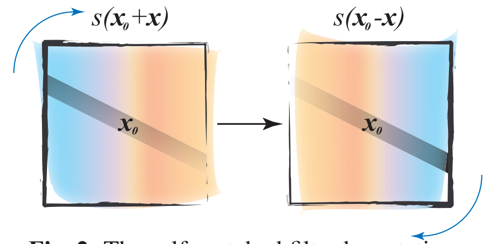

Fig. 2. The self-matched filter by rotating.

Proof: The response of the pseudo-matched filter is

$$r(\boldsymbol{x}_0) = \int_{\mathcal{D}} s^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x} + 2\int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)n(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x} + \int_{\mathcal{D}} n^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}$$

Assuming $\sigma_s^2 = \int_{\mathcal{D}} s^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x} \gg \sigma_n^2$ we obtain

$$E\left\{\left[\int_{\mathcal{D}} n^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}\right]^2\right\} \sim \sigma_n^4 \ll \sigma_s^2\sigma_n^2$$

$$E\left\{\left[\int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)n(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}\right]^2\right\} = \sigma_s^2\sigma_n^2$$

Hence the term $\int_{\mathcal{D}} n^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}$ can be neglected and the resultant SNR becomes

$$\mathbf{SNR}_p = \frac{E\left\{\left[\int_{\mathcal{D}} s^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}\right]^2\right\}}{E\left\{\left[2\int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)n(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x}\right]^2\right\}} = \frac{1}{4}\mathbf{SNR}_{opt}$$

■

Main shortcoming of pseudo-matched filtering is low discriminative power. Essentially, this filter computes a local signal variance, yielding large responses at all locally varying signals, such as edges, ridges, and corners. To overcome this problem we propose to exploit the rotational symmetry of the target features: vessels. As shown in Figure 2, rotating $s(\boldsymbol{x})$ by 180° around $\boldsymbol{x}_0$ does not change the signal, i.e.,

$$s(\boldsymbol{x}_0 - \boldsymbol{x}) = s(\boldsymbol{x}+\boldsymbol{x}_0) \tag{5}$$

Hence, we define a self-matched filter as follows

$$f(\boldsymbol{x}) = \tilde{s}(\boldsymbol{x}_0 - \boldsymbol{x}), \qquad \forall \boldsymbol{x} \in \mathcal{D} \tag{6}$$

According to the following lemma, the SNR of such a filter is half of that of the ideal matched filter and twice as large as the SNR of pseudo-matched filter.

Lemma 2 *If $\sigma_s^2 \gg \sigma_n^2$, the SNR achieved by self-matched filtering (6) is*

$$\mathbf{SNR}_{self} = \frac{1}{2}\mathbf{SNR}_{opt} \tag{7}$$

Proof: The proof is similar to that of Lemma 1. The response of the self-matched filter is

$$r(\boldsymbol{x}_0) = \int_{\mathcal{D}} s^2(\boldsymbol{x}+\boldsymbol{x}_0)d\boldsymbol{x} + \int_{\mathcal{D}} n(\boldsymbol{x}+\boldsymbol{x}_0)n(\boldsymbol{x}_0-\boldsymbol{x})d\boldsymbol{x} + \int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)\left[n(\boldsymbol{x}+\boldsymbol{x}_0)+n(\boldsymbol{x}_0-\boldsymbol{x})\right]d\boldsymbol{x}$$

The second term can be ignored due to $\sigma_s^2 \gg \sigma_n^2$. Furthermore, the miss-alignment of noise values in the last term reduces the noise energy by a factor of two, namely

$$E\left\{\left\{\int_{\mathcal{D}} s(\boldsymbol{x}+\boldsymbol{x}_0)\left[n(\boldsymbol{x}+\boldsymbol{x}_0)+n(\boldsymbol{x}_0-\boldsymbol{x})\right]d\boldsymbol{x}\right\}^2\right\} = 2\sigma_s^2\sigma_n^2$$

which yields (7). ■

In addition to achieving twice the SNR of pseudo-matched filtering, self-matched filtering has better discrimination for line features. When the signal in the local neighborhood is not symmetric, such as a corner, rotating it results in a very different signal and the self-matched filter produces a weak response. When the signal in the local neighborhood possesses a reflectional symmetry, such as a step edge, the self-matched filter produces a large but negative response. Only when the signal possesses a $180°$ rotational symmetry, such as a blood vessel, the response will be large and positive. Hence comparing the filter response with a positive threshold allows us to discriminate rotationally symmetric (by $180°$) features from other structures.

In the next section we show how to modify self-matched filter to discriminate light and dark vessels.

3. DISTINGUISHING LIGHT AND DARK VESSELS

The response of a Hessian filter is negative for light vessels and positive for dark ones, allowing one to distinguish them. Since self-matched filtering doesn't have this property (the response is always positive), we propose a simple solution that can help distinguish these types of vessels.

In a noiseless case, the center of a light vessel has a higher intensity than the mean intensity over its local neighborhood; the opposite is true for a dark vessel. Hence classification between dark and light vessels can be done by comparing the intensity of the center pixel with the mean intensity $\bar{s}$.

$$\begin{cases} \text{Light} & \bar{s} < s(\boldsymbol{x}_0) \\ \text{Dark} & \bar{s} > s(\boldsymbol{x}_0) \end{cases} \tag{8}$$

To reduce the effect of noise we can average intensity of several neighboring center pixels (instead of one).

Surprisingly, the same solution helps to address another potential problem. As shown in Figure 3(b), the self-matched filter produces a large output in between two parallel closely spaced vessels. This happens because the signal between the light vessels can be seen as a dark vessel, producing a large, positive response. Using (8) solves the problem (Figure 3(c)).

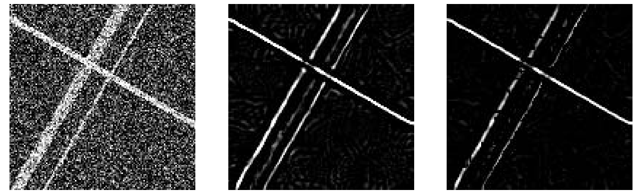

(a) (b) (c)

Fig. 3. The false vessel between parallel ones and its removal. (a) Parallel vessels (b) Self-matched filtering (c) Removing dark vessels

4. EVALUATION

To evaluate the performance of the proposed self-matched filter and compare it with the existing Hessian filter, we construct a simple post-processing chain that converts the filtering outputs into binary segmentation of vessel centerlines. Our pipeline [1] consists of hysteresis thresholding followed by morphological thinning to obtain a vessel centerline map. We use thinning because non-maximum suppression requires the local vessel orientation and cannot be applied to self-matched filtering results. For hysteresis thresholding the lower threshold was set as 0.4 times the upper threshold, which is a default value for the MATLAB function used.

We applied self-matched filtering to actual retinal images from DRIVE dataset [10], containing 40 images together with manually obtained ground truth vessel segmentation. We converted segmentation to vessel centerlines using morphological thinning. The window size of the self-matched filter was chosen to be 9×9, which is slightly larger than the maximum vessel width of 8 pixels [2]. For the Hessian filter based on Gaussian kernel $e^{-\frac{x^2}{t^2}}$, we used $t = 3.5$ and $t = 5.5$, which were chosen to slightly favor either thin or thick vessels.

Figure 4 shows the responses of the Hessian and self-matched filters on one of the retinal images. Visual inspection shows that compared to Hessian filter, the proposed self-matched filter yields stronger vessels and also is better at removing the distracting structures such as the field-of-view boundary[3], optic disc boundary and discolorations, but tends to produce more "holes" in the vessels.

Superiority in the visual sense translates into better detection results, as shown in Table 1 and Figure 6. The self-matched filter can reduce the area above the ROC curve (AAC) by up to a factor of 1.5, compared to Hessian filter. Note that, unlike in simulated images, in these experiment smaller scale

[1] More advance post-processing pipelines proposed were not considered since the emphasis of the paper is on the filtering step

[2] To estimate the vessel widths, we applied binary thinning operation, using 3×3 all-1 matrix as structure element, on all ground true vessel maps. If a vessel disappeared just after n'th thinning processes, its width was considered as $2n$.

[3] Note that the field-of-view boundary can be easily located and removed, and hence it has been excluded from all experimental results

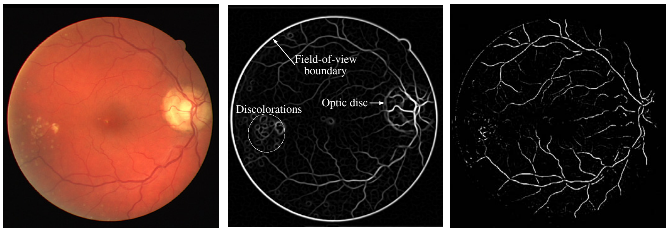

(a) (b) (c)

Fig. 4. The retinal image and the response of Hessian and self-matched filtering. (a) Retinal image (b) Hessian filtering (c) Self-matched filtering

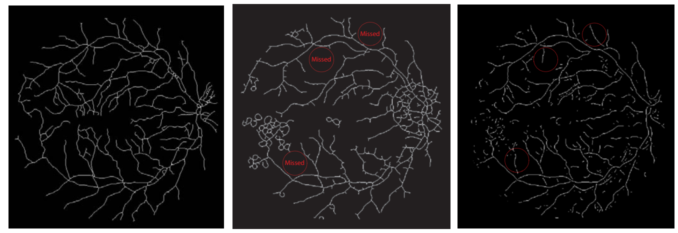

(a) (b) (c)

Fig. 6. Retinal image and its vessel maps detected by the Hessian and self-matched filters. (a) Ground true, (b) Hessian filter, (c) Self-matched filter

Table 1. Detection performance (AAC) of Hessian and self-matched filtering

Retinal Images	Hessian $t = 3.5$	Hessian $t = 5.5$	Self-matched	Retinal Images	Hessian $t = 3.5$	Hessian $t = 5.5$	Self-matched
1	0.136	0.255	0.156	11	0.326	0.447	0.246
2	0.190	0.278	0.175	12	0.275	0.365	0.191
3	0.284	0.364	0.248	13	0.285	0.394	0.208
4	0.318	0.419	0.253	14	0.224	0.329	0.180
5	0.313	0.401	0.247	15	0.267	0.390	0.208
6	0.297	0.392	0.250	16	0.261	0.367	0.149
7	0.341	0.451	0.268	17	0.306	0.382	0.253
8	0.336	0.453	0.267	18	0.271	0.391	0.178
9	0.317	0.422	0.231	19	0.210	0.304	0.136
10	0.294	0.398	0.235	20	0.267	0.366	0.180
				Average	0.276	0.378	0.213

Hessian filter ($t = 3.5$) performed better than the larger scale filter ($t = 5.5$) as shown in Figure 5. This can be explained by relatively few thick vessels present in these images. Another reason is the proximity of vessels which can cause miss detection when using large-scale Hessian filters. We have also tried widths $t = 3$ and $t = 4$, but have found them to perform slightly worse than $t = 3.5$.

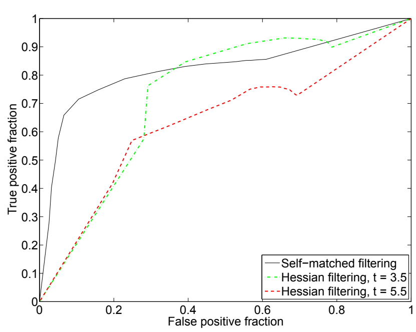

Fig. 5. ROC curves on test image 18 of DRIVE dataset.

5. CONCLUSION

We have proposed a novel filtering approach, called self-matched filtering, to improve the detection of blood vessels in retinal images. The proposed filter is a 180° rotated version of the signal in the local neighborhood, and has shown to achieve half the SNR of the optimal matched filter under normal conditions. However, since such a filtering does not require estimation of vessel orientations and widths, it yields favorable performance (both in quality and speed) compared with the alternative approach based on Hessian filtering. Moreover, the self-matched filtering can discriminate line features (vessels) from step edges, corners, and other distracting structures that might appear in the retinal images.

6. REFERENCES

[1] A. Can, C. Stewart, B. Roysam, and H. Tanenbaum, "A feature-based, robust, hierarchical algorithm for registering pairs of images of the curved human retina," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, no. 3, pp. 347–364, 2002.

[2] C. V. Stewart, C.-L. Tsai, and B. Roysam, "The dual-bootstrap iterative closest point algorithm with application to retinal image registration," *IEEE Trans. on Medical Imaging*, vol. 22, no. 11, pp. 1379–1394, Nov. 2003.

[3] T. Lindeberg, "Edge detection and ridge detection with automatic scale selection," *International Journal of Computer Vision*, vol. 30, pp. 117–156, 1998.

[4] M. Sofka and C. V. Stewart, "Retinal vessel extraction using multiscale matched filters, confidence and edge measures," *Technical Reports, #05-20, University of Innsbruck*, 2005.

[5] S. Chaudhuri, S. Chatterjee, N. Katz, M. Nelson, and M. Goldbaum, "Detection of blood vessels in retinal images using two-dimensional matched filters," *IEEE Trans. on Medical Imaging*, vol. 8, no. 3, pp. 263–269, Sept. 1989.

[6] S. M. Kay, *Fundamentals of Statistical Signal Processing*. New Jersey: Prentice-Hall, 1998.

[7] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, Sept. 1991.

[8] M. Jacob and M. Unser, "Design of steerable filters for feature detection using canny-like criteria," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1007–1019, Aug. 2004.

[9] A. Frangi, W. Niessen, K. Vincken, and M. Viergever, "Multiscale vessel enhancement filtering," *Medical Image Computing and Computer-Assited Intervention*, vol. 1496, pp. 130–137, 1995.

[10] J. Staal, M. Abramoff, M. Niemeijer, M. Viergever, and B. van Ginneken, "Ridge based vessel segmentation in color images of the retina," *IEEE Trans. on Medical Imaging*, vol. 23, no. 4, pp. 501–509, Apr. 2004.