Constrained-MSER Detection of Retinal Pathology

Gilbert Lim Yong San, Mong Li Lee, Wynne Hsu
National University of Singapore
{*limyongs, leeml, whsu*}*@comp.nus.edu.sg*

Abstract

With the increase in age and diabetes-related eye diseases, there is a rising demand for systems which can efficiently screen and locate abnormalities in retinal images. In this paper, we propose a framework that utilizes a variant of the Maximally Stable Extremal Region method, termed C-MSER, to systematically detect various retinopathy pathologies such as microaneurysms, haemorrhages, hard exudates and soft exudates. Experiments on three real-world datasets show that C-MSER is effective for online screening of diabetic retinopathy.

1. Introduction

Automated screening for diabetic retinopathy is a natural development given the worldwide prevalence of the elderly and diabetic demographics. Diabetic retinopathy can be manifested in many forms, and common pathologies include microaneurysms (MAs), haemorrhages (HAEs), hard exudates (HEs) and soft exudates (SEs), as illustrated in Figure 1.

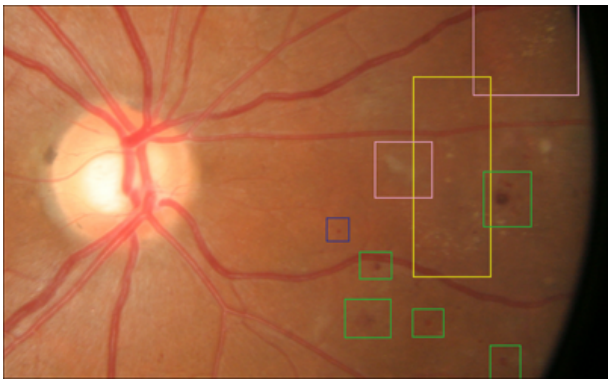

Figure 1: Examples of various types of pathologies. MA (blue), HAE (green), HE (yellow), SE (pink)

Attempts at automated screening of diabetic retinopathy date back to the 1980s, with the first acceptable systems appearing in the late 1990s, most notably a system for detecting microaneurysms in fluorescein angiographic images by the Aberdeen group [12]. Most recently, systems employing template matching using wavelets [10], Gaussian kernels [13] and double-ring filters [5] have exhibited good performance on MAs, but it remains unproven whether such techniques can be effectively generalized to other pathologies that do not conform to the relatively uniform appearance of MAs. Methods relying on machine learning and classification of sub-images of fixed size have also been attempted [2], but are generally unable to delineate the precise extent of the pathologies.

Clearly, an ideal fully automated screening system should be to be able to take a good-quality color fundus photograph of a retina and indicate whether pathologies exist in the image, and if so, where; in other words, the system must be able to handle all types of pathologies. In this paper, we propose the concept of a Constrained-Maximally Stable Extremal Region (C-MSER). The key idea is that visually-distinct regions conforming to pathologies may be *nested*, and it is useful to determine which of the nested regions corresponds to the most plausible description of an individual pathology. The C-MSER concept segments appropriate pathology candidates, and has produced classification results comparable to the state-of-the-art, where comparisons are available, for multiple types of pathologies, at both the lesion and image levels.

2. Background and Related Work

Much research in retinopathy detection has focused on MAs, which are tiny saccular bulges in the walls of capillary vessels appearing as distinct small, round and reddish objects. Their relatively uniform appearance and high correlation with retinopathy diagnosis has made them an ideal target for automated systems.

Work in this respect has most recently cumulated in the Retinopathy Online Challenge (ROCh) [6], which evaluates participants on a common test dataset. The best results obtained boasted successful detection of approximately 40% of all MAs at a rate of one false positive per image.

In any case, exceptional detection performance at the level of individual pathologies is not strictly necessary for screening purposes, with 100% sensitivity and 87% specificity reported at the image level based on a per-MA sensitivity of some 30% [7]. As such, deployable practical screening at the image level has been claimed to be feasible by several authors [9] [1]. However, we feel that there remains room for investigation, both in the identification of MA candidates and also in the identification of other types of pathologies which may not be as regular in terms of shape and other characteristics, and have therefore been paid relatively little attention.

3. Constrained-Maximally Stable Extremal Regions (C-MSER)

We observe that many pathologies appear to the human eye as a reasonably well-defined color blob that is visually separate from the background. In other words, the blob is relatively *stable* at some intensity, since its size does not change much when the image is binarized at successive thresholds in its green channel. In this section, we adapt the concepts expressed in [4] to define Constrained-Maximally Stable Extremal Regions (C-MSER) for the detection of various pathologies in retinal images.

Given a greyscale (green channel) image D, we define a region Q as a *contiguous subset* of D, i.e. for all pixels $p, q \in Q$ there is a sequence $p, a_1, \ldots, a_n, q$ such that p is 4-adjacent to a_1, a_i is 4-adjacent to a_{i+1}, and a_n is 4-adjacent to q. A *region boundary* is then defined as $\delta Q = \{q \in D \setminus Q : \exists p \in Q : q \text{ is 4-adjacent to } p\}$, i.e. the region boundary δQ of Q is the set of pixels adjacent to at least one pixel in Q, but is not itself in Q. A region Q is *nested* within another region Q', denoted by $Q \subset Q'$, if all the pixels in Q also belong to Q'.

An *extremal region* $Q_t \subset D$ at some threshold t is then a region such that for all $p \in Q_t, q \in \delta Q_t$, we have $I(p) \leq t < I(q)$ where $I(x)$ is the intensity of a pixel $x \in D$. Given a parameter Δ, the *stability* of a region Q_t is defined as $s(t) = |Q_{t+\Delta} \setminus \forall Q_{t-\Delta}|/|Q_t|$. Note that this definition may differ from the original definition for regions with multiple child regions, and is justified by the argument that visual stability is better expressed by the relative change of total region size at different thresholds, regardless of whether these regions happen to be separated, at lower thresholds.

Figure 2 shows an example of an extremal region, delineated by the red region Q_{109}. This region grows to 19 pixels at a threshold of 111 (Q_{111}), and shrinks to 15 pixels at a threshold of 107 (Q_{107}). Therefore, for Δ = 2, its stability is $(19 - 15)/16 = 0.25$.

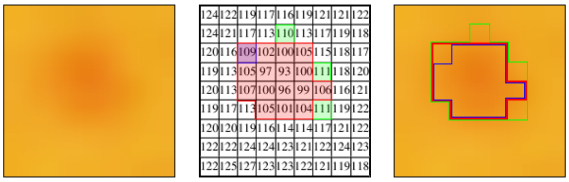

Figure 2: Example extremal regions:
Q_{109} (bounded in red), Q_{107} (blue), Q_{111} (green)

We use a max-tree to keep track of all the extremal regions for various thresholds. Each node in the max-tree corresponds to an extremal region. A node n_1 is a *child* of another node n_2 if and only if the region Q_{t1} corresponding to n_1 is nested within the region Q_{t2} corresponding to n_2, and there exists no other extremal region Q such that $Q_{t1} \subset Q \subset Q_{t2}$. This max-tree representation can be created in linear time [11] [8]. Based on the max-tree representation, an extremal region Q_t is a *maximally stable* extremal region (MSER) if and only if its stability is at a local minimum.

This definition is however not completely ideal for our purposes, since it is possible for a single pathology to correspond to multiple MSER. We therefore propose retaining only the most appropriate MSER from each such nested set, by introducing the idea of Constrained-Maximally Stable Extremal Regions (C-MSER). Figure 3 shows an example where MSER regions A and B correspond to a single MA.

We first use the max-tree to determine all the MSER. Then, an MSER n is a C-MSER if there is no other more-stable MSER n' along the same path as n such that $\frac{1}{r} \leq \frac{size(n)}{size(n')} \leq r$, where r is some user-defined threshold and $size(n)$ denotes the area of the region corresponding to node n. Suppose $r = 4$. It follows that region B is an MSER, but *not* a C-MSER as its size ratio is within the threshold r of region A (shaded in green in Figure 3), which has a better stability.

This proposed C-MSER approach has the ability to detect relevant distinct regions at multiple scales in a single pass, regardless of shape profile, and select an appropriate extent of the pathologies. Furthermore, C-MSER segmentation is linear in terms of time complexity, making it suitable for real-time screening.

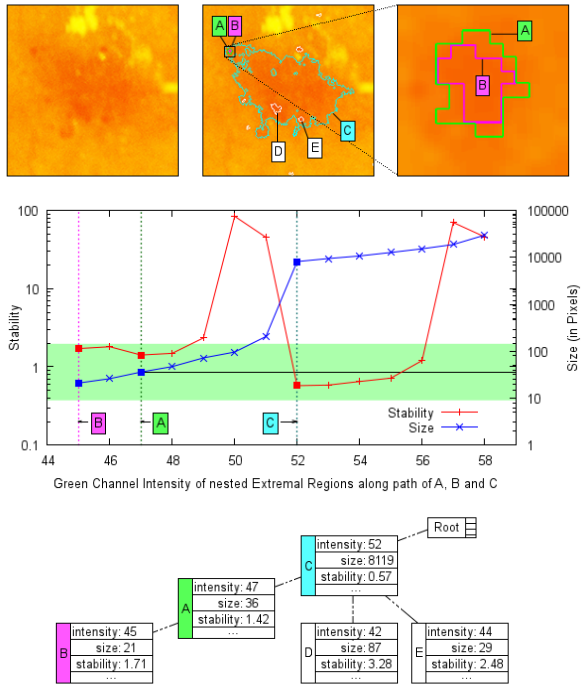

Figure 3: An example of nested features, with associated stability and size curves, and partial max-tree representation. Note that stability may not correspond to Δsize *along a single path*.

4. Experimental Results

In this section, we describe the results of experiments to evaluate the performance of C-MSER on various retinal image datasets. We set $\Delta = 2$ and $r = 4$ for all experiments. Bright pathologies were detected with the same procedure on the inverted-intensity images.

4.1. DIARETDB1 Dataset

We first evaluate the C-MSER method on the DIARETDB1[1] dataset [3], where the ground truths for four major classes of pathologies (MAs, HAEs, HEs and SEs) are supplied. For comparison, besides the colour locus-based results presented in [3], we also implemented a template matching method (TM) by using a filterbank of filters at various sizes and orientations that are suited to the different classes of pathologies.

We follow the methodology in [3] in training each method on the same 21 images, before testing on the remaining 68 images. An image is considered a true *positive* for some class of pathology if its ground truth for that class contains at least one pixel with a confidence level of 75% or above, and *negative* otherwise. For both C-MSER and template matching, we use a learnt discriminant table to identify individual candidates in the test images, then classify at the image level based on findings at the lesion level. The image level ROC curves obtained are shown in Figure 4. It can be seen that the C-MSER method outperforms the other two methods, especially in the case of HAEs and SEs, which often vary more strongly in their characteristics.

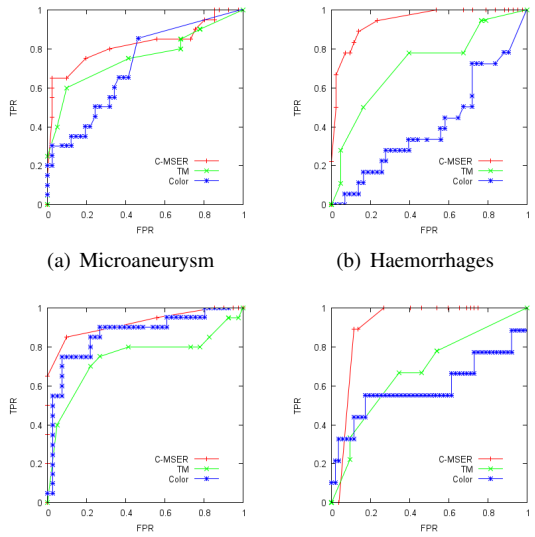

(a) Microaneurysm (b) Haemorrhages

(c) Hard Exudates (d) Soft Exudates

Figure 4: ROC plots of image-level classification for various types of pathologies on the DIARETDB1 dataset.

4.2. Retinopathy Online Challenge Dataset

Next, we employ the Retinopathy Online Challenge (ROCh)[1] training dataset, which consists of 50 images with associated MA ground truth data, to evaluate the performance of C-MSER for lesion-level MA detection. We implemented the existing state-of-the-art Hierarchical Gaussian [13] and Double Ring [5] methods, based on their published descriptions, for comparison. The same vessel masks and classification procedure (by constructing a discriminant table with 31 features as described in [13]) are used to evaluate all three methods.

We vary the stability threshold for C-MSER, and the correlation coefficient/response threshold for the other two methods, to plot the free-response ROC (FROC) graph for MA detection. Figure 5 shows the results. We observe that the C-MSER method outperforms the Hierarchical Gaussian and Double Ring methods when the average false positives per image is less than 8, indicating that C-MSER is suitable for screening purposes.

[1]We thank the organizers of the Retinopathy Online Challenge as well as the Machine Vision and Pattern Recognition Research Group at the Lappeenranta University of Technology for the provision of their retinal colour fundus image databases.

The seemingly low sensitivities for all three methods may be attributed to the large proportion of barely-visible MAs in the Retinopathy Online Challenge training dataset.

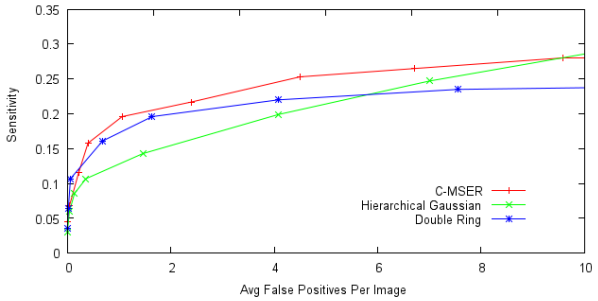

Figure 5: FROC plots of various MA detection methods on the Retinopathy Online Challenge dataset

4.3. DR Evaluation Dataset

Finally, we applied our proposed approach to the DR Evaluation (DRE) dataset, which consists of 417 images taken from a population-based Singapore Malay Eye Studies (SiMES) database that has been graded manually by experts for multiple classes of pathologies. Figure 6 shows the results. We observe that a majority of the pathologies can be detected at a relatively low false positive rate at the lesion level, with the relatively poor performance on soft exudates attributable to confusion with the far more numerous hard exudates.

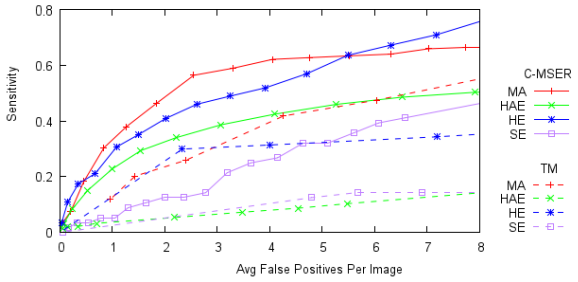

Figure 6: FROC plots of detection results on the DRE dataset

5. Conclusion

In this paper, we have described how we adapted the concept of MSER to identify multiple types of pathologies within a retinal image simultaneously, and evaluated its effectiveness on three retinal image datasets. The experimental results demonstrate that the proposed approach is able to achieve competitive results on the detection of MA, HAE, HE and SE in all datasets.

References

[1] M. Abràmoff, J. Reinhardt, S. Russell, et al. Automated early detection of diabetic retinopathy. *Ophthalmology*, 117(6):1147–1154, 2010.

[2] C. Agurto, V. Murray, E. Barriga, et al. Multiscale am-fm methods for diabetic retinopathy lesion detection. *Medical Imaging, IEEE Transactions on*, 29(2):502–512, 2010.

[3] T. Kauppi, V. Kalesnykiene, J. Kamarainen, et al. Diaretdb1 diabetic retinopathy database and evaluation protocol. *Proc. Medical Image Understanding and Analysis (MIUA)*, pages 61–65, 2007.

[4] J. Matas, O. Chum, M. Urban, et al. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767, 2004.

[5] A. Mizutani, C. Muramatsu, Y. Hatanaka, et al. Automated microaneurysm detection method based on double ring filter in retinal fundus images. *SPIE Medical Imaging 2009: Computer-Aided Diagnosis*, 7260(1):72601N, 2009.

[6] M. Niemeijer, B. Van Ginneken, M. Cree, et al. Retinopathy online challenge: Automatic detection of microaneurysms in digital color fundus photographs. *Medical Imaging, IEEE Transactions on*, 29(1):185–195, 2010.

[7] M. Niemeijer, B. van Ginneken, J. Staal, et al. Automatic detection of red lesions in digital color fundus photographs. *Medical Imaging, IEEE Transactions on*, 24(5):584–592, 2005.

[8] D. Nistér and H. Stewénius. Linear time maximally stable extremal regions. *Computer Vision–ECCV 2008*, pages 183–196, 2008.

[9] S. Philip, A. Fleming, K. Goatman, et al. The efficacy of automated disease/no disease grading for diabetic retinopathy in a systematic screening programme. *British Journal of Ophthalmology*, 91(11):1512–1517, 2007.

[10] G. Quellec, M. Lamard, P. M. Josselin, et al. Detection of lesions in retina photographs based on the wavelet transform. In *Engineering in Medicine and Biology Society, 2006. EMBS'06. 28th Annual International Conference of the IEEE*, pages 2618–2621. IEEE, 2008.

[11] O. A. Salembier, P. and L. Garrido. Anti-extensive connected operators for image and sequence processing. *IEEE Transactions on Image Processing*, 7(4):555–570, 1998.

[12] T. Spencer, J. A. Olson, K. C. McHardy, et al. An image-processing strategy for the segmentation and quantification of microaneurysms in fluorescein angiograms of the ocular fundus. *Computers and biomedical research*, 29(4):284–302, 1996.

[13] B. Zhang, X. Wu, J. You, et al. Hierarchical detection of red lesions in retinal images by multiscale correlation filtering. *SPIE Medical Imaging 2009: Computer-Aided Diagnosis*, 7260:72601L, 2009.