

AUTOMATIC SEGMENTATION OF CAPILLARY NON-PERFUSION IN RETINAL ANGIOGRAMS

Amit Agarwal, Jayanthi Sivaswamy

*CVIT, International Institute of Information Technology, Hyderabad, India.
amit@students.iiit.ac.in, jayanthi@iiit.net*

Alka Rani

*Aravind Eye Institute, Hyderabad, India.
alka_bala@rediffmail.com*

Taraprasad Das

*LV Prasad Eye Institute, Hyderabad, India.
tpd@lvpei.org*

Keywords: Capillary Non-Perfusion, Retina, Extrema Pyramid, Disease detection.

Abstract: Capillary Non-Perfusion (CNP) is a condition in diabetic retinopathy where blood ceases to flow to certain parts of the retina, potentially leading to blindness. This paper presents a solution for automatically detecting and segmenting CNP regions from fundus fluorescein angiograms (FFAs). CNPs are modelled as valleys, and a novel multiresolution technique for trough-based valley detection is presented. The proposed algorithm has been tested on 40 images and validated against expert-marked ground truth. Obtained results are presented as a receiver operating characteristic (ROC) curve. The area under this curve is 0.842 and the distance of ROC from the ideal point (0, 1) is 0.31.

1 INTRODUCTION

Diabetes is occurring in an ever increasing percentage in the world. Diabetes mellitus affects many organs of the body, and the eye is one of the organs that is affected relatively early (compared to the kidney). While diabetes affects all parts of the eye, the retina (retinopathy) is most commonly affected. Diabetic retinopathy progresses in phases. It starts with microaneurysms and superficial retinal hemorrhages (non-proliferative diabetic retinopathy; NPDR), progresses to accumulation of hard exudates in the posterior pole (diabetic maculopathy), and finally ends with new vessels in the surface of the retina and/or the optic disc (proliferative diabetic retinopathy; PDR). The underlying cause of the terminal event, the retinal new vessels, is retinal ischemia which manifests as areas of CNP that is most clearly seen in an FFA. These lesions appear as dark regions in the FFA images as shown in Fig. 1. If not treated in time, the CNP areas grow and spread across the entire retina. Large areas of non-perfusion lead to new vessel formation and bleeding into the vitreous cavity. These complications are responsible for severe visual loss in most patients with PDR (Kohner, 1993). FFA guides the choice and extent of laser treatment in diabetic maculopathy and PDR.

An automatic identification of important events in FFA is objective and very useful both for referral and treatment. Automated analyses of FFA images for the purpose of extracting important structures as well as lesions have received some attention. Image conditioning solutions that have been proposed include illumination correction using a parametric bi-cubic model for the illumination function (Cree et al., 1999) and noise suppression for a sequence of angiogram images based on bilateral filtering (Guo et al., 2005b). In FFA segmentation, stochastic models have been proposed to segment the fovea, arteries and veins from the central (macular) view of FFAs (Simó and de Ves, 2001) and among lesions, microaneurysms have received much attention. Several techniques ranging from morphological to model-based have been proposed for microaneurysm segmentation (Fleming et al., 2006), (Hafez, 2002) and (A. M. Mendonça, 1999). An automated technique for measurement of blood flow in capillaries has been attempted from angiograms, for determining the effect of cardio-pulmonary bypass surgery (Jagoe et al., 1992). The foveal region of the retinal image is processed to enhance the vascular structure and extract linear segments. The processed results from images taken before and after the bypass surgery are then compared (via a logical AND operation) to identify

the differences. However, to our knowledge, there are no reports in the literature of any technique to detect the cause of PDR namely the presence of the CNP regions anywhere in the retina. Detecting and segmenting CNPs is the focus of this paper.

The clinical procedure to detect CNPs is a visual scan of an FFA image. In order to estimate the amount of area damaged, the scan is generally done on the composite image of the retina obtained after suitable mosaicing of several retinal segments. Such a procedure suffers from several drawbacks: the variable skills and subjectivity of the observer, which also depend on the quality of the images; a lack of precise understanding of the area of retina affected which helps in deciding the nature and extent of laser treatment. Automated image analysis techniques can be used to address these issues but there are several challenges in devising solutions for CNP segmentation. FFAs suffer from non-uniform illumination due to the eye geometry, imaging conditions and presence of other media opacity such as cataract. Inter-patient and inpatient variability is also possible. The former is due to different pupil dilations and the latter is due to the time of image capture after injection of fluorescein dye. Another compounding factor is that the mean grey level of CNPs as well as their shape and size are variable, with the size ranging from very small to very large (from 100 to 55000 pixels). Often, the boundaries of CNPs are not well defined because of an inhomogeneous textured background. Thus, the only visually distinguishing characteristic of a CNP is that it is relatively darker than its surround.

In this paper, we propose a novel method to extract and quantify regions of CNP based on modeling CNPs as valleys in the image surface. The algorithm for CNP segmentation is developed and its details are presented in the next section. Section 3 provides implementation details and illustrative test results of the algorithm. Finally, some discussions and conclusions are presented in the last section.

2 VALLEY BASED CNP SEGMENTATION

2.1 Modelling CNP Regions

As discussed earlier, CNP occurs when the capillary network in a region of the human retina stops functioning and does not supply blood to the corresponding areas. In FFAs, regions receiving normal blood supply appear as bright white regions since they carry a fluorescent dye and regions lacking in blood (due

to abnormal supply of blood) appear as dark regions. Hence, regions of CNP appear as dull/dark lesions bounded by healthy vasculature.

A sample FFA image and an enlarged view of a CNP region and its surroundings is shown in Fig. 1. Also, included in this figure is the surface plot of the corresponding CNP region from which we can observe that the prominent vessels, the healthy capillary network and the CNP have very different topographic characteristics: While the major vessel appears as a ridge, the CNP appears as a valley with the healthy capillary network appearing as a plateau in the image. Hence, one can conclude that CNPs can be modelled as valleys. Watershed-based solution to valley detection (for example, (Gauch, 1999)) is possible, however, these result in oversegmentation or in the case of marker-based versions, require additional information. In the case of CNP detection, since the size of a CNP and the nature of its surround can be highly variable, obtaining such markers can be quite challenging. A better alternative is to identify the trough (lowest point on a curve) and use it to segment a CNP. Hence, we have taken a different approach to the problem and propose a technique that detects trough points and collates them across scales. We next present the details of our proposed algorithm for CNP segmentation comprising several steps.

2.2 CNP Detection Algorithm

The proposed CNP detection algorithm consists of these stages: Firstly, illumination correction (IC) is done to minimise the background intensity variation followed by denoising to eliminate noise that is frequently found in FFAs. Next, valley detection is performed to locate the seed points in the CNP regions which are used to extract the candidate CNP regions using a region growing algorithm. Finally, thresholding is done to reject false positives among the detected candidates. The processing in each of these stages are described next.

2.2.1 Illumination Correction:

Nonuniform illumination is a problem in retinal colour images as well as angiograms. A camera-model based solution for illumination correction in angiograms, obtained with non-confocal imaging, is given in (Cree et al., 1999) which assumes a macula-centric view of the retina. Our images are not necessarily macula-centric and are obtained from a laser-based confocal imaging system. We modified a quotient based approach proposed for face images (Wang et al., 2004) and model the non-uniform illumination

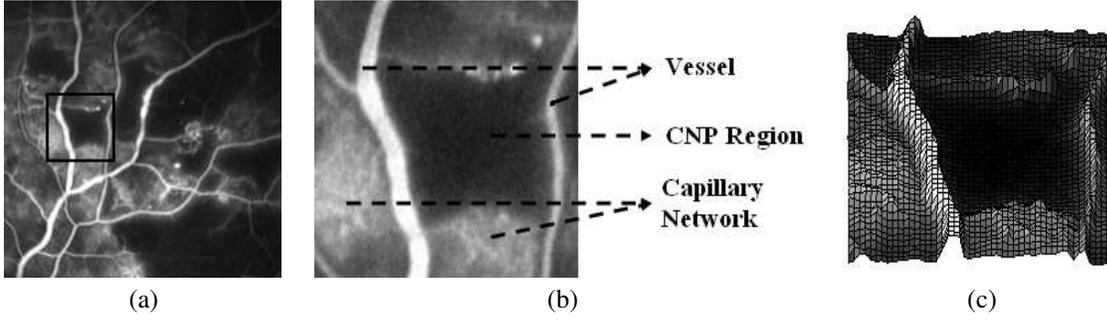


Figure 1: (a) A sample FFA image with CNP. (b) Enlarged view and (c) surface plot of the CNP region in (a).

as a multiplicative degradation function which is estimated by blurring the corrupted image. Let $I(x, y)$, $I_s(x, y)$ and $I_0(x, y)$ denote the given, smoothed and corrected images, respectively and l_0 be the desired level of illumination. The corrected intensity value at location (x, y) is found as

$$I_0(x, y) = \begin{cases} I(x, y) \times \frac{l_0}{I_s(x, y)} & \text{if } I_s(x, y) < l_0 \\ I(x, y) & \text{if } I_s(x, y) \geq l_0. \end{cases} \quad (1)$$

As can be observed from Eq. 1, a pixel where the estimated illumination is greater than the ideal illumination value is not corrected. This is to ensure that the regions which are inherently bright, like the optic-disk, haemorrhages, etc., are not wrongly classified as regions of excessive illumination and corrected accordingly. When the estimated illumination value is less than the ideal illumination value, multiplication by the fraction $\frac{l_0}{I_s(x, y)}$ ensures that regions with illumination less than the l_0 are elevated to the ideal illumination value. Moreover, contrast at such a pixel is improved by a factor of $\frac{l_0}{I_s(x, y)}$ thereby removing the need for subsequent brightness and contrast operations, as required in the case of quotient-image based technique. A sample FFA image and corresponding illumination corrected image is shown in Fig. 2.

2.2.2 Noise Removal:

The laser-based imaging produces fine-grain speckle type of noise in the angiograms as can be seen in Fig. 2. A bilateral filter-based approach proposed for color and gray scale images in (Tomasi and Manduchi, 1998) has been successfully applied to denoise images in an angiogram sequence (Guo et al., 2005a). The strength of bilateral filter based denoising is its ability to denoise without compromising edge quality. This is due to the filter's nonlinear characteristic which permits one to take into account the spatial distance as well the photometric similarity of a pixel to

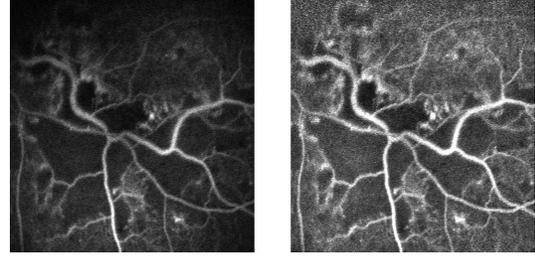


Figure 2: A sample FFA image and Illumination corrected image

its neighbors. The spatial context is provided by a domain filter while the photometric similarity is controlled by a range filter. We use a version of the bilateral filter for our noise removal task which is described next. Given an input pixel $I(P)$, the output pixel $I_0(P)$ is found as

$$I_0(P) = \frac{\sum_w I(Q) W_d(P, Q) W_r(P, Q)}{\sum_w W_d(P, Q) W_r(P, Q)}. \quad (2)$$

where P and Q are position vectors, w is the current neighbourhood and W_d , W_r are Gaussian kernels of the domain and range filters respectively. The edge preservation feature of the bilateral filter can be seen in the results of preprocessing (illumination correction + denoising) in Fig.6 (b).

2.2.3 CNP Segmentation:

Now we turn to the main task of detecting and segmenting CNP regions. Since we have modelled CNPs as valleys, a valley detection algorithm is needed to detect seed points in the CNP regions. As the CNPs vary widely in size, the valleys can be extended. Hence, a multiresolution approach is appropriate. The strategy we have adopted is to reduce the valleys to a single trough point via a pyramidal decomposition

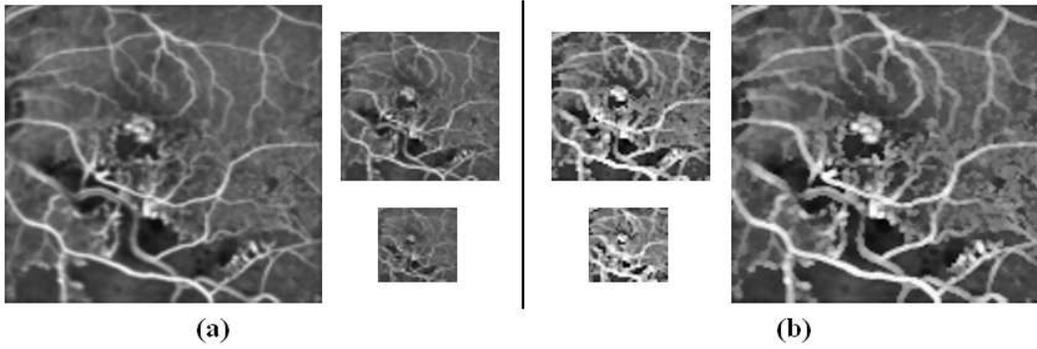


Figure 3: An FFA image at levels 3, 4, 5 in the (a) conventional and (b) proposed pyramidal decomposition.

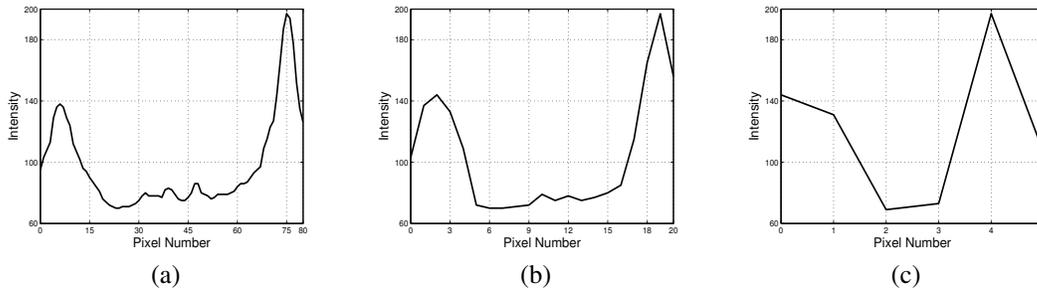


Figure 4: Intensity profile of a CNP and its surround, at levels (a) 1, (b) 3 and (c) 5 of the image pyramid.

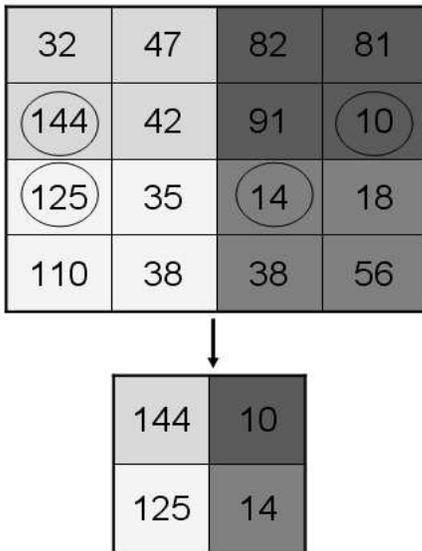


Figure 5: Example demonstrating Equation 3.

and then detect them using a trough detector at each level and collating them. Each of these steps are described next.

Extrema pyramid decomposition - A conventional pyramidal decomposition based on averaging and subsampling is inadequate for the problem at hand. This can be illustrated with an example shown in Fig. 3. It can be seen that the averaging process dulls the entire image and will therefore adversely affect CNP detection based on troughs. Another drawback with the averaging process is the difficulty in localising of the trough points in the full resolution image when performing the upsampling process after trough detection. In the problem at hand, the CNP regions are generally bigger and darker relative to the brighter regions which are thin. Averaging and down-sampling will result in the bright regions to disappear faster than the CNP regions, whereas for locating troughs, it would help to more or less retain the bright regions across several levels while accepting some loss in the CNP area. Hence, to preserve the relation between a CNP and its surround, and maintain the depth of the valley across levels, we need a method for pyramidal decomposition that will minimise the CNP regions at a much faster rate compared to the brighter surrounding regions. This calls for a controlled multi-resolution technique. One option is to generate a pyramid by retaining intensity maxima which will en-

sure the thinner bright regions are largely preserved during downsampling. However, this is detrimental to the relative contrast between a CNP and its surround as it elevates the average intensity of the CNP regions. The end result is a lowering of the depth of the troughs, which is undesirable. A better alternative is to generate the pyramid through an adaptive selection of pixels. The solution we propose is a technique for decomposition which is based on intensity extrema. Specifically, given an image I_l of size $M \times N$ a L -level decomposition is found as follows:

$$I_l(m,n) = \left\{ \begin{array}{ll} \min\{g_{i,j}(m,n)\} & \text{if } g_{i,j}(m,n) \leq t \\ \max\{g_{i,j}(m,n)\} & \text{otherwise} \end{array} \right\}. \quad (3)$$

$$\forall i, j = 0, 1$$

where $g_{i,j}(m,n) = I_{l-1}(2m+i, 2n+j)$, with $l = 2, \dots, L$ and t is a suitable threshold, taken to be the global mean in our experiments. An illustration of the equation is given in Fig. 5 for $t = 100$.

In an extrema pyramidal decomposition of an angiogram, the CNP regions diminish in size at a much faster rate than non-CNP regions across the levels. This is illustrated in Fig. 3 (b) where the thickness of vessels are more or less preserved but the CNPs are reduced to near-dots in the lowest level image. This effect is also seen from the intensity profiles shown in Fig. 4, along a horizontal line passing through a CNP region at different levels. The width of the valley reduces from 50 pixels at the first level, to about 1 pixel in the fifth level whereas the image has been downsampled by 16 between these levels. The relative brightness value (130) of the ridge and valley regions is preserved as a result of not performing a smoothing operation.

Trough detection - A trough is defined as the lowest point on a curve. Alternatively, the brightness at a trough is a local minimum. Since the context in which CNPs, and hence troughs, occur is variable in an angiogram, two parameters can be used to characterise a trough: μ , the mean brightness of the surround and P , the peak factor which represents the depth of the trough. These two parameters are used to develop the following trough detection algorithm in which the image is denoted by $I(x,y)$.

For every pixel (x,y) do the following:

1. Initialize a Boolean variable $isTrough = False$.
2. Check if $I(x,y)$ is a local minimum in a $M \times M$ neighborhood.
3. If yes, then calculate the mean (μ) of a $N \times N$ neighborhood, with $N > M$. Else, do nothing.
4. Let $T = \mu * P$ and check if $I(x,y) < T$.

5. If yes, then $isTrough = True$.

6. If $isTrough = True$, then mark $I(x,y)$ as a trough pixel. Else, do nothing.

The threshold T represents the depth of the valley from the mean μ . Since the image pyramid retains extrema, this threshold value has to be carefully chosen to ensure that enough seed pixels are captured in a valley while minimising the possibility of false alarms. A region with low μ is likely to be a CNP region and hence the required depth for that region is less whereas the same may not be true if μ is high and hence, a stricter condition is required in the latter case. Thus, choosing T proportional to μ is appropriate. Furthermore, since trough detection is carried out at multiple levels a peak factor has to be chosen for each level. A guiding factor in this choice is that due to retention of extremas, the likelihood of the local minima being a CNP region will be higher at upper levels. Hence, the peak factor should be progressively increased with the levels in the pyramid.

After performing trough detection at all levels, the results are combined with a simple logical *OR* operation. For locating the seed pixels in the original image, the fact that the extrema of four pixels is selected at every level is used iteratively.

CNP region extraction - The detected trough points can serve as seed points for region based approach to segmenting the CNP regions. Although geometric methods can potentially yield better results, as an initial experiment we chose to use a simple region growing technique for extracting the CNP regions as it was computationally simpler. Given the variability of the appearance of the CNP regions within and across images, the traditionally used intensity-based homogeneity criterion for region growing is not suitable. Instead, by noting that CNP regions are smooth, the better alternative is to perform the pixel aggregation in the variance space.

In our experiments, the range for the variance was taken to be ± 4 . In order to reject false candidates, a final thresholding operation was performed. A threshold based on the global mean intensity was applied since the global mean is always lowered with the presence of CNPs.

3 IMPLEMENTATION AND RESULTS

The proposed algorithm was implemented as follows. In the illumination correction stage, the ideal illumination I_0 in (1) was set to be roughly half the maximum grey value in the image or 120. The smoothed

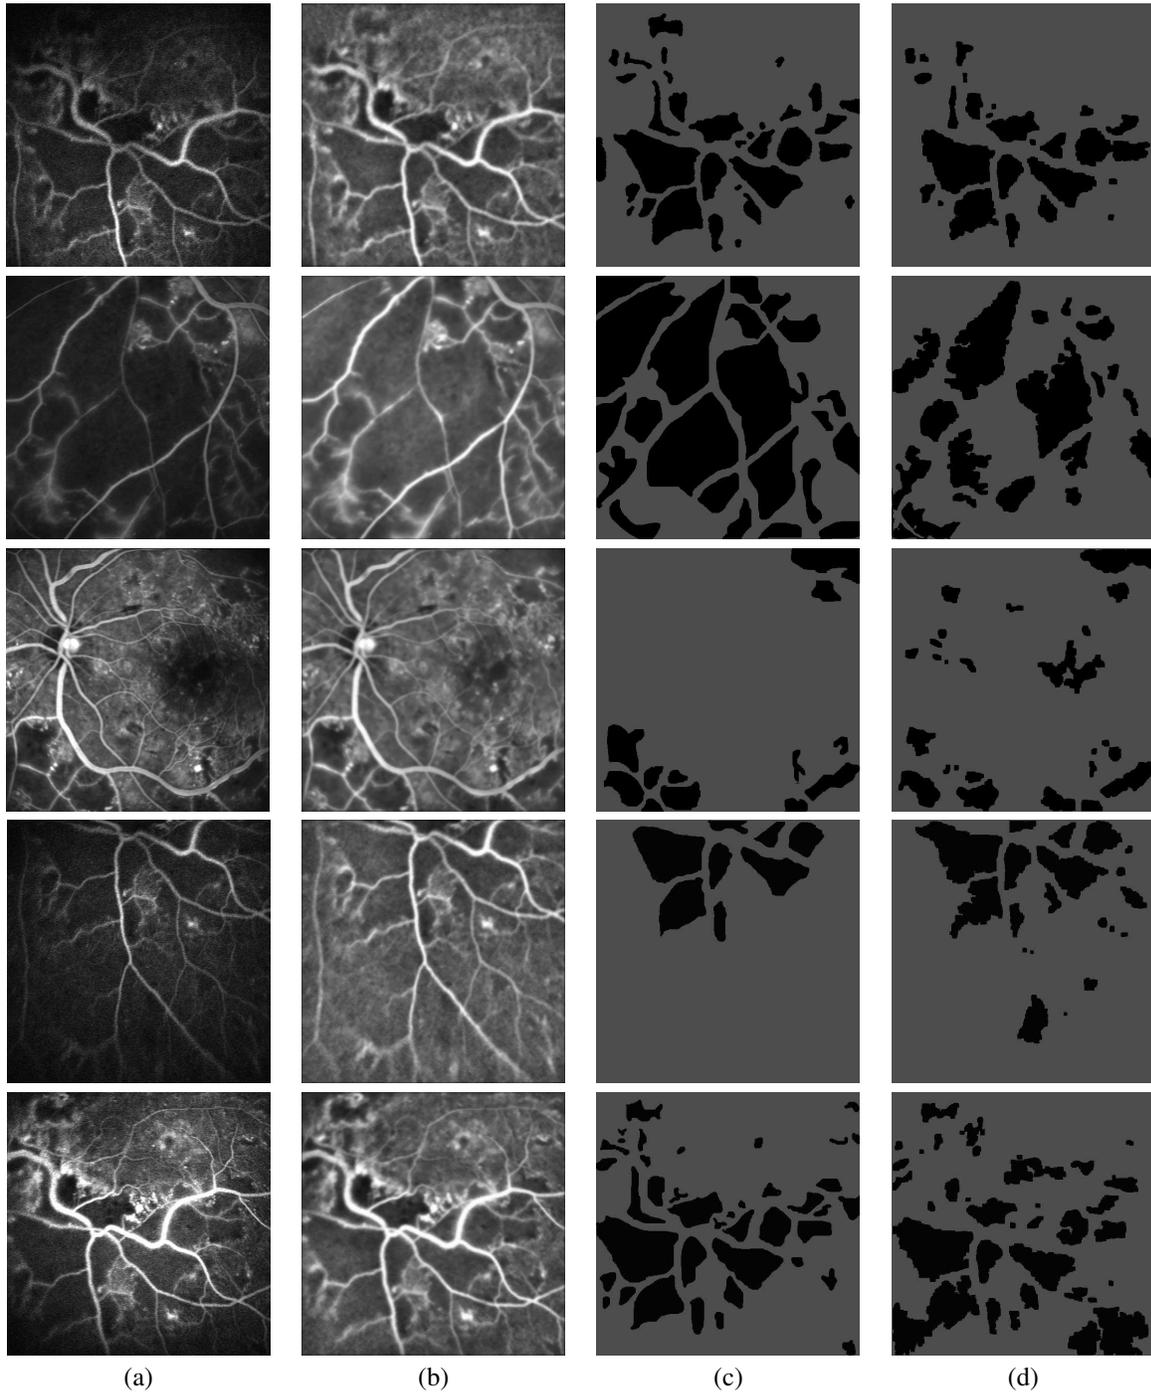


Figure 6: Results: (a) Sample images, (b) preprocessed image, (c) corresponding ground truths and (c) segmented results with CNP regions shown in black.

image was obtained with a 30×30 Gaussian mask on a subsampled (by 4) version of the original image for faster processing. The resulting image was upsampled, one level at a time, while blurring the image

using a 5×5 Gaussian mask at each upsampled level. For denoising, a filter kernel size of 9×9 was used and σ for the domain and range filters were fixed at 3 and 10 respectively. For valley detection, a 5-level

pyramid was generated; M , N were fixed at 5 and 7 respectively and the peak factor was incremented by 0.02 at each level in the pyramid. In region growing, the variance was calculated over a 5×5 neighbourhood.

The proposed CNP segmentation algorithm was tested on 40 images which contained many CNPs. These were acquired from the digital confocal scanning laser ophthalmoscope of Heidelberg Retina Angiograph. The images were of retinal segments for which the ground truth, in the form of boundaries of CNPs, were prepared manually by a retina expert (a co-author). Some sample test images along with corresponding ground truth and results of our CNP segmentation algorithm, with a peak factor of 0.41, are shown in Fig. 6 (a), (b) and (c) respectively. CNP regions are shown in black in both ground truth and segmented results. The five sample test images indicate the variability in images in terms of quality, size of CNPs and presence of other structures such as optic disk, macula and microaneurysms. A quantitative assessment of the algorithm was done using a ROC curve and not a FROC curve since the area of CNP is of clinical interest. A comparison between computed and marked CNP segments was done on a pixel by pixel basis. By using the peak factor as a control parameter, the obtained ROC curve, shown in Fig. 7, was found to have an area under the curve (AUC) of 0.842 and a distance (D_i) to the ideal point (1,0) of 0.35. The ideal values for AUC and D_i are 1 and 0 respectively.

4 CONCLUDING REMARKS

An unsupervised algorithm for automatically segmenting CNPs from FFA images has been presented. Its overall performance is quite good as indicated by the ROC curve and the AUC, D_i metrics. Since there is no reported work on this problem it is not possible to do any benchmarking. A visual inspection of segmented results indicates that the algorithm successfully detects CNPs of all sizes, however, it tends to undersegment large CNPs because the IC stage intensifies the variability within CNPs. A failure analysis indicates that the macula region gets mislabeled as a CNP (as seen in the bottom row of Fig. 6) since the two have similar characteristics, and CNPs in the image periphery tend to be missed since the valley model is weak in this region.

The main focus of the presented work was on detection of CNPs. Hence, while the performance our current implementation is quite good, there is scope for improvement of the algorithm's performance: ge-

ometric techniques such as fast marching method in (Malladi and Sethian, 2006) can be used to more accurately extract the CNP region boundaries while a pixel-based classifier will help improve the rejection of the false alarms. Likewise, incorporation of a macula detection stage will help the algorithm from falsely classifying macula as a CNP region.

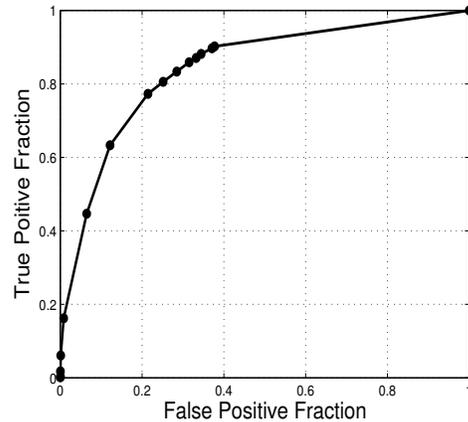


Figure 7: ROC plot.

Finally, it should be noted that the ground truth generation process for CNPs is a laborious one. Retina experts found it challenging to draw precise boundaries because they often appear to be intricate or ill-defined. This points to the need for the use of multiple expert-markings for a fuller evaluation of the algorithm. Such an exercise might also shed light on the degree of observer bias in CNP segmentation.

REFERENCES

- A. M. Mendonça, A. J. Campilho, J. M. N. (1999). Automatic segmentation of microaneurysms in retinal angiograms of diabetic patients. In *Proc. International Conf. on Image Analysis and Processing*, pages 728 – 733.
- Cree, M. J., Olson, J. A., C.McHardy, K., Sharp, P. F., and Forrester, J. V. (1999). The preprocessing of retinal images for the detection of fluorescein leakage. *Phys. Med. Biol.*, 44:293 – 308.
- Fleming, A. D., Philip, S., Goatman, K. A., Olson, J. A., and Sharp, P. F. (2006). Automated microaneurysm detection using local contrast normalization and local vessel detection. In *IEEE Trans. on Medical Imaging*, volume 25, pages 1223 – 1232.
- Gauch, J. M. (1999). Image segmentation and analysis via multiscale gradient watershed hierarchies. In *IEEE Trans. on Image Processing*, volume 8, pages 69 – 79.

- Guo, X.-X., Lu, Y.-N., Xu, Z.-W., Liu, Z.-H., Wang, Y.-X., and Pang, Y.-J. (2005a). Noise suppression of fluorescein angiogram sequences using bilateral filter. In *Proc. International Conf. on Machine Learning and Cybernetics*, volume 9, pages 5366 – 5371.
- Guo, X.-X., Lu, Y.-N., Xu, Z.-W., and Pang, Y.-J. (2005b). An eawa filter for denoising of filtering of fluorescein angiogram sequences. In *Proc. International Conf. on Computer and Information Technology*, pages 614 – 618.
- Hafez, M. (2002). Using adaptive edge technique for detecting microaneurysms in fluorescein angiograms of the ocular fundus. *Proc. Mediterranean Electrotechnical Conf.*, pages 479 – 483.
- Jago, R., Arnold, J., Blauth, C., Smith, P., K.M.Taylor, and Wootton, R. (1992). Measurement of capillary dropout in retinal angiograms by computerised image analysis. *Pattern Recognition Letters*, 13:143 – 151.
- Kohner, E. M. (1993). Diabetic retinopathy. *BMJ*, 307(6913):1195 – 1199.
- Malladi, R. and Sethian, J. A. (2006). Fast methods for shape extraction in medical and biomedical imaging. In *R Malladi (ed) Geometric Methods in Bio-medical Image Processing*. Berlin, Springer, pages 49–61.
- Simó, A. and de Ves, E. (2001). Segmentation of macular fluorescein angiographies. a statistical approach. *Pattern Recognition*, 34:795 – 809.
- Tomasi, C. and Manduchi, R. (1998). Bilateral filtering for gray and color images. In *Proc. International Conf. on Computer Vision*, pages 839–846.
- Wang, H., Li, S. Z., and Wang, Y. (2004). Generalized quotient image. In *Proc. Conf. on Computer Vision and Pattern Recognition*, volume 2, pages 498 – 505.