

Accurate Vessel Segmentation with Progressive Contrast Enhancement and Canny Refinement

Xin Yang^{1,2}, K.T.Tim Cheng² and Aichi Chien³

¹Dept. of Electronics Information Engineering, HUST, Wuhan, China

²Dept. of Electrical Computer Engineering, UCSB, CA, USA

³Division of Interventional Neuroradiology, UCLA, Medical School, CA, USA

Abstract. Vessel segmentation is a key step for various medical applications, such as diagnosis assistance, quantification of vascular pathology, and treatment planning. This paper describes an automatic vessel segmentation framework which can achieve highly accurate segmentation even in regions of low contrast and signal-to-noise-ratios (SNRs) and at vessel boundaries with disturbance induced by adjacent non-vessel pixels. There are two key contributions of our framework. The first is a progressive contrast enhancement method which adaptively improves contrast of challenging pixels that were otherwise indistinguishable, and suppresses noises by weighting pixels according to their likelihood to be vessel pixels. The second contribution is a method called canny refinement which is based on a canny edge detection algorithm to effectively re-move false positives around boundaries of vessels. Experimental results on a public retinal dataset and our clinical cerebral data demonstrate that our approach outperforms state-of-the-art methods including the vesselness based method [1] and the optimally oriented flux (OOF) based method [2].

1 Introduction

The segmentation of vascular structures plays a significant role in diagnosis assistance, quantification of vascular pathologies, treatment and surgery planning. For instance, segmenting arteries and their bifurcations in the Circle of Willis, and quantifying their changes over a span of time can facilitate cerebral aneurysm detection and development analysis. In neurosurgical procedures, vessels, giving indication of where the blood supply of a lesion is drawn from and drained to, often serve as landmarks and guidelines to the lesion during surgery. The more accurate the vascular segmentation is, the more precise a computer-guided procedure can be made.

With growing streams of data generated by modern imaging modalities, such as computed tomography angiography (CTA) and magnetic resonance angiography (MRA), automatic vessel segmentation to minimize laborious and error-prone manual operations is in great demand. There have been numerous dedicated research efforts on this subject over years. Some of most successful ones apply filters (e.g. Hessian-based filters [1], optimally oriented flux (OOF) [2],

steerable filters [19], and learned filters [3, 4, 5, 6, 7, 8]) to individual pixels and classify a pixel as a part of a vessel or not based on its filter response. However, these filters mainly rely on image gradients or high-order derivatives, thus they can hardly provide accurate responses at regions with very low contrast and a poor signal-to-noises ratio (SNR). The top row of Figs. 1 (b) and (c) display a sub-region of a retinal vessel image of Fig. 1 (a) and a contrast-enhanced version, respectively. Due to low image contrast, several small vessels in Fig. 1 (b)-top can barely be distinguished. Applying contrast enhancement to this region could slightly improve the visibility of small vessels while greatly increase noises resulting in a low SNR (as shown in Fig. 1 (c)-top). As a result, most existing methods fail to achieve a high true-positive rate and a low false-positive rate in those regions (as shown in Fig. 1 (d)-top). Another limitation of existing methods is that vascular filters usually give similarly weak responses for pixels around vascular borders, either vessel or non-vessel pixels, resulting in inaccuracy in localizing the true boundary of a vessel tube. As shown in Fig. 1 (d)-bottom, most pixels in the neighborhood of vessel boundaries are incorrectly classified (as denoted in red) which could result in inaccurate quantification of vascular pathologies and diagnosis.

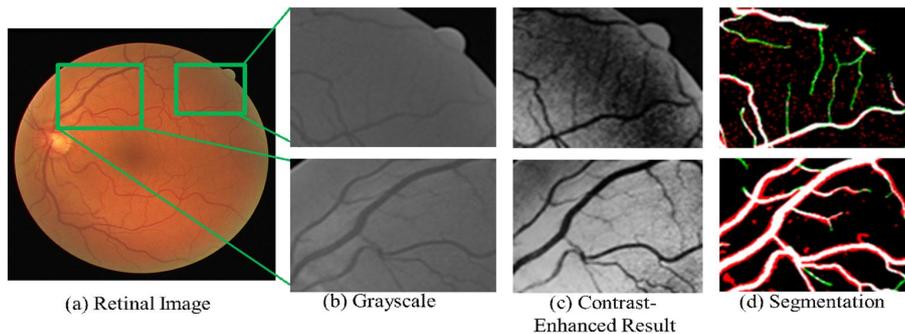


Fig. 1. Limitations of existing vessel segmentation methods. (a) An exemplar retinal image. (b) and (c) are grayscale images of two sub-regions and their contrast-enhanced results respectively. (d) Segmentation results for the contrast-enhanced images based on vesseness (i.e. the Frangi’s method [1]), one of the most popular methods for vessel segmentation. White, green and red colors indicate true positives, false negatives and false positives, respectively. Two major limitations of the Frangi’s method can be observed: 1) for regions with a low contrast and SNR, it fails to detect most of small vessels (green pixels in (d-top)) and incorrectly classifies many noises as vessels (red pixels in (d-top)); and 2) it fails to precisely localize boundaries of vessels (d-bottom). Although we use results of the Frangi’s method for illustration, these two limitations are common for most existing methods.

In this paper, we present an automatic vessel segmentation framework, with the primary focus on achieving high accuracy in two challenging scenarios: in

regions with low contrast and low SNR and at vessel boundaries. Specifically, there are two main contributions of the proposed framework:

1. We propose a progressive contrast enhancement method that iteratively excludes a subset of pixels, which have been identified as vessel pixels with high confidence in previous iterations, from contrast enhancement in the next iteration. Comparing to existing methods which process all pixels within a particular region, the proposed approach, adjusting the contrast only for the remaining pixels in each iteration, places more emphasis on challenging pixels which are difficult to be classified in previous iterations. As a result, our approach can better capture subtle vessel information in low contrast regions. To further suppress noises in low SNR regions, we weight the intensity of every pixel based on a function of shape responses to reduce the impact of noises in the contrast enhancement procedure. The idea behind this strategy is that the shape information is complementary to the intensity information and it is less likely that a non-vessel pixel with high noise could have both its shape response and its intensity value similar to those of a vessel pixel.

2. We propose a simple yet effective method, called canny refinement, for precisely localizing vessel pixels, particularly at vessel boundaries. Our method employs canny edge detection to identify pixels on the boundaries of vessels. Then a ro-bust and effective function is designed based on canny edges to determine whether a pixel is between two boundaries of a vessel or is outside a vessel. Based on the output of the function, the system method can refine the filtering results and minimize false positives which are outside a vessel. The rest of the paper is organized as follows. Section 2 reviews the related work. Section 3 presents details of the proposed method. In Section 4, we compare the performance of our method with two state-of-the-art methods. Section 5 concludes the paper.

2 Related Work

The broad application of vessel segmentation has stimulated the development of several categories of approaches, each of which has distinct strengths. Active contour within the level set framework [11, 12, 13], which is capable of handling topology changes and is adaptable to shapes of complex vessel structures, has proven to be effective for vessel segmentation. Several enhancements have been made for further performance improvement. Most recent efforts [9, 10] have been focusing on simplifying and automating the parameter settings to achieve optimized performance for a wide range of data content and quality.

Another category of approaches applies vessel enhancement filters to individual pixels and then classifies each pixel, as either a vessel or a non-vessel pixel, by thresholding the filtering score [1, 2, 3, 5, 6, 7, 8, 20, 21]. Our framework belongs to this category. A number of vessel enhancement filters have been developed in recent years. Some of them utilize the second-order derivatives to distinguish specific tubular shape of vessels, which have a locally prominent low curvature orientation (i.e. the vessel direction) and have planes of a high intensity cur-

vature (i.e. the cross-sectional planes) [1, 20, 21, 22, 23]. The Hessian matrix is the most common tool to capture tubular structure information. Eigenvalues of the Hessian matrix can discriminate between plane-, blob- and tubular-like structures, and corresponding eigenvectors indicate the vessel orientations. A representative example of the Hessian-matrix based method is the vesselness filter proposed in Frangi et al. [1] which has been widely used in practice, owing to its intuitive geometric formulation. The Weingarten matrix is a less popular alternative to the Hessian matrix. Filters based on the Weingarten matrix include those proposed in [22] and [23].

Instead of analyzing the second-order derivatives, another category of methods exploit the local distribution of the gradient vectors. For instance, the method in [3] analyzes the eigenvalues of the gradient vectors' covariance matrix. Bauer and Bischof [24] leveraged a vector field obtained from the gradient vector flow (GVF) diffusion. Law and Chung proposed the use of optimally oriented flux (OOF) [2] which relies on the measure of gradient flux through the boundary of local spheres. Comparing to the Hessian-based filters, OOF could be more accurate and less sensitive to disturbances from adjacent structures.

It has been pointed out in recent literature [5, 6, 7, 8] that real vascular structures, which do not necessarily conform to an ideal tubular shape model, can drastically impact the performance of methods relying on handcrafted shape filters. Several efforts have been made to learn filters to describe convoluted appearances and structures of vessels. For instance, Agam et al. [3] estimated the eigenvalue distribution of the gradient vectors' covariance matrix via Expectation Maximization. Support Vector Machines operating on the Hessian's eigenvalues have been used to discriminate between vascular and nonvascular pixels [4]. In [8], rotational features were computed at each pixel using steerable filters and fed to an SVM to classify pixels as vessel pixels or not. Inspired by [8], a series of improvements [5, 6, 7] were made which include more filters (i.e. vesselness [1] and OOF [2]), in addition to the steerable filters, and leverage more advanced machine learning techniques. A comprehensive survey of vessel segmentation methods can be found in [15, 16]. The problem, however, is that both handcrafted and learned filters mainly rely on image gradients or high-order derivatives, thus their responses are sensitive to noises and often too weak to discriminate vascular and nonvascular pixels in low contrast regions. Today's angiograms inevitably contain noises and exhibit inhomogeneous contrast. The intensity of some vessels (particularly narrow vessels) could differ from the background by as little as four grey levels, yet the standard deviation of background noise is around 2.3 grey levels. As a result, most, if not all, existing filters are ineffective in low contrast and/or low SNR regions. In addition, vascular filters usually produce weak responses around vascular borders, yielding difficulties in precisely localizing the exact boundary of a vessel tube. Imprecise boundary localization could consequently result in inaccurate quantification of pathologies and diagnosis. This paper focuses on addressing these two challenging problems. Specifically, we proposed two techniques: progressive contrast enhancement and canny refinement, which can be used together with existing filtering based

methods and greatly boost their segmentation performance in low contrast, low SNR regions and at vascular boundaries.

3 Our Method

Fig. 2 illustrates our vessel segmentation framework, which consists of three main components: vessel enhancement filtering (the orange block), canny refinement (the green block) and progressive contrast enhancement (the blue blocks). Given an input image, vessel enhancement filtering is first applied to every image pixel to obtain the likelihood of each pixel being a vessel pixel. In our implementation, we employ vesselness [1] and OOF filters [2], which are known as two of the best filters to date. Canny refinement is then applied to revise the filtering results: the filtering responses of those pixels classified to be outside a vessel by canny refinement are adjusted to zero (i.e. non-vessel pixel with the highest confidence). Based on the revised responses, pixels which can be classified with high confidence as either vessel or background pixels are added to the final segmentation results and removed from the image. The method adjusts the contrast of the remaining pixels by shape-weighted contrast enhancement and then restarts the above-mentioned procedure on the remaining pixels. Such procedure repeats until no more fine vessels can be detected or the number of iterations reaches a limit. In the following, we provide technical details and describe strengths of canny refinement and progressive contrast enhancement.

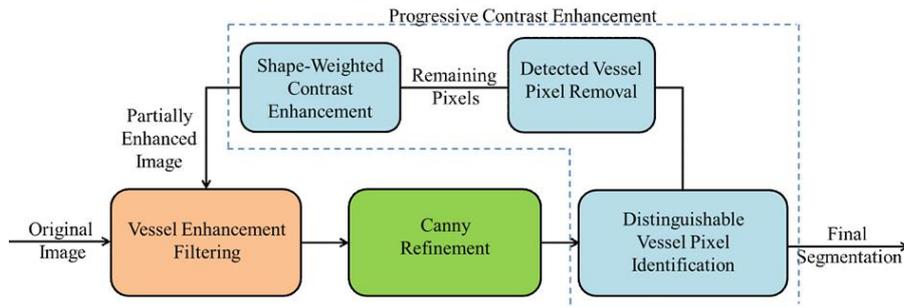


Fig. 2. Framework of vessel segmentation with progressive contrast enhancement and canny refinement

3.1 Canny Refinement

Canny [25] has been widely regarded as the best solution for robust edge detection and precise localization of edge pixels. Vascular boundaries which can be approximated as step edges should be accurately localized by canny. Fig. 3 (a) displays the vessel segmentation based on vesselness filtering overlaid with

detected canny edges (the blue pixels). Clearly, many canny edge pixels correctly locate at real vessel boundaries, forming "classification planes" which separate true positives (the white pixels) from false positives (the red pixels).

However, canny provides only the location of edges but could not determine whether a pixel adjacent to an edge is inside or outside of a vessel tube. Therefore, solely relying on the edge location cannot remove false positives. To address this problem, we construct a verification map based on canny edges. Each entry of the map is a value of quadruples 1, 0, -1, null (as shown in Fig. 3 (b)), i.e. 1 (green) and -1 (white) indicate pixels inside and outside a vessel tube, respectively. A 0 (red) de-notes a pixel at the boundary and null (back) indicate pixels far from any edges and thus are unnecessary to be examined in the current iteration. Based on the verification map, the method can refine the filtering results, i.e. pixels with small filtering response values and are labeled as -1 in the verification map are re-classified as negatives.

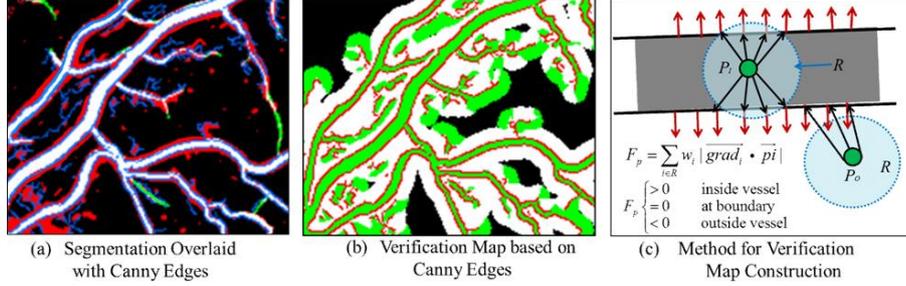


Fig. 3. (a) Segmentation based on vesselness filtering overlaid with detected canny edges. White, red and blue colors indicate true positives, false positives and canny edge pixels respectively. (b) Verification map based on canny edges. Green, white and red colors denote pixels reside inside a vessel, outside a vessel and at vessel boundary respectively. (c) Illustration of our method for verification map construction.

We design a robust and effective method to construct the verification map. The key idea of our method is outlined as follows. For every non-edge pixel P , we construct a vector from P to a nearby canny edge pixel E . Then we compute the dot product between the gradient orientation vector of pixel E . If P resides inside a vessel and vessel pixels are generally darker than the background, then the dot product is greater than zero; otherwise, the dot product is negative. Based on the sign of the dot product, we can determine whether a pixel is inside or outside a vessel. A verification map based on a single edge pixel is usually sensitive to noises. To improve the robustness, for every pixel P we consider a set of the canny edge pixels near P and sum up the weighed dot product according to Eq. (1) (as illustrated in Fig.3 (c)).

$$F_p = \sum_{E_i \in R} w_{E_i} \left| \overrightarrow{\text{grad}}_{E_i} \bullet \overrightarrow{PE_i} \right| \quad (1)$$

The weights w_{E_i} is sampled from a Gaussian distribution centered at P as Eq. (2), where (x_P, y_P) , (x_{E_i}, y_{E_i}) are coordinates of pixels P and E_i , and σ is the standard deviation of the Gaussian distribution.

$$w_{E_i} = \exp\left(\frac{(x_P - x_{E_i})^2 + (y_P - y_{E_i})^2}{2\sigma^2}\right) \quad (2)$$

Based on F_P we can construct a verification map V_p according to Eq. (3).

$$V_p = \begin{cases} 1, & F_p > 0 \text{ and } P \notin E \\ 0, & P \in E \\ -1, & F_p < 0 \text{ and } P \notin E \\ \text{null}, & R_p \cap E = \emptyset \end{cases} \quad (3)$$

3.2 Progressive Contrast Enhancement

In this section, we first briefly overview conventional contrast enhancement approaches and their limitations for vessel segmentation, followed by details of our progressive contrast enhancement method.

Histogram Equalization for Contrast Enhancement Heterogeneous contrast, resulting from the contrast agent inhomogeneity, noises and image artifacts, is a common problem in many medical image modalities. Histogram equalization is a common tool to increase contrast by stretching out the overall intensity range of an image. More specifically, it maps one distribution (i.e. original histogram of a given image) to another distribution (i.e. a wider and more uniform distribution of intensity values) based on a transformation function so that the intensity values can spread over the entire range. The transformation function is built based on the cumulative distribution function (CDF) defined as Eq. (4),

$$\text{cdf}_x(i) = \sum_{j=0}^i p_x(j), \quad 0 \leq i \leq L \quad (4)$$

where $p_x(i)$ is the probability of an occurrence of gray level i in image x , L is the total number of gray levels in the image (typically 256). The desired image y should have a flat histogram with a linearized CDF across the entire range, for a constant K .

$$\text{cdf}_y(i) = iK \quad (5)$$

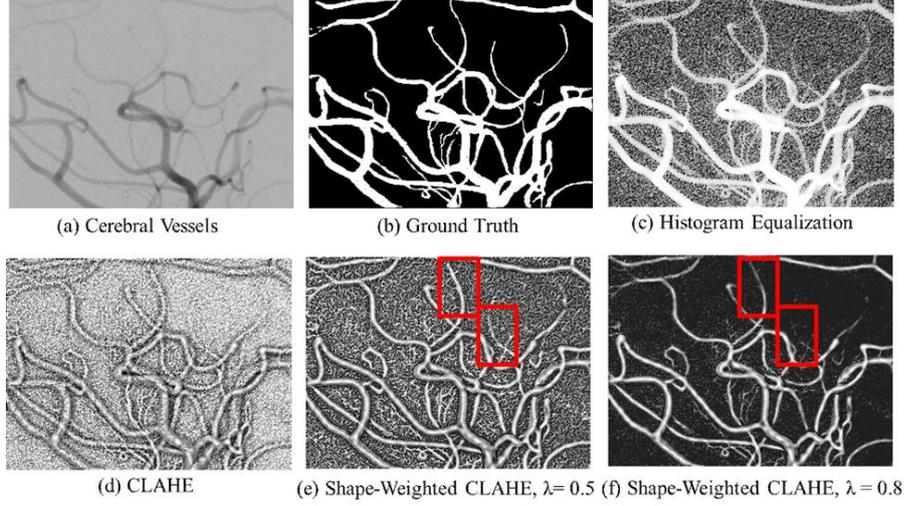


Fig. 4. Illustration of contrast enhancement results based on different methods. (a) Original image with cerebral vessels. (b) Ground truth obtained by manual label. (c) - (d) Contrast enhancement results based on histogram equalization and CLAHE with clip limit being 30, region size being 5050. (e) - (f) Contrast enhancement results based on shape-weighted CLAHE with λ being 0.5 and 0.8, respectively. Shape information is obtained by vesselness filtering.

According to Eqs. (4) and (5), the intensity transformation function can be derived as

$$T(i) = \text{round} \left(\frac{cdf_x(i) - cdf_{x_{min}}}{(M \times N) - cdf_{x_{min}}} \times (L - 1) \right) \quad (6)$$

where $M \times N$ gives the total number of pixels in image x .

However, histogram equalization often fails to provide satisfactory results for medical images with inhomogeneous contrast. Regions that are much lighter or darker than the rest of the image cannot be sufficiently enhanced. In addition, it could over-amplify noises in relatively homogeneous regions. Fig. 4 (c) displays the contrast-enhanced result for Fig. 4 (a) based on histogram equalization. Clearly, background noises are greatly amplified. Contrast Limited Adaptive Histogram Equalization (CLAHE) [15] is a popular solution to address these problems. It adjusts contrast locally by deriving a local transformation function from a neighborhood region of each pixel, and clips the histogram at a predefined value before computing the CDF to prevent over-amplification of noises. However, for images with a very low SNR, CLAHE still cannot effectively suppress noises, resulting in a noisy background and rough vessel boundaries (as shown in Fig.4 (d)). More importantly, although CLAHE performs equalization locally, it is inevitable that a local region contains both large vessels with good contrast to the background and fine vessels with low contrast (as shown in Fig. 1 (a))-

top). For these regions, results are usually dominated by large vessels, resulting in insufficient enhancement for small vessels (as shown in Fig. 1 (b)-top). Our progressive contrast enhancement can more successfully suppress noises through weighting each pixel’s intensity by its shape filtering response and focus mainly on enhancing contrast of challenging pixels (i.e. small vessel pixels) in each iteration.

Shape-Weighted Contrast Enhancement CLAHE utilizes only the intensity of an image, thus is very sensitive to noises which have similar intensity values as vessels. The local geometric structure around each pixel is a discriminating feature useful for distinguishing vessel pixels from random noises. In addition, the local shape information is complementary to the intensity information, thus it is less likely that noise pixels have similar values of both shape responses and intensity as vessel pixels. Based on these observations, we propose to use local shape information $S(x)$ obtained from vessel enhancement filtering (i.e. vesselness or OOF) to weight the corresponding pixel intensities $I_{norm}(x)$ before performing CLAHE, as shown in Eq. (7). Intensities are normalized to the range of $[0, 1]$ and for images in which vessel pixels are darker than the background, we reverse each pixel’s intensity value by subtracting the original normalized intensity from 1. We set a parameter λ to adjust the impact of the weights. Larger results in greater impact of weighting and vice versa.

$$I_{new}(x) = \begin{cases} I_{norm}(x) \times (S(x))^\lambda, & \text{vessel pixels are brighter than background} \\ (1 - I_{norm}(x)) \times (S(x))^\lambda, & \text{vessel pixels are darker than background} \end{cases} \quad (7)$$

Figs. 4 (e) and (f) show the enhanced results based on shape-weighted CLAHE with λ being 0.5 and 0.8, respectively. Pixels in homogeneous regions generally have small shape responses and thus most noises in those regions can be prevented from being amplified. Increasing λ from 0.5 to 0.8 can suppress more noises, while may also decrease the contrast for regions around small vessels (as indicated by red rectangles in Figs. 4 (e) and (f)), resulting from inaccurate shape responses due to low contrast at these regions. In our implementation, we set λ to 0.8 as the default value which produced the best empirical results.

Progressive Contrast Enhancement on Challenging Pixels It’s common that both large and fine vessels appear in the same regions (as illustrated by red rectangles in Fig. 5). Within such a region, large vessels have better contrast to the background (either darker or brighter) than small vessels. As a result, the intensity range is often dominated by large vessels, resulting in insufficient enhancement for small vessels in such a region. Reducing the region size of CLAHE could help limit the size differences in a region, but may also reduce the robustness of CLAHE.

To address this problem, we propose to progressively increase contrast for vessels of different sizes. In each iteration, we detect distinguishable vessel pixels

that can be easily classified as vessels with high confidence and remove them from further consideration in future iterations. In each iteration, shape-weighted CLAHE is applied only to those remaining pixels, which usually contain smaller vessels which have not been detected in previous iterations. After the contrast enhancement to this subset of pixels in the image, more pixels in fine vessels can be detected and removed from consideration in future iterations. We repeat this iterative procedure until no more fine-vessel pixels can be detected or a limit on the iteration count is reached.

To classify vessel and background pixels with high confidence in each iteration, we set two strict thresholds. That is, pixels whose filtering responses are greater than a high threshold T_H or smaller than a low threshold T_L are classified as vessel and back-ground pixels respectively; the remaining pixels whose responses fall within the range of T_H and T_L are considered as unknown and their labeling will be done in future iterations. To guarantee that the true-positive and true-negative rates are both high for the classification, we exploited multiple settings of the parameters for running CLAHE to generate several enhanced results and performed vessel enhancement filtering and classification for every resulting image. Pixels which are classified as vessel pixels in all of the resulting images are considered as robust and distinguishable vessel pixels. They are then labeled as vessel pixels and excluded from consideration in future iterations.

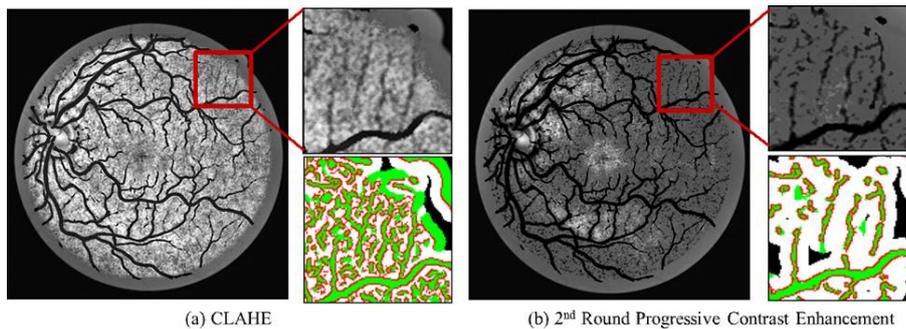


Fig. 5. Illustration of contrast enhancement results based on (a) CLAHE and (b) the 2^{nd} round progressive contrast enhancement. A region including both challenging pixels (i.e. small vessels) and part of a larger vessel are highlighted by a red rectangle. Its enlarged version and the verification map are displayed on the right. Clearly, our progressive contrast enhancement can provide much fewer noises, better visibility of small vessels and hence a more accurate verification map for small vessels. The parameter settings for CLAHE are the same for both (a) and (b), i.e. clip limit = 30 and region size = 5050.

Fig. 5 compares the contrast enhancement results based on CLAHE and the 2nd round progressive contrast enhancement, respectively. We highlight a small region containing both small vessels and part of a larger vessel by a red rectangle

and display the verification map of this region. Clearly, the result obtained by our progressive contrast enhancement provides better visibility of small vessels and contains fewer noises. As a result, the verification map is more accurate than that of CLAHE.

4 Experimental Results

In this section, we provide quantitative evaluation of our method using a public retinal dataset, DRIVE [18], and clinical cerebral data. We first describe our datasets and the evaluation metric, followed by the results and analysis.

4.1 Datasets

DRIVE [18] is a public-available dataset of 2D RGB retinal scans to enable comparative studies on segmentation of blood vessels in retinal images. Each image was captured using 8 bits per color plane at 768×584 pixels and was JPEG compressed. The entire dataset contains 40 images which are divided into a training set and a testing set, both containing 20 images. For each testing case, two ground truths obtained by manual segmentation are provided. We test our approach on the testing images of DRIVE. Figs. 7 (a) and (b) show an exemplar image and one ground truth image from DRIVE.

We also evaluate our method on 2D clinical cerebral vessel data, approved by [re-moved for anonymous submission]. The image was obtained by digital subtraction angiography (DSA), represented using 8bits grayscale TIFF format, including 560×414 pixels. For quantitative evaluation, we asked two experts to manually label the image, yielding two ground truth images. Figs. 8 (a) and (b) illustrate our cerebral vessel image and one of its ground truth images. The primary focus of this paper is to improve the segmentation performance in challenging scenarios. Thus, the ground truth data must be able to facilitate evaluation and performance comparison for the challenging cases. For this purpose, we further divide all pixels in each ground truth image into two parts: vessel pixels which can be correctly classified by all baseline methods we have implemented are labeled as easy pixels and pixels which are incorrectly labeled by at least one baseline method are marked as challenging pixels. Specifically, we implemented two baseline methods: vesselness based method and OOF based method (details about baselines can be Sec. 4.3). The threshold for binary classification is adjusted so that the precision is above 95%. Figs. 7 (c) and (d), and Figs. 8 (c) and (d) illustrate the easy and challenging vessel pixels in the ground truth images, respectively. Obviously, challenging pixels are mainly located around vessel boundaries and at small vessels which have very low contrast to its surrounding background.

4.2 Evaluation Metric

We use recall and precision to evaluate the segmentation performance. Recall is defined as the number of true positives which are identified as vessel pixels

in both ground truth and segmented image divided by the total number of vessel pixels in the ground truth. Precision is defined as the number of true positives divided by the total number of pixels that are identified as vessel pixels in segmented images. As mentioned in Section 4.1, we focus our evaluation on challenging pixels, thus we exclude easy vessel pixels from the precision-recall calculation, as shown in Eqs. (8) and (9),

$$Recall = \frac{TP - TP_{easy}}{\text{challenging vessel pixels in ground truth}} \quad (8)$$

$$Precision = \frac{TP - TP_{easy}}{\text{challenging vessel pixels in segmented image}} \quad (9)$$

We plot the Recall-Precision curve to demonstrate the overall segmentation performance when varying the threshold parameter for binary classification. The larger is the area under the curve, the better the performance of the method.

4.3 Experimental Setup

Any existing vessel enhancement filter can be used in the first step of our segmentation framework (the orange block of Fig. 2). In this work, we experimented with two filters: multi-scale vesselness and multi-scale OOF, due to their widely-acknowledged good performance for delineating tubular structures. We utilized the ITK implementation for multi-scale vesselness filtering and relied on [17] for the implementation of OOF. For each filter, we manually adjusted parameters to obtain the best performance and used the same parameter settings throughout the entire evaluation process. For both filters, we used identical parameters for multi-scale processing - the minimum and maximum standard deviations for Gaussian are set to 0.5 and 5, respectively, and the total number of scales is set to 10.

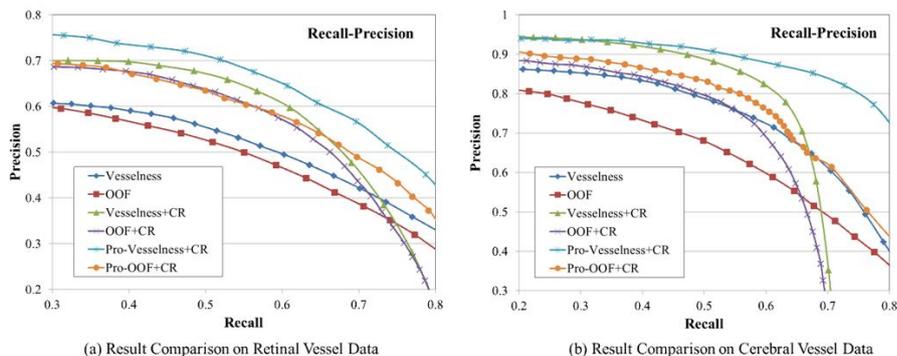


Fig. 6. Recall-Precision curves obtained for (a) the retinal vessel data and (b) the cerebral vessel data. Our method with canny refinement and progressive contrast enhancement outperforms vesselness and OOF over the entire range.

In our experiments, we compared six methods: vesselness, OOF, vesselness with canny refinement (i.e. vesselness+CR), OOF with canny refinement (OOF+CR), vesselness with the two proposed techniques (Pro-vesselness+CR), and OOF with the two proposed techniques (Pro-OOF+CR).

4.4 Results

Figs. 6 (a) and (b) show the comparison results on the retinal and cerebral data respectively. First, we evaluate the effectiveness of canny refinement. We compare the performance of vesselness and vesselness+CR, as shown by red and light green curves in Figs. 6 (a) and (b). When the recall is relatively small (e.g. below 75% in (a) and below 70% in (b)), canny refinement can greatly improve the precision by removing false positives arising from random noises and the disturbing objects adjacent to vessel boundaries. However, when the recall is greater than a certain value, canny refinement could adversely decrease the precision. This is mainly because canny edge detection may incur errors, e.g. missing true edges of vessels and mistakenly detecting edges on noises, especially in regions with poor contrast and a low SNR. Incorrect canny edges may lead to errors in the verification map, yielding incorrect removal of true vessel pixels. As a result, reducing the threshold cannot improve the recall any more while reduce the precision. Similar results can be observed by comparing the results of OOF (blue curves) and OOF+CR (purple curves) for both datasets.

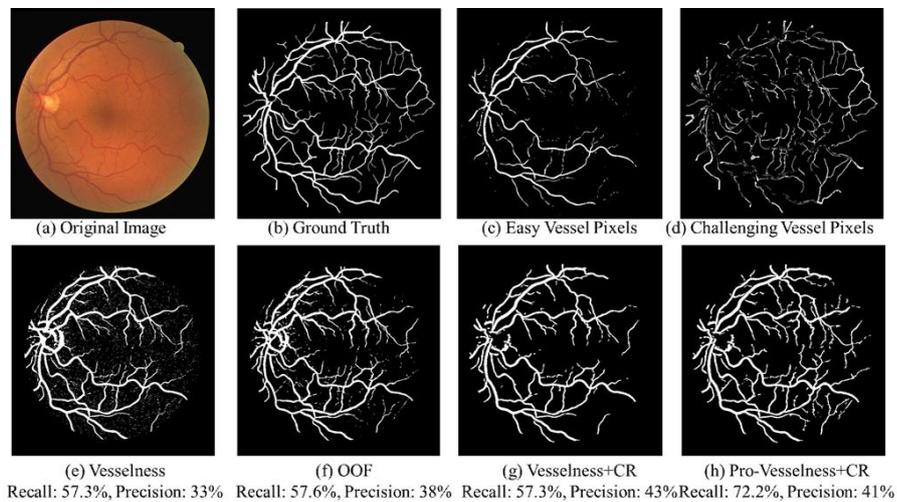


Fig. 7. Illustration of segmentation results on an exemplar retinal image. (a) Original image. (b) A ground truth image. (c) and (d) indicate easy and challenging vessel pixels on the ground truth image. (e) - (h) show the segmentation results of four methods. Pro-vesselness+CR achieves the best performance.

Next, we examine the effectiveness of progressive contrast enhancement. We compare the performance of vesselness (the red curves), vesselness+CR (the light green curves) and Pro-vesselness+CR (the light blue curves). Clearly, Pro-vesselness+CR outperforms the other two methods over the entire range. In particular, for large recall (i.e. greater than 70%) progressive contrast enhancement can help greatly boost the performance of vesselness+CR and maintain superior performance to vesselness. This result demonstrates that progressive contrast enhancement can effectively improve the contrast and SNR in low quality regions, and in turn increase the detection rate of vessels in those regions. Similar results can be also observed for methods based on OOF filters.

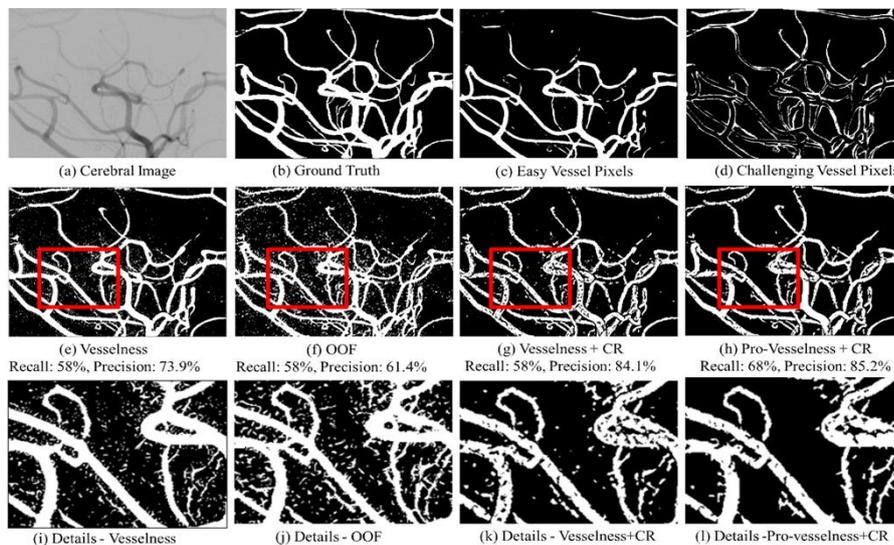


Fig. 8. Illustration of segmentation results on a cerebral image. (a) Original image. (b) A ground truth image. (c) and (d) indicate easy and challenging vessel pixels on the ground truth image. (e) - (h) display the segmentation results of four methods. (i)-(l) show segmentation details of a region. Pro-vesselness+CR achieves the best performance.

The second row of Figs. 7 and 8 illustrate the segmentation results for a retinal and a cerebral image respectively. We omit the results for OOF+CR and Pro-OOF+CR since their results are similar to those of vesselness+CR and Pro-vesselness+CR. For the first three methods ((e)-(g)), i.e. vesselness, OOF and vesselness+CR, we manually tune the threshold so that the recalls for all the three methods are similar (57% 58%). We then compare their precisions (as shown under the segmentation results). For the retinal image, vesselness+CR achieves 10% and 5% greater precision than vesselness and OOF respectively. For the cerebral image, it is 10% and 22.7% higher than those of vesselness and

OOF respectively. We further applied progressive contrast enhancement to the results of vesselness+CR (as shown in (h)), which improves the recall by another 10% 15% while maintaining the precision.

5 Conclusion

In this paper, we present a framework for accurate vessel segmentation in two challenging scenarios: in regions with poor contrast and a low SNR, and at vessel boundaries. We propose and validate two techniques: progressive contrast enhancement and canny refinement. Progressive contrast enhancement involves an iterative procedure where each iteration emphasizes only on challenging pixels (usually pixels of small vessels) which were not distinguishable in previous iterations. Experimental results demonstrate that by excluding large vessel pixels detected in previous iterations from contrast enhancement, small vessel pixels can be better highlighted by CLAHE. In addition, progressive contrast enhancement can effectively suppress noises spread in a homogeneous background by weighting pixels according to their shape responses.

This paper also demonstrates that canny refinement which constructs a verification map based on canny edges can successfully minimize false positives around boundaries of vessels. Experimental results on a retinal dataset and a cerebral data demonstrate that the two proposed techniques can greatly improve the performance of state-of-the-art filtering-based segmentation methods, such as vesselness and OOF.

Acknowledgement. This work was supported in part by the Society of Interventional Radiology (SIR) Foundation Dr. Ernest J. Ring Academic Development Grant, The Aneurysm and AVM Foundation (TAAF) Cerebrovascular Research Grant, and a UCLA Radiology Exploratory Research Grant.”

References

1. A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale Vessel Enhancement Filtering. In Proc. of MICCAI. Volume 1496 of Lecture Notes in Computer Science, 1998.
2. M. Law and A. Chung. Three Dimensional Curvilinear Structure Detection using Optimally Oriented Flux. In Proc. of ECCV, pp. 368-382, 2008.
3. G. Agam and C. Wu. Probabilistic Modeling-Based Vessel Enhancement in Thoracic CT Scans. In Proc. of CVPR, 2005.
4. A. Santamaría-Pang, C. M. Colbert, P. Saggau, and I. Kakadiaris. Automatic Center-line Extraction of Irregular Tubular Structures Using Probability Volumes from Multi-photon Imaging. In Proc. of MICCAI, 2007.
5. R. Rigamonti and V. Lepetit. Accurate and Efficient Linear Structure Segmentation by Leveraging ad hoc Features with Learned Filters. In Proc. of MICCAI. Volume 7510 of Lecture Notes in Computer Science, 2012.
6. C. Becker, R. Rigamonti, V. Lepetit, and P. Fua. Supervised Feature Learning for Curvilinear Structure Segmentation. In Proc. of MICCAI, 2013.

7. R. Rigamonti, A. Sironi, V. Lepetit, and P. Fua. Learning Separable Filters. In Proc. of CVPR, 2013.
8. G. Gonzalez, F. Fleuret, and P. Fua. Learning Rotational Features for Filament Detection. In Proc. of CVPR, 2009.
9. X. Yang, K.-T. Cheng, and A.-C. Chien. Geodesic Active Contour with Adaptive Configuration for Cerebral Vessel and Aneurysm Segmentation. In Proc. of ICPR, 2014.
10. P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D Active Contour Segmentation of Anatomical Structures: Significantly improved efficiency and reliability. *Neuroimage*, 2006 Jul. 1, 31(3):1116-28.
11. D. Nain, A.J.Yezzi, and G. Turk. Vessel Segmentation using a Shape Driven Flow. In Proc. of MICCAI. 2004. pp. 51-59.
12. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic Active Contours. *Int. J. Comput. Vis.* 1997. 22, 61-79.
13. S.C. Zhu and A. Yullie. Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation. *IEEE Trans. on Pattern Anal. Mach. Intell.* 1996. 18(9), 884-900.
14. D. Lesage, E.D. Angelini, I. Bloch, and G. Funka-Lea. A Review of 3D Vessel Lumen Segmentation Techniques: Models, Features and Extraction Schemes. 2009. *Med. Image Anal.*, vol. 13, pp: 819-845, 2009.
15. K. Zuiderveld. Contrast Limit Adaptive Histogram Equalization. *Graphics Gems IV*, pp: 474-485, 1994.
16. C. Kirbas and F. Quek. A Review of Vessel Extraction Techniques and Algorithms. *Journal of ACM Computing Surveys*, vol. 36 issue 2, pp: 81-121, June 2004.
17. Optimally Oriented Flux implementation. <https://github.com/fethallah/ITK-TubularGeodesics>.
18. J.J. Staal, M.D. Abramoff, M. Niemeijer, M.A. Viergever, and B. van Ginneken. Ridge based Vessel Segmentation in Color Images of the Retina. *IEEE Transactions on Medical Imaging*, vol. 23, pp. 501-509, 2004.
19. W.T. Freeman and E. H. Adelson. The Design and Use of Steerable Filters. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 13, pp. 891-906, 1991.
20. H. Shikata, E.A. Hoffman, M. Sonka. Automated Segmentation of Pulmonary Vascular Tree from 3D CT Images. In Proc. SPIE Med. Imaging, vol. 5369, pp. 107-116, 2004.
21. Q. Lin. Enhancement, Detection, and Visualization of 3D Volume Data. Ph.D. Thesis, Dept. EE, Linkoping University, SE-581 83 Linkoping, Sweden, Dissertations No. 824, May 2003.
22. N. Armande, P. Montesinos, Monga. A 3D thin nets extraction method for medical imaging. In Proc. of ICPR, p. 642, 1996.
23. V. Prinet, O. Monaga, C. Ge, S.L. Xie, S.D. Ma. Thin network extraction in 3D images: application to medical angiograms. In Proc. of ICPR, pp. 386-390, 1996.
24. C. Bauer, H. Bischof. A Novel Approach for Detection of Tubular Objects and Its Application to Medical Image Analysis. In Proc. DAGM Symp. Pattern Recognit. 2008 pp. 163-172.
25. J. Canny. A Computational Approach To Edge Detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8(6):679-698, 1986.