Leveraging Multiscale Hessian-Based Enhancement With a Novel Exudate Inpainting Technique for Retinal Vessel Segmentation

Roberto Annunziata, Andrea Garzelli, Lucia Ballerini, Alessandro Mecocci, and Emanuele Trucco

Abstract—Accurate vessel detection in retinal images is an important and difficult task. Detection is made more challenging in pathological images with the presence of exudates and other abnormalities. In this paper, we present a new unsupervised vessel segmentation approach to address this problem. A novel inpainting filter, called neighborhood estimator before filling, is proposed to inpaint exudates in a way that nearby false positives are significantly reduced during vessel enhancement. Retinal vascular enhancement is achieved with a multiple-scale Hessian approach. Experimental results show that the proposed vessel segmentation method outperforms state-of-the-art algorithms reported in the recent literature, both visually and in terms of quantitative measurements, with overall mean accuracy of 95.62% on the STARE dataset and 95.81% on the HRF dataset.

Index Terms—Exudates, inpainting, retina, vessel segmentation.

I. INTRODUCTION

ANALYZING the vascular tree structure is useful for: monitoring arteriolar narrowing [1], characterizing plus disease in retinopathy of prematurity with tortuosity measurements [2]–[4], and the diagnosis of hypertension and cardiovascular diseases through accurate vessel width estimation [5], [6]. However, since the manual segmentation of retinal vessels is extremely time consuming, automated segmentation becomes crucial. Accurate vessel segmentation is a very difficult task for several reasons: 1) the presence of lesions, exudates, haemorrhages; 2) the variability of the vessel width and length; 3) the low contrast between the vessels and the background; 4) the central reflex on large vessels; 5) the presence of small regions affected by noise; and 6) the occlusion between vessels. Several methods have been developed for retinal blood vessel segmentation. One of the main weaknesses of previously reported methods is that their results tend to degrade when applied to pathological eyes. In particular, they produce a high number of false positives in the presence of exudates, haemorrhages, and other confounding retinal structures. These limitations have motivated the development of the framework described here, which improves on previous algorithms when applied to abnormal (and also normal) images and manages the presence of exudates and similar retinal features in a more robust manner.

The main contribution of this paper is an *ad-hoc* exudate inpainting technique. Several works have been presented to solve the problem of filling holes in digital images by propagating surrounding structures to synthesize a visually plausible image (for instance, [8]). The goal of our inpainting technique has a different focus. Indeed, neither texture synthesis nor a visually plausible image is needed. Our goal is to fill structures such as exudates in retinal images so that, when vessel enhancement is applied, the number of nearby false positives is greatly reduced. This goal is only achieved if exudates are filled in a smooth way that reduces or eliminates possible edges. A multiple-scale Hessian-based enhancement is applied to detect retinal vessels. This technique is fast and has proven to be effective when detecting vessels of normal eyes. However, it also enhances other retinal structures and, therefore, becomes unsuitable for a general retinal vessel segmentation framework. The key idea of the proposed method is to apply Hessian-based enhancement after exudate inpainting. This reduces false vessel detection. Although simple in principle, the accuracy achieved by our method is comparable or higher than those reported in the literature. Moreover, it yields the best performance on pathological images, the target of most automated retinal image analysis tools. Indeed, a vessel segmentation algorithm is usually the first step for the automated detection of eye diseases. In order to be used in clinical practice, these methods should be robust enough to analyze pathological and nonpathological images without requiring user interaction. We propose a fully automated algorithm. Our results suggest that joint detection of vessels and other retinal structures could finally solve the problem of accurate and reliable retinal vessel segmentation suitable to various practical scenarios.

This paper is organized as follows. Section II gives an overview of the state-of-the-art vessel segmentation methods, Section III describes the datasets we used. Section IV presents in detail the proposed method. Experimental results are provided in Section V. Finally, we conclude the paper in Section VI.

R. Annunziata was with the Department of Information Engineering and Mathematical Sciences, Universitá degli Studi di Siena, 53100 Siena, Italy. He is now with the VAMPIRE, CVIP Group, School of Computing, University of Dundee, Dundee DD1 4HN, U.K. (e-mail: r.annunziata@dundee.ac.uk).

A. Garzelli and A. Mecocci are with the Department of Information Engineering and Mathematical Sciences, Universitá degli Studi di Siena, 53100 Siena, Italy (e-mail: andrea.garzelli@unisi.it; alemecoc@alice.it).

L. Ballerini and E. Trucco are with the VAMPIRE, CVIP Group, School of Computing, University of Dundee, Dundee DD1 4HN, U.K. (e-mail: l.ballerini@dundee.ac.uk; e.trucco@dundee.ac.uk).

II. Related Work

Many retinal vessel segmentation methodologies have been proposed: supervised and unsupervised. A recent detailed review of these methods can be found in [7]. In general, the performance of supervised methods is higher than that of unsupervised ones. On the other hand, supervised methods require a preliminary training phase which is time consuming because it needs a training set of manually segmented images for each camera setup.

Supervised segmentation methods use ground truth data for classifying each image pixel, based on given features. For instance, k-nearest neighbor was used by Staal *et al.* [9] to classify feature vectors obtained by a ridge detector. In [10], six features are computed using multiscale analysis of Gabor wavelet transform. The approach adopts two kinds of classifiers: Gaussian Mixture Model Bayesian and Linear Minimum Squared Error.[1] Ricci and Perfetti [12] used line operators and by support vector machine classification. A recent supervised approach is based on mathematical morphology and moment invariant features, followed by a neural network classifier [13]. Fraz *et al.* [14] employed an ensemble of bagged decision trees and a feature vector based on the orientation analysis of gradient vector field, morphological transformation, line strength measures, and Gabor filter responses. Finally, combining hand-crafted features with learned context filters has been recently shown to improve performance on challenging curvilinear structures such as corneal nerve fibers and neurites [15].

Unsupervised methods include techniques based on matched filtering, morphological processing, vessel tracking, multiscale analysis, and model-based algorithms [7]. In [16], matched filtering is used; it is based on 2-D linear structural element with a Gaussian cross section, rotated through many orientations. A thresholding technique is then applied to obtain the segmented vessels. In [17], a different approach is proposed, based on a multithreshold probing scheme. Mathematical morphology in combination with matched filtering for centerline detection is exploited by Mendonça and Campilho [18]. Martinez-Perez [19] proposed a method in which the vascular tree is obtained by using a multiscale feature extraction approach. The local maxima of the gradient magnitude over different scales, the maximum principal curvature of the Hessian matrix [20], and a region growing scheme are combined to segment the retinal image. In [21], a similar enhancement step is applied, but a different pre-processing step is proposed to decrease the disturbance of bright structures before vessel extraction. Azzopardi *et al.* [22] introduced a method based on the combination of shifted filter responses with promising results. Recently, a novel hand-crafted feature scale and curvature invariant ridge detector (SCIRD), has been proposed to achieve multiple invariances when segmenting tortuous and fragmented structures [23].

Although a lot of work has been done on automated retinal vessel segmentation, very little exists on vessel detection approaches in which other retinal structures are taken into account during vessel detection. For instance, a divergence vector field is proposed by Lam and Yan [24] and adapted to handle bright lesions. Recently, Lam *et al.* [25] proposed a model-based approach with differentiable concavity measure to handle both healthy and unhealthy retinal images.

Exudate detection methods reported in the literature range from region growing methods [26] for candidate detection to morphological reconstruction for obtaining a precise localization of the exudate boundaries [27]. Complex machine learning methods can also be used along with different sorts of features [28]. However, we present here a simple method for exudate detection, since our goal is not accurate exudate segmentation for pathology detection and analysis, but their removal before Hessian-based vessel enhancement.

In this paper, we propose a new pipeline for retinal vessel segmentation whose main component is an *ad-hoc* exudate inpainting filter aimed at reducing false detection generated by the strong contrast around exudates. This makes our retinal vessel segmentation framework more general than previously reported methods, since it is suitable for both healthy and unhealthy eyes affected by exudate regions and similar structures.

III. Materials

To evaluate the performance of the vessel segmentation approach described in the next section, two publicly available datasets are used: the low-resolution structured analysis of the retina (STARE)[2] dataset [16] and the high-resolution fundus (HRF)[3] image dataset [29].

The STARE database contains 20 retinal images captured by a TopCon TRV-50 fundus camera at $35°$ field of view (FOV). The images were digitized to 700×605 pixels, 8 bits per color channel. The FOV in the images are approximately 650×550 pixels. Unlike other datasets, STARE covers several abnormal cases, using ten retinal images. There are patients who have serious problems such as large regions of exudates, multiple haemorrhages, and vessel occlusions that can affect vessel segmentation. Only the first observer's manual segmentations were used to validate our method, a common choice for this dataset (e.g., in [10] and [22]).

The HRF image dataset contains retinal images taken with a CANON CF-60UVi fundus camera, with an attached CANON EOS-20D digital camera at $60°$ FOV. Each image is digitized to 3504×2336 pixels, 8 bits per color channel and compressed in JPEG format. This high resolution is comparable to the common resolution in clinical use. The dataset contains 45 images divided into three subsets: healthy fundus, diabetic retinopathy (DR), and glaucoma. The retinal images of DR patients present pathological changes, such as neovascular nets, haemorrhages, bright lesions, and spots after laser treatment. Patients with glaucoma present symptoms of focal and diffuse nerve fiber layer loss. These last two subsets allow evaluation of segmentation methods on pathological retinas. Each subset has 15 images with FOV masks and manual segmentation gold standard.

---

[1]Currently the VAMPIRE software suite [11] implements a version of Soares' algorithm as the best compromise between speed and accuracy.

[2]STARE http://www.ces.clemson.edu/
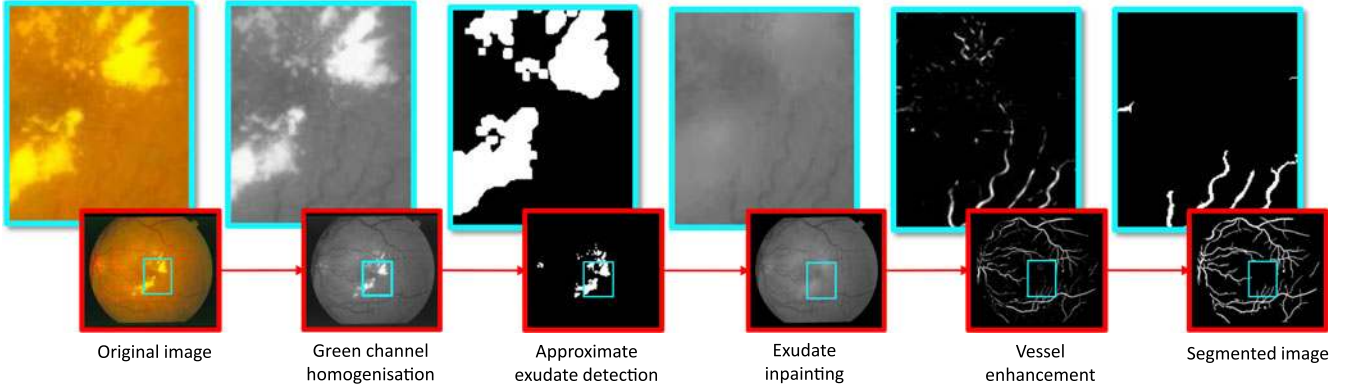[3]HRF http://www5.informatik.uni-erlangen.de/research/data/fundus-images

Fig. 1.   Stepwise illustration of the proposed technique.

IV. PROPOSED METHOD

An overview of the proposed approach is depicted in Fig. 1. The following steps can be identified:

(1) image preprocessing for exudate detection;
(2) exudate inpainting;
(3) multiscale Hessian eigenvalue analysis for vessels enhancement;
(4) percentile-based thresholding.

Only the green channel of the RGB original image was used as it offers the best vessel-background contrast.

A. Image Preprocessing for Exudate Detection

Typically, exudates appear much brighter than vessels (see Fig. 5(a), for example). However, nonuniform illumination and inhomogeneities make unfeasable a simple gray-level thresholding for identifying them. Indeed, exudate pixels in a retinal image may have the same gray level of poorly contrasted thin vessel pixels. To address these issues, a preprocessing phase similar to [13] is applied. This phase consists of

(1) nonuniform illumination correction;
(2) image homogenization;

We use a large median filter (69×69 for STARE and 139×139 for HRF) for background estimation. This filter has been selected because it is particularly effective at roughly removing blood vessels without blurring edges of larger regions in the background. This median filter is applied to the region of interest (ROI). The ROI has been previously expanded to avoid border artefacts [see Fig. 2(a) and (f)]. Then, the estimated background I_{med} [see Fig. 2(b) and (g)] is subtracted from the green channel of the original image I to obtain the difference image D:

$$D(x,y) = I(x,y) - I_{\mathrm{med}}(x,y). \tag{1}$$

The illumination corrected image I_C [see Fig. 2(c) and (h)] is obtained by linearly stretching the gray-levels of D to cover the whole range of possible intensity values ($[0, 255]$, for an 8-bit per pixel image). The homogenization step is carried out as follows [13]: The histogram of the brightness-corrected image I_C is displaced toward the middle of the gray scale, by modifying

pixels intensity according to

$$g_{\mathrm{Output}} = \begin{cases} 0, & \text{if } g < 0 \\ 255, & \text{if } g > 255 \\ g, & \text{otherwise} \end{cases} \tag{2}$$

where

$$g = g_{\mathrm{Input}} + 128 - g_{\mathrm{Input}_M} \tag{3}$$

and g_{Input} and g_{Output} are the gray-level values of the input and the output images (I_C and I_H, respectively). The value denoted by g_{Input_M} is the mode in the histogram of I_C. The homogenization step is based on the idea that the background consists of much more pixels than the foreground (vessels in this case), so the intensity value corresponding to the mode of the histogram represents the background value. Then, a 3×3 median filter is applied to the homogenized image I_H for removing residual noise [see Fig. 2(d) and (i)].

Finally, thanks to the previous homogenization, a simple threshold can be used to obtain exudate masks since the gray level of the vessel is now much lower than that of the exudate. We experimentally observed that selecting a threshold in the range $[160, 170]$ does not change the final performance for both datasets. We tuned this threshold taking into account the tradeoff between having a percentage of undetected exudates and false positives.

B. Exudate Inpainting

We propose a novel inpainting filter (*Algorithm 1*), called *neighborhood estimator before filling* (NEBF) to fill detected exudate regions.

Algorithm 1 NEBF

$ExudMask \leftarrow \text{dilate}(ExudMask)$;
$TmpInp \leftarrow OrgImg(ExudMask \neq 0) = 0$;
while all exudates are not inpainted **do**
    $ExudMask \leftarrow \text{erode}(ExudMask)$;
    $TmpInp \leftarrow \textbf{call } ExudInp(TmpInp, ExudMask)$;
end while
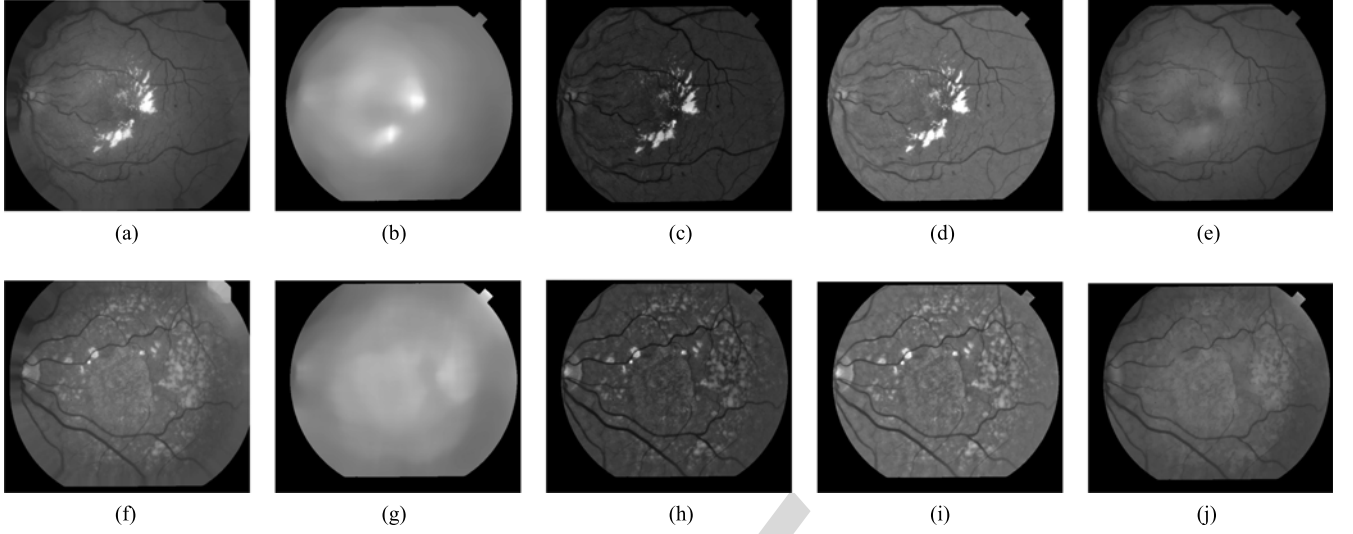$ImgInp \leftarrow TmpInp$;

Fig. 2. Exudate detection and inpainting on images 1 and 3 from STARE dataset: (a), (f) ROI expansion of the green channel. (b), (g) Background estimation. (c), (h) Nonuniform illumination correction. (d), (i) Homogenized image. (e), (j) Exudate inpainting.

Algorithm 2 $I = ExudInp(I, ExudMask)$

$PxToFill \leftarrow ExudMask$ - erode($ExudMask$);
$\forall \; \mathbf{p} \in PxToFill \mid PxToFill(\mathbf{p}) \neq 0$

$$I(\mathbf{p}) = \; mean \; I(\mathbf{q})$$
$$\mathbf{q} \in \mathbf{N_p}$$
$$I(\mathbf{q}) \neq 0$$

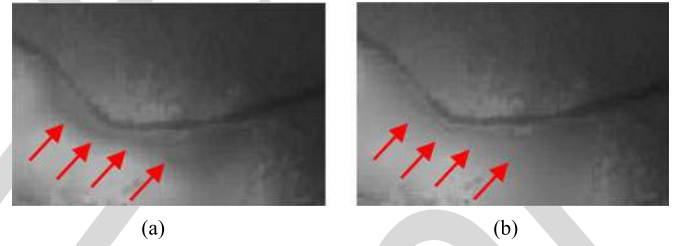$N_p = \{\mathbf{q} \in \mathbb{N}, |\mathbf{q} - \mathbf{p}| \leq r\}$

Fig. 3. Exudates inpainted by the NEBF filter. (a) Without neighborhood estimation. (b) With neighborhood estimation. Note smoother edges within exudate region in (b).

The algorithm proceeds iteratively in a radial way towards the exudate's core. Using a conservative threshold to detect exudates (set to reduce false positives, therefore allowing more false negatives) typically undersegments each individual exudate, leaving a narrow border of undetected pixels. For this reason, we dilate the detected exudate mask after thresholding with a circle of radius 3 for STARE and 6 for HRF. We then proceed radially toward the exudate's core. The structuring element for the erosions in the Algorithms 1 and 2 is a circle with radius of 1 pixel for STARE and 3 pixels for HRF. Our goal is to fill exudates in a very smooth way. As Algorithm 2 shows, this is accomplished by averaging both background and already estimated values falling in the eight-connected neighborhood of each pixel (the radius $r = 3$ for STARE and $r = 7$ for HRF). During the averaging process, detected exudate pixels (set to 0) are not taken into account. In fact, iteration by iteration, the influence of background pixels decreases, while that of estimated pixels increases.

Notice that NEBF is applied to the original image (not to the homogenized one used only to detect exudates).

Estimating the neighborhood before filling is a key advantage of NEBF, since it reduces greatly radial strips and edges creation within filled regions (see Fig. 3). These artefacts would lead to many false positives in the following enhancement and segmentation steps. As we show in Section V-C (see Fig. 9), our method is more suitable than a state-of-the-art inpainting technique. In fact, we aim at creating smooth inpainted regions rather than filling exudates in a visually plausible way.

A *linear-opening-by-reconstruction* [30] is used to remove smaller and poorly contrasted exudates not detected in the previous steps. It can be mathematically stated as

$$\min(\gamma_B(I), I) \tag{4}$$

where $\gamma_B(I)$ is defined as the morphological opening of I using B as structuring element (line 15×1 for STARE and 30×1 for HRF).

Notice that linear-opening-by-reconstruction preserves all the structures except small undetected exudates. As a result, thin vessels' color and morphometric characteristics are not altered.

Fig. 2 shows all the steps of our exudate inpainting on two retinal images in STARE dataset.

C. Multiscale Hessian Eigenvalue Analysis for Vessels Enhancement

Hessian-based methods have proven effective in retinal vessel enhancement [19]–[21], [31]. The key idea is to extract principal directions in which the local second-order structure of
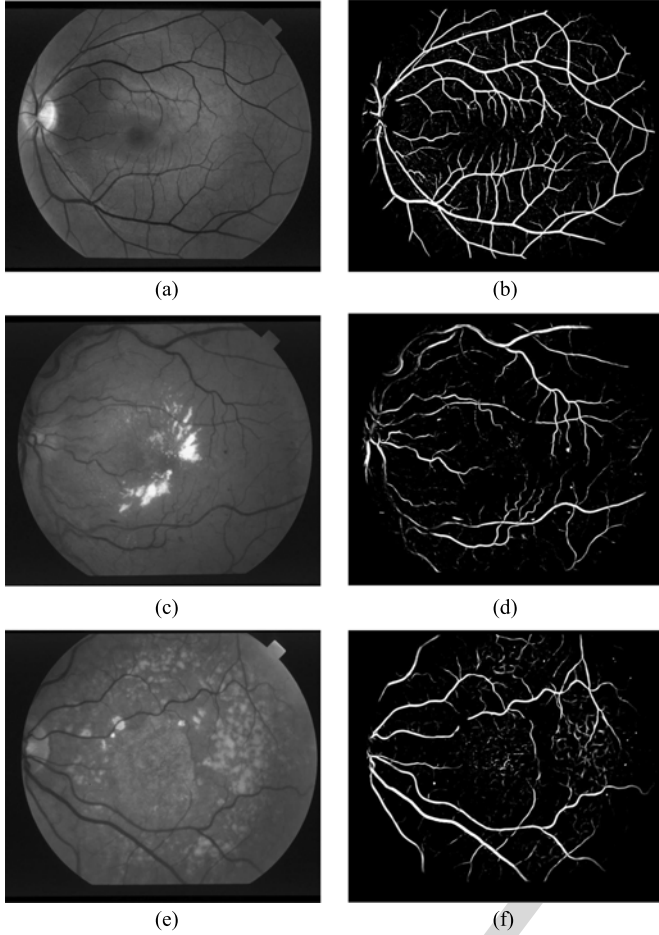
Fig. 4. Vessels enhancement applied to a normal case and to two abnormal cases from the STARE dataset.

the image can be decomposed. Analyzing a vessel, the largest eigenvalue (λ_2) of the Hessian matrix is relative to the smallest eigenvector, which is ideally orthogonal to vessel's walls. On the contrary, the smallest eigenvalue (λ_1) corresponding to the largest eigenvector is aligned with the vessel.

We carry out eigenvalue analysis at multiple spatial scales ($s \in \{2, 3, 4\}$ for STARE and $s \in \{2, 3, 4, 5, 6\}$ for HRF) to enhance vessels regardless their width.

The largest eigenvalue over scales, $\lambda_{\max}$, is obtained as

$$\lambda_{\max} = \max_s \frac{\lambda_2(s)}{s}. \tag{5}$$

Notice that normalizing by the scale factor s in (5) leads to unbiased comparison among scales [19].

We found that using only the largest eigenvalue is sufficient for vessel enhancement, unlike previously reported methods [19]–[21], [31], which make use of both λ_1 and λ_2.

The first row in Fig. 4 shows vessel enhancement carried out on a normal case: our multiscale approach is capable to enhance both wide and thin vessels; furthermore, false positives near the optic nerve's border are greatly reduced. The second and the third rows in Fig. 4 show vessel enhancement on two abnormal cases with large and small exudates: our preprocessing step greatly reduced false positives at exudate borders.

D. Percentile-Based Thresholding

We employ a simple thresholding algorithm based on the percentile to obtain the final vessel segmentation. Indeed, we threshold the cumulative histogram of the Hessian enhanced image. The chosen threshold is the intensity value keeping a certain percentage of pixels, whose value is estimated on the training dataset using the observers segmentation as reported in [10].

Finally, small nonvessel isolated connected components are removed by area thresholding. Notice that our segmented vascular trees have great connectivity; therefore, our framework is not very sensitive to this specific postprocessing parameter.

V. EXPERIMENTAL EVALUATION

A. Performance Measures

In order to quantify performance, we use *Sensitivity* (*Se*), *Specificity* (*Sp*), *Positive Predictive Value* (*PPV*), *Negative Predictive Value* (*NPV*), and *Accuracy* (*Acc*). These measures are defined as

$$Se = \frac{TP}{TP + FN} \tag{6}$$

$$Sp = \frac{TN}{TN + FP} \tag{7}$$

$$PPV = \frac{TP}{TP + FP} \tag{8}$$

$$NPV = \frac{TN}{TN + FN} \tag{9}$$

$$Acc = \frac{TP + TN}{TP + FN + TN + FP} \tag{10}$$

where TP (true positives), FP (false positives), FN (false negatives), and TN (true negatives) are obtained by considering only pixels within the FOV. *Se* and *Sp* measures are the ratio of well-classified vessel and non-vessel pixels, respectively. PPV is the ratio of correctly classified vessel pixels. NPV is the ratio of correctly classified nonvessel pixels. Finally, *Acc* is the proportion of true results (both true positives and true negatives) in the population of pixels. We also measured performance using of receiver operating characteristic (ROC) curves by varying the percentile threshold. The area under the ROC curve (AUC) is also used. Following previous work (e.g., [21], [22], [25]), we computed all performance measures for each image and then reported averages. Accuracy and AUC are used to rank methods in terms of overall performance and other performance measures are used to highlight differences among methods for specific tasks (e.g., Se—the ability to detect vessel pixels; Sp—the ability to reduce FP).

B. Experimental Setup

We adopt the same experimental setup as most of the previous works, separating datasets into training and testing sets for a fair comparison. The parameters of our method have been manually tuned using a subset of training images, taking into account the resolution of each dataset. For the STARE dataset,

TABLE I
PERFORMANCE COMPARISON OF VESSEL SEGMENTATION METHODS ON THE STARE DATASET

	Method	Se	Sp	NPV	PPV	AUC	Acc
	STARE						
Unsupervised	Hoover *et al.* [16]	0.6747	0.9565	–	–	0.759	0.9275
	Jiang and Mojon [17]	–	–	–	–	0.9298	0.9009
	Mendonça *et al.* [18]	0.6996	0.973	–	–	–	0.9479
	Martinez-Perez *et al.* [19]	0.7506	0.9569	–	–	–	0.941
	Al-Rawi *et al.* [32]	–	–	–	–	0.9467	0.909
	Ricci and Perfetti [12]	–	–	–	–	0.9602	0.9584
	Al-Diri *et al.* [33]	0.7521	0.9681	–	–	–	–
	Lam *et al.* [25]	–	–	–	–	0.9739	0.9567
	Yu *et al.* [21]	0.7112	0.9709	–	–	–	0.9463
	Azzopardi *et al.* [22]	0.7716	0.9701	–	–	0.9563	0.9497
	No inpainting	0.6911	0.9813	0.9648	0.8085	–	0.9511
	Proposed method	**0.7128**	**0.9836**	**0.9677**	**0.8331**	**0.9655**	**0.9562**
Supervised	Staal *et al.* [9]	–	–	–	–	0.9614	0.9516
	Soares *et al.* [10]	0.7207	0.9747	–	–	0.9671	0.948
	Ricci and Perfetti [12]	–	–	–	–	0.968	0.9646
	Marín *et al.* [13]	0.6944	0.9819	0.9659	0.8227	0.9769	0.9526
	Fraz *et al.* [14]	0.7548	0.9763	–	–	0.9768	0.9534

the medial filter size used in the homogenization step has been set to 69×69 following [13]; the average filter size used in the NEBF has been set to 7×7 to achieve a good compromise between speed and exudate smoothing; the size of the structuring element in (4) has been set to 15×1 since 15 pixels is approximately the maximum vessel width; the number of scales in the Hessian eigenvalue analysis has been set considering the minimum and the maximum vessel width; the threshold 165 for the exudate detection has been set to reduce the amount of false positives at the expenses of true positives that were in general very small [for this reason, we apply the linear opening by reconstruction in (4)]. Notice that the homogenization step which corrects for nonuniform illumination changes and centers the histogram of each image allows us to set a single threshold across all datasets. For the HRF dataset, since the maximum vessel width is approximately 30 pixels, we doubled all filters width (and height) to test the generalization performance of our method on a different set of images.

C. Vessel Segmentation Results

1) STARE: Our approach is tested on the STARE dataset [16] using the first observer's manual segmentation as ground truth.

Table I shows comparison with state-of-the-art methods.

Performance measures in Table I show that our unsupervised method outperforms most of the state-of-the-art unsupervised and supervised algorithms. The algorithms presented by Ricci and Perfetti [12] and Lam *et al.* [25] reported an accuracy higher than our algorithm. However, Ricci and Perfetti built their classifier by using a training set comprising samples randomly extracted from test images while we use a leave-one-out strategy to set the percentile threshold. Indeed, Lam *et al.* [25] reimplemented their method reporting an accuracy of 0.9422. Lam *et al.*'s [25] approach greatly reduces the detection of false positives hence increasing accuracy. Our approach achieves the same goal with a much simpler and faster algorithm. In fact, time to

TABLE II
PERFORMANCE COMPARISON OF VESSEL SEGMENTATION METHODS ON THE PATHOLOGICAL IMAGES OF THE STARE DATASET

	Method	Acc
	STARE - Abnormal Images	
Unsup	Jiang and Mojon [17]	0.9337
	Mendonca and Campilho [18]	0.9426
	Lam *et al.* [24]	0.9474
	Line (Impl. in [25])	0.9352
	Lam *et al.* [25]	0.9556
	No inpainting	**0.9449**
	Proposed method	**0.9565**
Sup	Soares *et al.* [10]	0.9425
	Marin *et al.* [13]	0.9510

Accuracy values are from [25] and [13].

run each STARE image for Lam *et al.*'s technique is approximately 13 min as reported in [25], while our procedure takes about 1 min. We implemented the proposed method in MATLAB, running on a PC with an AMD A4-3300M APU at 1.90 GHz and 6-GB RAM. In a prototype implemented using C\C++ without any optimization, the time decreased to less than 25 s.

Notice that reported performance of supervised methods is generally higher than unsupervised ones. On the other hand, supervised approaches need a preliminary time-consuming training phase that requires manually segmented images. Our performance is comparable or superior to that of state-of-the-art supervised methods, without the need for this time-consuming stage. To facilitate comparisons, we report PPV and NPV as done by Marín *et al.* in [13], a top-ranking supervised method. Table I shows that the proposed method detects more of the vasculature ($Se = 0.7128$ versus 0.6944 for Marín *et al.*) while reducing the FP count (a 1% increase in PPV for the proposed method). This table also reports our method performance without the *ad-hoc* exudate inpainting stage ("No inpainting"). We observe that using exudate inpainting, PPV increases from 0.8085 to 0.8331
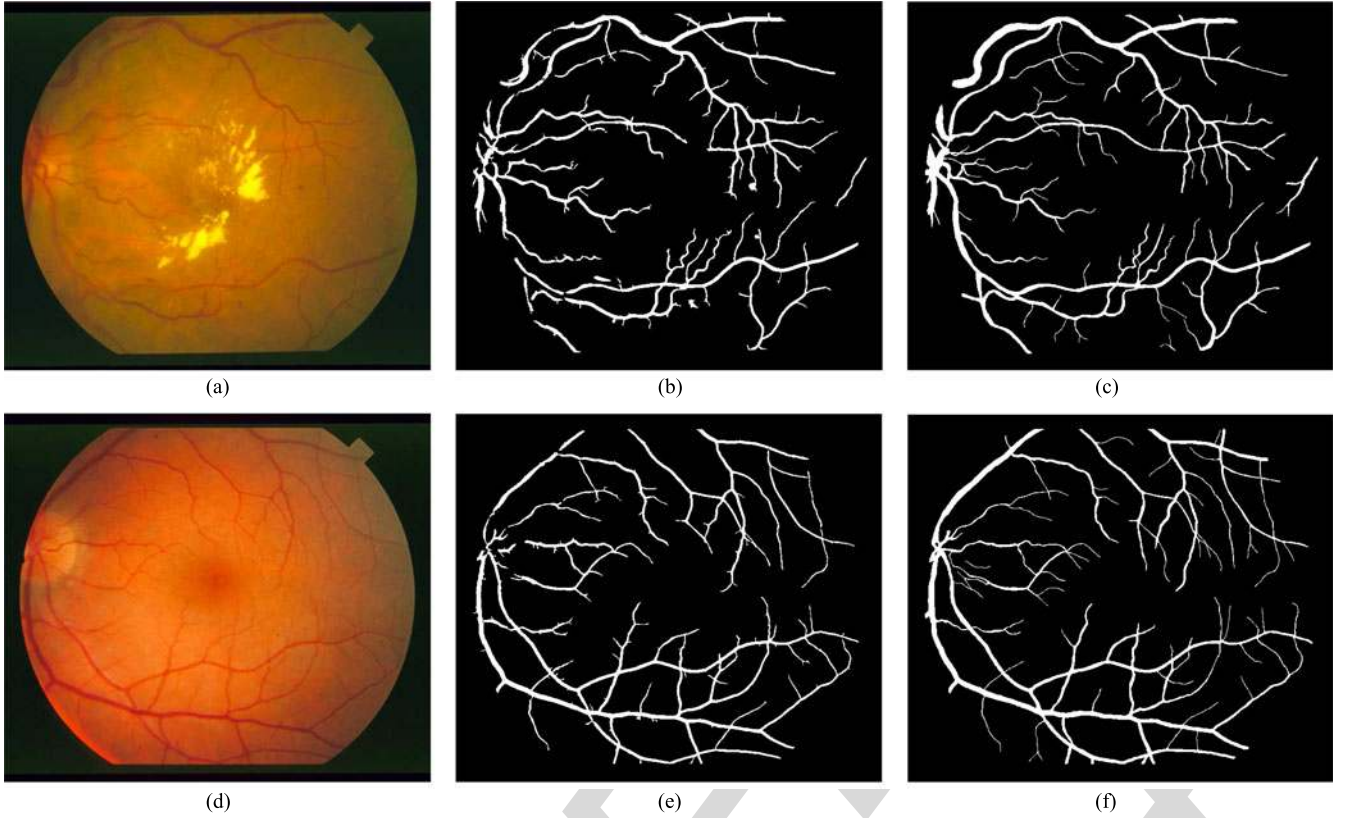
Fig. 5. Segmentation of a pathological and normal image of STARE. (a) Original image "im0001.ppm." (b) Vessel segmentation of (a) using our method. (c) Manual segmentation (a) by observer 1. (d) Original image "im0081.ppm." (e) Vessel segmentation of (d) using our method. (f) Manual segmentation of (d) by observer 1.
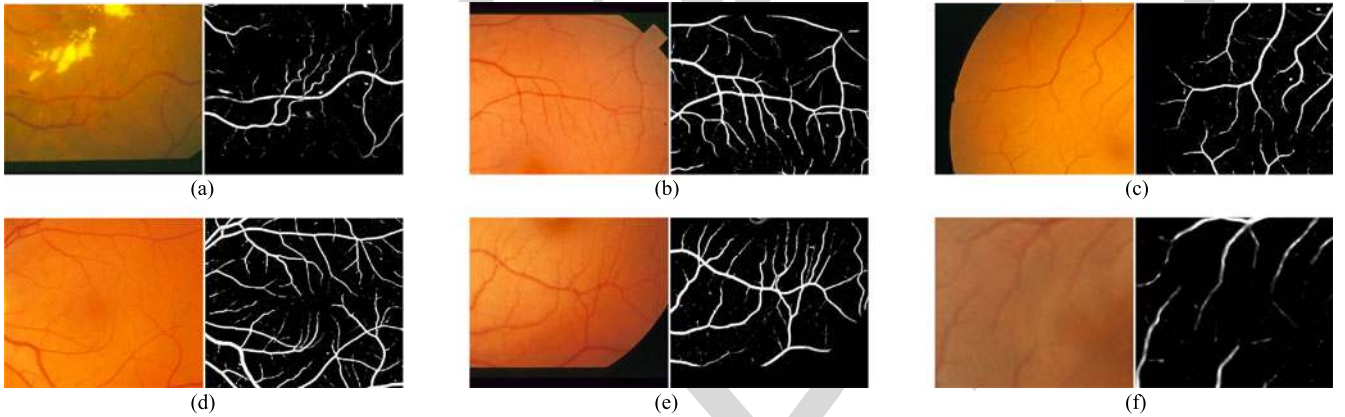

Fig. 6. Qualitative evaluation in detecting thin and poorly contrasted vessels in six challenging areas of image of the STARE dataset.

showing a strong reduction of false positives with a slightly better NPV.

Table II shows results on the ten abnormal images of the STARE database.[4] The accuracy of our method is the best among all the unsupervised and supervised methods on this subset of images. Indeed our accuracy on abnormal images is, on average, comparable to that of the whole dataset (i.e., 0.9565 versus 0.9562). Notice that the average accuracy of our pipeline without the exudate inpainting stage ("No inpainting") is 0.9449, which indicates the effectiveness of our exudate inpainting procedure as preprocessing step.

Fig. 5 shows a normal and abnormal case from the STARE dataset, our automatic segmentation, and their respective ground truth. Notice that due to our exudate inpainting step, few false positives are created near exudate regions.

Fig. 6 shows qualitative evaluation of our method on six challenging areas: most of the poorly contrasted and tortuous vessels are detected.

[4]As stated at http://www.ces.clemson.edu/ ahoover/stare/diagnoses/all-mg-codes.txt, abnormal images are: im0001.ppm, im0002.ppm, im0003.ppm, im0004.ppm, im0005.ppm, im0044.ppm, im0139.ppm, im0291.ppm, im0319.ppm and im0324.ppm.
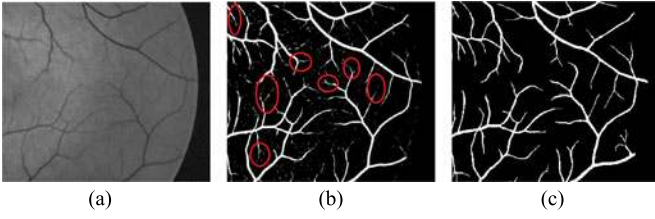
Fig. 7. Thin structured missed by observer 1 and detected by our algorithm. (a) Original image. (b) Vessel enhancement. (c) Manual segmentation by observer 1.
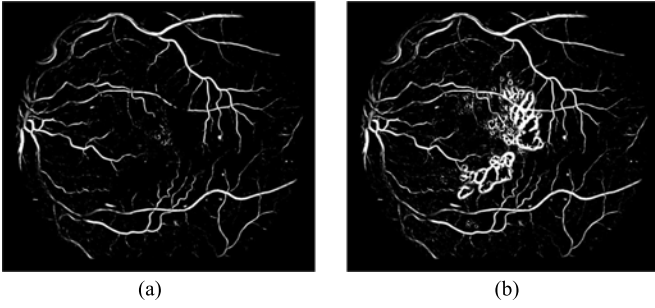

Fig. 8. Comparison of vessels enhancement with and without the exudate inpaiting technique.

Fig. 9. Comparison between NEBF and Criminisi *et al.* method for exudate inpainting. (a) Criminisi *et al.* inpainting, (b) Enhancement image after Criminisi *et al.* inpainting. (c) NEBF inpainting. (d) Enhancement image after NEBF inpainting.

The proposed method achieves top-rank performance when using Observer 1 as ground truth. However, manual segmentation of thin vessels is a challenging task even for experienced human observers as reported elsewhere (e.g., [21]). Further visual inspection of our enhanced images reveals that these structures are often detected by our algorithm but missed by Observer 1 as shown in Fig. 7. Notice that this aspect yields a slight decrease of measured accuracy since thin vessels missed by the observer and segmented by our method are regarded as false positives.

Fig. 8 shows a comparison of vessel enhancement with and without the inpainting technique applied to the fundus image in Fig. 5(a).

Furthermore, Fig. 9 shows a comparison between the NEBF and the inpainting technique proposed by Criminisi *et al.* [8] applied to the fundus image in Fig. 5(a). As can be seen in Fig. 9(a), estimated exudates are efficiently inpainted in a way that is visibly consistent with the background. However, that technique can create artefacts that can potentially lead to false positives. Instead, Fig. 9(c) and (d) shows the results obtained applying the NEBF. A few artefacts or edges are visible within the inpainted exudates.

2) HRF: We evaluated our method performance on HRF dataset containing higher resolution images [29]. Table III shows a comparison with the previously reported methods for the HRF: Yu *et al.* [21] and Odstrcilik *et al.* [29]. Our method shows a lower *Se* with respect to others, as we employ a simple yet fast algorithm for vessel detection. Nevertheless, higher *Sp* makes the proposed approach the best in terms of overall accuracy on the whole dataset. This higher accuracy level is entirely driven by a much lower number of FPs, thus confirming the key role of our inpainting strategy. A further analysis on each subset (i.e. Healthy, DR, and Glaucoma) reveals that our approach outperforms the others in terms of accuracy mostly on the un-
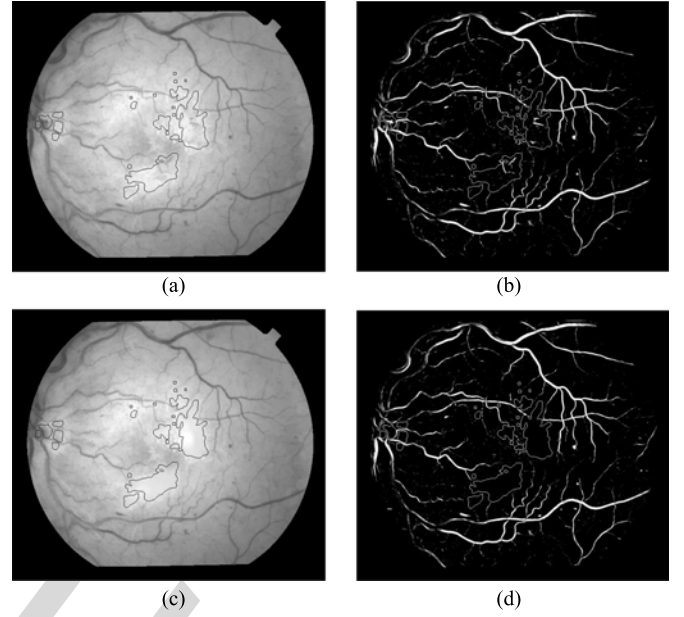
TABLE III
PERFORMANCE COMPARISON OF VESSEL SEGMENTATION METHODS ON THE HRF DATASET

			HRF			
Data set	**Methods**	**Se**	**Sp**	**NPV**	**PPV**	**Acc**
H	Odstrcilik *et al.* [29]	0.7861	0.9750	–	–	0.9539
	Yu *et al.* [21]	0.7938	0.9767	–	–	0.9566
	Our method	0.6820	0.9935	0.9614	0.9271	**0.9587**
DR	Odstrcilik *et al.* [29]	0.7463	0.9619	–	–	0.9445
	Yu *et al.* [21]	0.7604	0.9625	–	–	0.9460
	Our method	0.6997	0.9787	0.9729	0.7428	**0.9554**
G	Odstrcilik *et al.* [29]	0.7900	0.9638	–	–	0.9497
	Yu *et al.* [21]	0.7890	0.9662	–	–	0.9518
	Our method	0.7566	0.9785	0.9783	0.7567	**0.9603**
ALL	Odstrcilik *et al.* [29]	0.7741	0.9669	–	–	0.9494
	Yu *et al.* [21]	0.7811	0.9685	–	–	0.9515
	Our method	0.7128	0.9836	0.9709	0.8089	**0.9581**

healthy patients where exudates produce a high number of FPs in other methods. Fig. 10 shows segmentation results for the same healthy, DR and Glaucoma cases as reported by Odstrcilik *et al.* [29] ("06_dr" is also reported by Yu *et al.* [21]).

VI. CONCLUSION

An improved Hessian-based approach for unsupervised retinal blood vessel segmentation using an *ad-hoc* exudate inpainting technique has been described. The application of the proposed exudate inpainting technique, followed by a simple enhancement and thresholding method, yields results comparable to state-of-the-art techniques that use specialized, sophisticated enhancement, and classification algorithms. Our approach
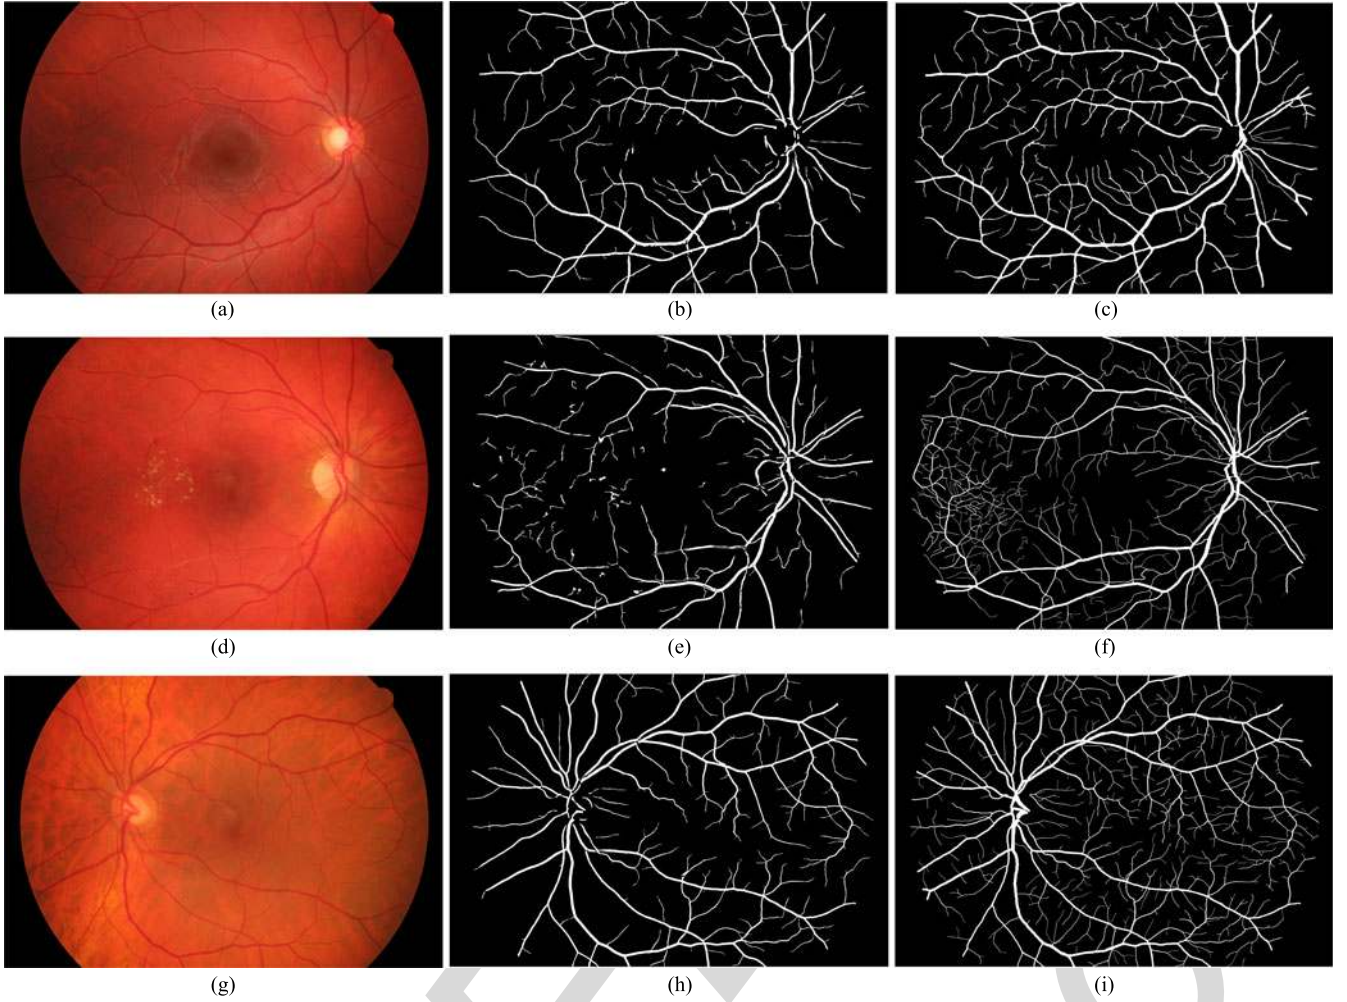
Fig. 10. Comparison of our method segmentation results (second column) with corresponding ground truth (third column) (HRF). (a) Original image "13_h.jpg" from the Healthy dataset. (b) Segmentation results of (a). (c) Manual segmentation of (a). (d) Original image "06_dr.jpg" from the Diabetic Retinopathy dataset. (e) Segmentation results of (d). (f) Manual segmentation of (d). (g) Original image "12_g.jpg" from the Glaucoma dataset. (h) Segmentation results of (g). (i) Manual segmentation of (g).

performs better than previously reported methods on the several challenges of retinal vessel detection.

Experimental results demonstrate the excellent performance of our segmentation method both in pathological and nonpathological retinas included in the STARE and HRF datasets.

The NEBF filter has revealed great effectiveness in removing isolated exudates. However, short vessels passing through an exudate may be lost due to the NEBF filter application, thus preventing the estimation of biomarkers such as tortuosity. This problem could be solved by taking into account connectivity and shape information or employing a more accurate exudate detection technique.

Our results suggest that joint detection of vessel and other retinal structures combined together could finally solve the problem of accurate and reliable retinal vessel segmentation suitable to various practical scenarios. In future work, we plan to extend this idea to other structures such as drusen and haemorrhages to make vessel segmentation even more robust against false positives generated by the presence of such structures.

REFERENCES

[1] H. Li, W. Hsu, M. L. Lee, and T. Y. Wong, "Automatic grading of retinal vessel caliber," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 7, pp. 1352–1355, Jul. 2005.

[2] J. J. Capowski, J. A. Kylstra, and S. F. Freedman, "A numeric index based on spatial frequency for the tortuosity of retinal vessels and its application to plus disease in retinopathy of prematurity," *Retina*, vol. 15, no. 6, pp. 490–500, 1995.

[3] A. Lisowska, R. Annunziata, G. K. Loh, D. Karl, and E. Trucco, "An experimental assessment of five indices of retinal vessel tortuosity with the ret-tort public dataset," in *Proc. 36th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Aug. 2014, pp. 5414–5417.

[4] 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 *Proc. 1st Int. Workshop Ophthalmic Med. Image Anal.*, 2014, pp. 113–120.

[5] T. Y. Wong, A. Shankar, R. Klein, B. E. Klein, and L. D. Hubbard, "Prospective cohort study of retinal vessel diameters and risk of hypertension," *BMJ*, vol. 329, no. 7457, p. 79, 2004.

[6] T. Y. Wong, F. A. Islam, R. Klein, B. E. Klein, M. F. Cotch, C. Castro, A. R. Sharrett, and E. Shahar, "Retinal vascular caliber, cardiovascular risk factors, and inflammation: The multi-ethnic study of atherosclerosis (MESA)," *Investigative Ophthalmology Visual Sci.*, vol. 47, no. 6, pp. 2341–2350, 2006.

[7] M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A. Rudnicka, C. Owen, and S. Barman, "Blood vessel segmentation methodologies in retinal images: A survey," *Comput. Methods Programs Biomed.*, vol. 108, no. 1, pp. 407–433, Oct. 2012.

[8] A. Criminisi, P. Perez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1200–1212, Sep. 2004.

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

[10] J. Soares, J. Leandro, R. Cesar, H. Jelinek, and M. Cree, "Retinal vessel segmentation using the 2-D gabor wavelet and supervised classification," *IEEE Trans. Med. Imag.*, vol. 25, no. 9, pp. 1214–1222, Sep. 2006.

[11] E. Trucco, L. Ballerini, D. Relan, A. Giachetti, T. MacGillivray, K. Zutis, C. Lupascu, D. Tegolo, E. Pellegrini, G. Robertson, P. Wilson, A. Doney, and B. Dhillon, "Novel VAMPIRE algorithms for quantitative analysis of the retinal vasculature," in *Proc. Biosignals Biorobotics Conf.*, Feb. 2013, pp. 1–4.

[12] E. Ricci and R. Perfetti, "Retinal blood vessel segmentation using line operators and support vector classification," *IEEE Trans. Med. Imag.*, vol. 26, no. 10, pp. 1357–1365, Oct. 2007.

[13] D. Marin, A. Aquino, M. E. Gegundez-Arias, and J. M. Bravo, "A new supervised method for blood vessel segmentation in retinal images by using gray-level and moment invariants-based features," *IEEE Trans. Med. Imag.*, vol. 30, no. 1, pp. 146–158, Jan. 2011.

[14] M. M. Fraz, P. Remagnino, A. Hoppe, B. Uyyanonvara, A. R. Rudnicka, C. G. Owen, and S. A. Barman, "An ensemble classification-based approach applied to retinal blood vessel segmentation," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 9, pp. 2538–2548, Sep. 2012.

[15] R. Annunziata, A. Kheirkhah, P. Hamrah, and E. Trucco, "Boosting hand-crafted features for curvilinear structures segmentation by learning context filters," in *Proc. Med. Image Comput. Comput. Assisted Interventions*, 2015, accepted for publication.

[16] A. Hoover, V. Kouznetsova, and M. Goldbaum, "Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response," *IEEE Trans. Med. Imag.*, vol. 19, no. 3, pp. 203–210, Mar. 2000.

[17] X. Jiang and D. Mojon, "Adaptive local thresholding by verification-based multithreshold probing with application to vessel detection in retinal images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 1, pp. 131–137, Jan. 2003.

[18] A. Mendonca and A. Campilho, "Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction," *IEEE Trans. Med. Imag.*, vol. 25, no. 9, pp. 1200–1213, Sep. 2006.

[19] M. Martinez-Perez, A. Hughes, S. Thom, A. Bharath, and K. Parker, "Segmentation of blood vessels from red-free and fluorescein retinal images," *Med. Image Anal.*, vol. 11, no. 1, pp. 47–61, 2007.

[20] A. Frangi, W. Niessen, K. Vincken, and M. Viergever, "Multiscale vessel enhancement filtering," in *Proc. Med. Image Comput. Comput.-Assisted Interventions*, 1998, pp. 130–137.

[21] H. Yu, S. Barriga, C. Agurto, G. Zamora, W. Bauman, and P. Soliz, "Fast vessel segmentation in retinal images using multiscale enhancement and second-order local entropy," *Proc. SPIE*, vol. 8315, pp. 83151B-1–83151B-12, Feb. 2012.

[22] G. Azzopardi, N. Strisciuglio, M. Vento, and N. Petkov, "Trainable COSFIRE filters for vessel delineation with application to retinal images," *Med. Image Anal.*, vol. 19, no. 1, pp. 46–57, Jan. 2015.

[23] R. Annunziata, P. Kheirkhah, Hamrah, and E. Trucco, "Scale and curvature invariant ridge detector for tortuous and fragmented structures," in *Proc. Med. Image Comput. Comput. Assisted Interventions*, 2015, accepted for publication.

[24] B. Lam and H. Yan, "A novel vessel segmentation algorithm for pathological retina images based on the divergence of vector fields," *IEEE Trans. Med. Imag.*, vol. 27, no. 2, pp. 237–246, Feb. 2008.

[25] B. Lam, Y. Gao, and A.-C. Liew, "General retinal vessel segmentation using regularization-based multiconcavity modeling," *IEEE Trans. Med. Imag.*, vol. 29, no. 7, pp. 1369–1381, Jul. 2010.

[26] C. Sinthanayothin, J. F. Boyce, T. H. Williamson, H. L. Cook, E. Mensah, S. Lal, and D. Usher, " Automated detection of diabetic retinopathy on digital fundus images," *Diabetic Med.*, vol. 19, no. 2, pp. 105–112, Feb. 2002.

[27] T. Walter, J.-C. Klein, P. Massin, and A. Erginay, "A contribution of image processing to the diagnosis of diabetic retinopathy-detection of exudates in color fundus images of the human retina," *IEEE Trans. Med. Imag.*, vol. 21, no. 10, pp. 1236–1243, Oct. 2002.

[28] C. I. Sánchez, M. Niemeijer, I. Isgum, A. Dumitrescu, M. S. A. Suttorp-Schulten, M. D. Abrámoff, and B. van Ginneken, " Contextual computer-aided detection: Improving bright lesion detection in retinal images and coronary calcification identification in CT scans," *Med. Image Anal.*, vol. 16, no. 1, pp. 50–62, Jan. 2012.

[29] J. Odstrcilik, R. Kolar, A. Budai, J. Hornegger, J. Jan, J. Gazarek, T. Kubena, P. Cernosek, O. Svoboda, and E. Angelopoulou, "Retinal vessel segmentation by improved matched filtering: Evaluation on a new high-resolution fundus image database," *IET Image Process.*, vol. 7, no. 4, pp. 373–383, 2013.

[30] J. Serra, *Image Analysis and Mathematical Morphology*. Orlando, FL, USA: Academic, 1983.

[31] N. M. Salem and A. Nandi, "Segmentation of retinal blood vessels using scale-space features and k-nearest neighbour classifier," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2006, vol. 2, p. II.

[32] M. Al-Rawi and H. Karajeh, "Genetic algorithm matched filter optimization for automated detection of blood vessels from digital retinal images," *Comput. Methods Programs Biomed.*, vol. 87, no. 3, pp. 248–253, Sep. 2007.

[33] B. Al-Diri, A. Hunter, and D. Steel, "An active contour model for segmenting and measuring retinal vessels," *IEEE Trans. Med. Imag.*, vol. 28, no. 9, pp. 1488–1497, Sep. 2009.

Authors' photographs and biographies not available at the time of publication.