

LESION BOUNDARY SEGMENTATION USING LEVEL SET METHODS

Elizabeth M. Massey, James A. Lowell

University of Lincoln, Brayford Pool, Lincoln LN6 7TS UK

Foster Findlay Associates Limited, Newcastle Technopole Kings Manor, Newcastle Upon Tyne NE1 6PA UK

bmassey@lincoln.ac.uk, james.lowell@gmail.com

Andrew Hunter, David Steel

University of Lincoln, Brayford Pool, Lincoln LN6 7TS UK

Consultant Ophthalmologist, Sunderland Eye Infirmary, Queen Alexandra Road, Sunderland SR2 9HP UK

ahunter@lincoln.ac.uk, David.Steel@chs.northy.nhs.uk

Keywords: computer vision, retinal lesion segmentation, segmentation, level set methods.

Abstract: This paper addresses the issue of accurate lesion segmentation in retinal imagery, using level set methods and a novel stopping mechanism - an *elementary features* scheme. Specifically, the curve propagation is guided by a gradient map built using a combination of histogram equalization and robust statistics. The stopping mechanism uses elementary features gathered as the curve deforms over time, and then using a *lesionness* measure, defined herein, 'looks back in time' to find the point at which the curve best fits the real object. We implement the level set using a fast upwind scheme and compare the proposed method against five other segmentation algorithms performed on 50 randomly selected images of exudates with a database of clinician marked-up boundaries as ground truth.

1 INTRODUCTION

The diagnosis of diabetic retinopathy is based upon visually recognizing various clinical features. Retinal lesions are among the first visual indicators suggestive of diabetic retinopathy. The threat to visual loss increases with the frequency of retinal lesions combined with their encroachment into the macula (one optic disc diameter around the fovea). To enable early diagnosis, it is therefore necessary to identify both frequency and position of retinal lesions in relation to the fovea. This paper focuses on the segmentation of retinal lesions and presents an application of level set methods and a novel elementary features scheme for ensuring an accurate boundary detection solution. While most applications of level set methods have yielded excellent results, many assume a fairly noise-free surface. We propose to apply level set methods to retinal images, which are noisy and have a slight surface curve especially near the edges. We present a novel stopping mechanism which uses elementary features gathered over time as the curve deforms and then a calculated *lesionness* measure to find the point in time at which the curve best fits the lesion candidate.

This paper is presented as follows: Sections 2 and

3 provide some background information and discuss the current literature, respectively, on region growing schemes as a basis for comparison. Section 4 describes the level set method used followed by a description of the algorithm and the process framework. Section 5 discusses the evaluation results and provides comparison and observations about the proposed method. Section 6 concludes the paper.

2 BACKGROUND

2.1 Exudate Appearance

Retinal exudates are an interesting challenge for segmentation algorithms as they vary in appearance, conforming to one of three structures: dot exudates, fluffy exudates and circumscribed plaques of exudate. Dot exudates consist of round yellow spots lying superficially or deep in the sensory retina (Porta and Bandello, 2002). Fluffy exudates are more pale than dot exudates and tend to lie more superficially in the sensory retina. Plaque exudates vary in size more than the other two groups and represent a more diffuse accumulation of lipoprotein. In addition to their vari-

ous appearances, exudates can be arranged in different configurations. Exudates may surround leaking capillaries and microaneurysms in a circular (circinate) pattern or lie scattered, taking on no particular shape. Exudates are usually reflective and may appear to have a rigid, multifaceted contour, ranging in colour from white to yellow (Chen, 2002).

With varying shapes, sizes, patterns and contrast, exudate segmentation is a demanding problem, complicated by lighting variation over the image, natural pigmentation, the intrinsic colour of the lesion, and decreasing colour saturation at lesion boundaries (Goldbaum et al., 1990). Figure 2 shows examples of the various shapes retinal exudates can exhibit.

3 Previous Work

3.1 Segmentation Algorithms

Several authors have presented algorithms for the segmentation of exudates in fundus images, attaining varied results. Ward *et al.* (Ward et al., 1989) introduced a semi-automated exudate detection and measurement method, in which an operator selected a threshold value to segment exudates from a shade-corrected retinal background.

Sinthanayothin *et al.* (Sinthanayothin et al., 2002) presented a recursive region-growing algorithm applied to a contrast enhanced image. To reduce the effects of uneven illumination over the fundus, images were pre-processed to enhance local contrast. The intensity component of the IHS (Intensity Hue Saturation) model was decoupled from colour and the fundus images converted from RGB (Red Green Blue) and normalised to IHS; see equations 1 to 3.

$$I = \frac{1}{3}(R + G + B) \quad (1)$$

$$S = 1 - \frac{3}{R + G + B}[\min(R, G, B)] \quad (2)$$

$$H = \cos^{-1} \left\{ \frac{\frac{1}{2}[(R - G) + (R - B)]}{[(R - G)^2 + (R - B)(G - B)]^{\frac{1}{2}}} \right\} \quad (3)$$

Local contrast enhancement was calculated using sub-windows of 49×49 centered on each pixel (i, j) of the IHS's intensity component. The contrast enhancement transformation is defined by equation 4 and uses the sigmoidal function in equation 5. The maximum and minimum intensity values of the intensity image are denoted by f_{max} and f_{min} respectively, with the mean and standard deviation of the intensity

within the sub-window W being denoted by $\langle f \rangle_w$ and $\sigma_w(f)$ (equations 6 and 7) respectively.

$$f(i, j) \rightarrow g(i, j) = 225 \frac{[\Psi_w(f) - \Psi_w(f_{min})]}{[\Psi_w(f_{max}) - \Psi_w(f_{min})]} \quad (4)$$

$$\Psi_w(f) = \left[1 + \exp \left(\frac{\langle f \rangle_w - f}{\sigma_w} \right) \right]^{-1} \quad (5)$$

$$\langle f \rangle_{w(i, j)}(f) = \frac{1}{M^2} \sum_{(k, l) \in W(i, j)} f(k, l) \quad (6)$$

$$\sigma_w^2(f) = \frac{1}{M^2} \sum_{(k, l) \in W(i, j)} (f(k, l) - \langle f \rangle_w)^2 \quad (7)$$

By using a sigmoidal function, areas with poor initial contrast (small σ value) were greatly enhanced, leaving areas with good initial contrast (large σ value) largely unaltered. To reduce the noise produced by contrast enhancement, a 2D Gaussian smoothing filter was first applied to the image. Recursive region growing works by dividing the image into regions of similar gray levels. Starting with a pixel at coordinates (x, y) the intensity values of its four neighbouring pixels $(x, y + 1)$, $(x, y - 1)$, $(x - 1, y)$ and $(x + 1, y)$ are compared. If the intensity difference between the current pixel and a neighbouring pixel is less or equal to a threshold (in this case 10 intensity levels) it is added to the current region. A region is defined when no more pixels qualify to join the current region. At this point, Sinthanayothin replaces the original intensity values within the region with the median intensity of the area. To split the image into exudate and non-exudate regions, a threshold is set to median intensity of the image background (based on the region with the most pixels). Regions with a median intensity above and below this threshold were classified as exudates and non-exudates respectively. Sinthanayothin stated that the algorithm would not detect faint exudated regions, nor distinguish between other similar coloured lesions.

Wang *et al.* (Wang et al., 2000) defines a feature space \mathbf{F} to include colour and exposure information and represents the red, green and blue (R,G,B) channels as spherical coordinates described by equations 8, 9 and 10.

$$L = (R^2 + G^2 + B^2)^{\frac{1}{2}} \quad (8)$$

$$\Theta = \text{Arccos} \left(\frac{G}{R} \right) \quad (9)$$

$$\Phi = \text{Arccos} \left(\frac{B}{L} \right) \quad (10)$$

The brightness of the image is denoted by \mathbf{L} and the transformed colour coordinates denoted by φ and Θ . The feature space \mathbf{F} is therefore represented by $\mathbf{F}(f_L, f_\Theta, f_\varphi)$. A training set for each of two groupings is obtained by selecting small sub-windows inside exudate and background regions. The means of each sub-window are calculated and stored as feature centers for the two groups $\mathbf{C}_{lesion}(f_L, f_\Theta, f_\varphi)$ and $\mathbf{C}_{bkgrnd}(f_L, f_\Theta, f_\varphi)$ respectively. For each pixel in the fundus image the illumination and colour information are extracted $\mathbf{X}(X_L, X_\Theta, X_\varphi)$ and the minimum distance discriminant $D_i(\mathbf{X})$ (equation 11) is calculated for each group $D_{lesion}(\mathbf{X})$ and $D_{bkgrnd}(\mathbf{X})$.

$$D_i(\mathbf{X}) = (\mathbf{X} - \mathbf{C}_i)^T (\mathbf{X} - \mathbf{C}_i) \quad (11)$$

Osareh *et al.* (Osareh et al., 2001) introduced a fuzzy C-Means clustering algorithm based on the work of Young Won Lim *et al.* (Lim and Lee, 1990) to segment a colour retinal image into homogenous regions. To compensate for the wide variation of colour in the fundus, the images are converted from RGB to IHS, normalised and finally locally contrast enhanced (described above in Sinthanayothin *et al.*). Fuzzy C-Means clustering allows pixels to be segmented and grouped with varying degrees of membership. Young Won Lim's algorithm consisted of two stages - coarse and fine segmentation. The coarse segmentation divides the image into a number of regions determined by thresholds automatically generated by scale-space filtering Hue Saturation and Intensity (HSI) histograms. The number of significant peaks in the histogram determine the number of regions. Pixels not segmented by the coarse segmentation phase are processed in the fine segmentation phase by fuzzy partitioning. This is done by minimizing the squared error loss function in equation 12

$$L = \sum_{j=1}^m \sum_{i=1}^n [\mu_j(x_i)]^b \|x_i - c_j\|^2 \quad (12)$$

where x is a data vector, ($i = 1, \dots, n$), is the number of pixels, c_j denotes the centre of fuzzy clusters ($j = 1, \dots, m$) where m is the number of clusters determined in the coarse segmentation stage. The fuzzy membership of x_i to cluster j is denoted by $\mu_j(x_i)$, and the overlap between the fuzzy cluster regions is controlled by a weighting exponent b . Osareh et al (Osareh et al., 2001), states that the segmentation by FCM is a conservative process finding all but the faintest (ambiguous) exudate regions. False positive non-exudate segmented regions were also found by the algorithm caused by cluster overlapping, noise, and uneven colour distribution.

Contrast Gradient Region Growing (CG), introduced in (Lowell, 2005), uses a traditional region growing method employing a pixel intensity aggregation scheme for region growth, while using a Gaussian smoothed gradient image to iteratively calculate a gradient contrast between a grown (core) inner boundary and a dilated outer boundary. The algorithm starts with a seed (peak) point, determined by identifying isolated points where the intensity value of the isolated point is quite different from its neighbouring pixels. A small 5×5 sub-window is morphologically run over the fundus image, applying a maximum filter within each sub-window, producing peak points. The core region is then grown by appending (selecting) the brightest neighbouring (boundary) pixels on each iteration. The inclusion of the brightest pixel from boundary into the 'selected' pixels is repeated until a stopping criteria is met. This growing process continues, halting when the grown region loses its *compactness*. If the compactness metric exceeds a threshold of 30, region growing is terminated. The final boundary is then located by using a combination of *diameter* and *contrast* to determine the point of growth at which the object's contrast gradient is most significant.

Since the pioneering work of Osher and Sethian (Osher and Sethian, 1988) Geometric Deformable Models, or Level Sets, have had a significant impact on the imaging community due to their capability to preserve the topological information in an image. The applications of level sets are varied, including object tracing/image sequence analysis, shape from shading, color image segmentation, object shape recovery, 3-D reconstruction and modeling and surface defect evolution simulation (Srolovitz, 2005). For a state-of-the-art review of level sets and their applications see (Suri et al., 2002).

The literature on retinal image object segmentation using level sets focuses mainly on segmenting structures rather than pathologies. Excellent work by Wang *et al.*, (Wang et al., 2004) show the power of evolving a curve to map prominent structures in an image. Deschamps *et al.* (Deschamps et al., 2004) used level sets combined with embedded boundary methods to simulate blood flow and segment major vessels. Lowell, *et al.* (Lowell et al., 2004) used active contours, the fore-runner to level sets, to find the optic nerve head. The work described herein is based on the seminal paper from (Osher and Sethian, 1988) and the numerical implementation takes insights from Sapiro, chap.2 (Sapiro, 2001).

4 Level Set Method

For our work in lesion segmentation, level set methods provide the capability to determine not just the coarse shape of an object, but are extremely useful to tease out the fine delicate boundary fissures and curves that give a deeper look into the overall shape of a lesion candidate.

4.1 Curve Propagation

Beginning with the definition of level sets from Osher and Sethian (Osher and Sethian, 1988)

$$\phi_t + F |\nabla\phi| = 0, \text{ given } \phi(\bar{x}, t = 0) \quad (13)$$

then,

$$\frac{\partial\phi}{\partial t} = F |\nabla\phi| \quad (14)$$

and

$$\phi_t + F_0 |\nabla\phi| + \vec{U}(x, y, t) \vec{\nabla}\phi = \epsilon K |\nabla\phi| \quad (15)$$

where: ϕ_t is the propagating function at time t ,

$F_0 |\nabla\phi|$ is the motion of the curve in the direction normal to the front,

$\vec{U}(x, y, t) \vec{\nabla}\phi$ is the term that moves the curve across the surface,

$\epsilon K |\nabla\phi|$ is the speed term dependent upon curvature.

For our purposes, $\vec{U}(x, y, t) \vec{\nabla}\phi$ is the gradient map, described in section 4.3 and $\epsilon K |\nabla\phi|$ is approximated using a central differencing scheme.

4.2 Numerical Implementation

We consider curve movement of the form:

$$\frac{\partial C}{\partial t} = \beta \tilde{N} \quad (16)$$

where $\beta = \beta(k)$, that is, β is a function of the Euclidean curvature. For simplicity we use $\beta(k) = 1 + \epsilon k$ as our velocity function.

Let ϕ_i^n be the value of ϕ at a point (pixel) i at the time n . An algorithm to describe the evolution of the curve over a given time step is

$$\phi_{ij}^{n+1} = \phi_{ij}^n - \Delta t [\max(-\beta_{ij}, 0) \Delta^+ + \min(-\beta_{ij}, 0) \Delta^-] \quad (17)$$

where u_{ij}^n is the 'current' level set zero, Δt is the time step (or scaling factor) and the $[\max \dots \min]$ describes the *normal* component, and where

$$\Delta^+ = [\max(D_x^-, 0)^2 + \min(D_x^+, 0)^2 + \max(D_y^-, 0)^2 + \min(D_y^+, 0)^2]^{1/2} \quad (18)$$

$$\Delta^- = [\max(D_x^+, 0)^2 + \min(D_x^-, 0)^2 + \max(D_y^+, 0)^2 + \min(D_y^-, 0)^2]^{1/2} \quad (19)$$

and $D_x^-, D_x^+, D_y^-, D_y^+$ are the forward and backward difference approximations in the x and the y direction, respectively.

4.3 Gradient Map

The boundary of a lesion can be characterised by the point of strongest intensity contrast between itself and the background retina. By determining the gradient of image I_{orig} , this maximum rate of change can be exploited. Equation 17 propagates the curve ϕ over the surface u . Optimally, what we want is to propagate to an object edge and then stop when the curve has correctly formed to the (correct) perimeter pixels. To do this we must provide an edge stopping function. Since the retinal images are inherently noisy, and the edge pixels of retinal lesions can look very much like background pixels, we want a mechanism that smooths out the noise but preserves the edges. Isotropic filters (such as Gaussians) smooth the image, but also lose important detail. Anisotropic filters address the issue of edge preservation. Foundational work in anisotropic diffusion by Perona and Malik (Perona and Malik, 1990) gives the following classical description:

$$\frac{\partial I(x, y, t)}{\partial t} = \text{div}(g(\|\nabla I\|, \sigma) \nabla I), \quad (20)$$

where $\|\nabla I\|$ is the gradient magnitude, and $g(\|\nabla I\|)$ is an edge-stopping function and σ is a scale parameter. The g function is chosen to satisfy $g(x, \sigma) \rightarrow 0$ when $x \rightarrow \infty$, so that diffusion is 'stopped' across the edges. See also (Black and Sapiro, 1999) as an excellent source for a description of anisotropic diffusion.

Perona and Malik suggested the following edge-preserving g function (Perona and Malik, 1990)

$$g(x) = \frac{2x}{2 + \frac{x^2}{\sigma^2}}. \quad (21)$$

The function $g(x)$ acts as a 'weighted' function in that, small gradient values x will receive high weight and high gradient values will have low influence on the diffusion solution. In other words, areas of high gradient will be 'smoothed' less, thus preserving edges. While this function can be implemented in an iterative fashion, and $x = (x, y)$, a pixel location is smoothed over time, we applied the function statically to create our gradient map

$$g_I(x, y) = \frac{2 * (I_n)}{(2 - (I_n)^2)} \quad (22)$$

where: I_n is a histogram equalized, normalised grayscale (green channel) image $I(x, y)$ and $\sigma = 1$.

The anisotropic nature of equation 21 increases the contrast of the high amplitude pixels thus allowing the curve a strong edge to detect during propagation.

4.4 Stopping Criteria

Once the gradient map is generated from the original (gray-scale) image the curve propagates for a given number of iterations. Finding the ‘best’ stopping point for the curve is relative to the object boundary. In cases such as figure 1 the boundary is not well defined, even with a properly contrasted gradient map, and especially in the case of bright lesions, the ‘boundary’ can be much the same color as the background. It is for these reasons that we need to use a mechanism that is robust to conditions of noise and illumination variance.




Figure 1: Curve Fitting: a) Gradient Map b) Match Results

A traditional use of level sets is to track a curve to an object’s boundary. In our case, it is more interesting to ‘peek ahead’ by allowing the curve to move past the optimal boundary and then ‘look back’ and measure how well-formed the accumulated region is as a lesion. We define the term *lesionness* as a combination of compactness ($c = p^2/a$), where $p = \text{perimeter}$ and $a = \text{area}$ (Gonzalez and Woods, 2001) and perimeter size constancy *shp* and use it as our ‘stopping’ mechanism. These measurements and others are explained in section 4.5.2. A significant problem of curve deformation algorithms is the formulation of a suitable stopping criteria. Optimally, the curve will stop its movement once the point of highest gradient (boundary) is reached. If the intensity difference between a lesion and background is marginal, using gradient thresholds may not be enough to halt the advancement of the curve front. As stated previously, when segmenting lesions several challenges exist and specifically, in terms of finding a suitable stopping criteria. First, lesions vary in size, so defining termination rules based on area would likely produce unsatisfactory results. Second, bright lesions can vary widely in shape; as figure 2 shows a few examples of lesion shapes from the current dataset.




Figure 2: Exudate Shapes: a) Round b) Inverted Heart c) Flame d) Rectangular.

4.5 Process & Algorithm

The elementary features algorithm encapsulated in a four phase framework: 1) Pre-processing 2) Processing and Measurements 3) Initial Value Determination 4) Reassessment.

4.5.1 Pre-processing

The single channel, 59x59 pixel image I_{orig} is used to generate a gradient map as discussed in section 4.3. The initial level set begins as a small circle of $radius = 1$ and propagates outward according to equation 17 and as the curve deforms measurements at each change (in form) are taken.

4.5.2 Processing and Measurements

The starting point of the curve is determined using the simple peak detection algorithm described in Contrast Gradient Region Growing (above). The curve is then allowed to propagate past the optimal point (boundary) of the object. The purpose of this is to avoid the underestimation problem inherent in traditional region growing methods, and take advantage of ‘forward/backward looking’ measures.

We are looking for measurements that can give indicators of how well-formed a region is as a candidate lesion. Thus, elementary features include 1) perimeter *length* of curve (p), 2) the number of iterations the curve held its size (shp), 3) the compactness value (c), 4) the number of iterations the curve held that compactness value (chp), and 5) the maximum gradient contrast. Using morphological operations of dilation, equation 23, and erosion, equation 24, two ‘rings’, an

inner and an outer ring, are generated about the curve. The contrast between these two rings is calculated.

$$\delta_{CE}(C_0) = C_0 \oplus CE \quad (23)$$

$$\varepsilon_{CE}(C_0) = C_0 \ominus CE \quad (24)$$

where: CE is a 3×3 structuring element.

After the curve has moved for a number of iterations (we use 180) it is possible that the curve has evolved past the optimal point describing the object boundary. Because of this possibility, the gathered measurement values are then used to ‘look back in time’ to find the point at which the curve best fit the object boundary.

4.5.3 Initial Value Determination

Two measurements that show curve stabilization (slowing down) are shp and chp . When the curve reaches an ‘edge’ its propagation rate slows down and over a number of counted iterations the perimeter size and compactness values tend to remain the same. We count the number of iterations during these stabilizing points and find that they tend to coincide with the other important measures. Figure 3 shows an example of how the measurements work together to find the optimal fit. In figure 3, the x-axis is *iterations* and the y-axis is *counts*. The two iteration points encircled by ellipses are the *minCommon* (17) and *maxCommon* (37) points of commonality between the curve size constancy shp and compactness constancy chp measures. If the curve holds its size and compactness for more than two iterations we say the curve is beginning to stabilize, or slow down, and this is an indication of an edge. We want to know which of the other measures fall within these common stabilizing points and use those relevant measures to make an initial estimation (IV) of the best fit. In figure 3 denoted by the square is the max gradient contrast point $max(gc) = 8$. The circle at iteration 14 shows the point where the curve held its size (perimeter) for the longest $max(shp) = 14$ (i.e. the curve moved very slowly during this time), the triangle at iteration 21 shows the point where the compactness value is at minimum $min(cp) = 21$ (where the shape of the curve was *most* compact). Encircled by the same ellipse as the maximum common point (37) is where the curve held its compactness value the longest $max(chp) = 37$. Notice that most of the significant elements tend to cluster down toward the lower numbers. This visual cue is an indication that the lesion is small. Computationally, we use this ‘clustering’ in the reassessment phase of the process.

To get an initial value estimate we add together the element (measurement) values that fall within the




Figure 3: Measures of Curve Elements

common range and divide by the number of elements in the range

$$IV = \frac{\sum_{e \in [M, N]} x}{x \in [M, N]} \quad (25)$$

where: $e = elements$; $x = number\ of\ elements$; $M = minCommon$; $N = maxCommon$.

4.5.4 Reassessment

If we sort all the elements along with the min and max common points we see that most of the elements tend toward the lower numbers (e.g., < 20). Taking the differences of all sorted elements de , less than the maximum difference, and using that average as a reassessment against the initial value IV we get a second value estimate.

$$SV = \frac{\sum_{de < max(de)} size(de)}{size(de)} \quad (26)$$

In this example, the difference between $IV = 29$ and $SV = 15$ is nearly as great as one of the values, so, we tend toward the smaller number to avoid overestimation. In this case, the difference is 14, so we choose the $min(IV, SV)$ and take that as our final curve. Figure 4 shows the curve plots at the various elemental values.




Figure 4: a) Plots of various elemental points, b) Final curve.

The final curve has a length, given by the Euclidean distance between the (x, y) coordinates of each pixel included on the zero level set curve.

5 Evaluation

A comparison is made between the presented algorithm and five segmentation approaches - fuzzy C-Means clustering, recursive region growing, adaptive recursive region growing, contrast gradient region growing and a colour discriminant function. Table 1 shows the results of our evaluation.

Table 1: Algorithm Performance Metrics.

Model	Sens.	Spec.	Accuracy	Error
ELS	96.94	98.97	98.87	29.35
CG	96.24	98.71	98.59	36.59
AR	91.13	92.53	92.45	196.15
Fuzzy	88.29	94.18	93.89	158.95
RRG	47.72	90.99	88.85	290.1
DC	64.67	75.77	75.21	644.75

Where:

ELS - Elementary Features Scheme;

CG - Contrast Gradient;

AR - Adaptive Recursive;

Fuzzy - Fuzzy C-means;

RRG - Recursive Region Grow;

DC - Colour Discriminant.

All algorithms were implemented and evaluated against a reference standard dataset of 50 randomly selected lesion images. Each image is provided with boundary markups by an expert ophthalmologist using custom designed software. The images are provided by the Sunderland Eye Infirmary with permission to be used in this research.

The benchmark comparison with the aforementioned techniques was achieved by measuring the number of common pixels shared between the reference standard and the algorithm's segmented area. For each reference standard region r , true positive TP , false negative FN , false positive FP and true negative TN statistics were calculated for each segmentation approach. The values in Table 1 were measured using pixel-wise sensitivity, specificity, accuracy and error-rate:

$$\text{Sensitivity} = \frac{TP}{TP + FN} \quad (27)$$

where: TP =pixels matched to reference standard

$$\text{Specificity} = \frac{TN}{TN + FP} \quad (28)$$

where: TN =background pixels in reference standard

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + TN + FN} \quad (29)$$

where: FP =pixels added over the reference standard boundary and FN =pixels missed in reference standard

$$\text{Error} = FN + FP \quad (30)$$

The following observations are made on the performance of the ELS algorithm:

Accuracy As shown in table 1 the ELS method outperforms the CG algorithm in all areas and especially in showing a reduction in error. Experiments show that the CG algorithm tends to underestimate the lesions in general, as denoted by the sensitivity measure. Figure 5, shows how the CG algorithm underestimates the same lesion object that the ELS algorithm overestimates. The CG algorithm tends to underestimate the lesions in general due to the smoothed gradient image used to determine the boundary contrast. As such, low contrast pixels get merged into the retinal pigmentation.




Figure 5: a) CG Enhanced Image b) CG Result c) ELS Gradient Map d) ELS Result.

Robustness The ELS algorithm does not depend on a single criteria, such as *compactness* to find a solution, rather a multiple of measurements are taken as the curve propagates. Since the measurements are not dependent on specific thresholds, the *true* measures of the data can be taken into account during the initial value calculation and reassessment phases.

Geometric The ELS algorithm is not dependent on a single pixel value at a specific point in time, rather the curve moves in relation to curvature and direction of the normal. Thus, global as well as local

information is used during curve propagation. Tracking the zero level set, as we do here, overcomes topological problems (such as discontinuities) that would hamper, even halt, traditional curve propagation algorithms. The current implementation allows the curve to split and merge as necessary during propagation. This capability, inherent in the method, provides for future research into discovering the *distribution* of bright blobs in retinal images. The pattern of distribution or clustering of bright exudates, or other bright objects helps clinicians to diagnose not only the bright spot, but also other anomalies around it.

6 Conclusions

Algorithms for the automated segmentation and classification of candidate lesions have been presented. Although a number of algorithms have been published for lesion segmentation, many are unreliable due to marginal colour and intensity difference between diabetic lesions and the background retina. This limited contrast has an adverse effect on alternate algorithms causing poor lesion boundary estimations.

Experimental comparisons have been conducted on five segmentation approaches - Contrast Gradient, Fuzzy C-Means clustering, recursive region growing, adaptive recursive region growing, and a colour discriminant function. All algorithms were evaluated against a randomly-selected image set with ophthalmic lesion boundary demarcation. The results shown in Section 5 demonstrate the advantage of allowing the curve propagation (region growing) to run past the optimal boundary point, thus providing a 'peek ahead' to adjacent areas. Then using gathered elementary features to 'look back in time' to determine the best fitting curve.

REFERENCES

- Black, M. and Sapiro, G. (1999). Edges as outliers: Anisotropic smoothing using local image statistics. *Scale-Space Theories in Computer Vision*, pages 259–270.
- Chen, H.-C. (2002). *Vascular Complications of Diabetes; current issues in pathogenesis and treatment*, chapter 10, pages 97–108. Blackwell Publishing.
- Deschamps, T., Schwartz, P., Trebotich, D., Colella, P., Saloner, D., and Malladi, R. (2004). Vessel segmentation and blood flow simulation using level-sets and embedded boundary methods. *Computer Assisted Radiology and Surgery. Proceedings of the 18th International Congress and Exhibition*, 1268:75–80.
- Goldbaum, M., Katz, N., Nelson, M., and Haff, L. (1990). The discrimination of similarly colored objects in computer images of the ocular fundus. *Investigative Ophthalmology & Visual Science*, 31:617–623.
- Gonzalez, R. C. and Woods, R. E. (2001). *Digital Image Processing*. Prentice Hall, Upper Saddle River, NJ.
- Lim, Y. W. and Lee, S. U. (1990). On the color image segmentation algorithm based on the thresholding and the fuzzy c-means technique. *Pattern Recognition*, 23:935–952.
- Lowell, J. (2005). *Automated Retinal Analysis*. PhD thesis, University of Durham.
- Lowell, J., Hunter, A., Steel, D., Basu, A., Ryder, R., Fletcher, E., and Kennedy, L. (2004). Optic nerve head segmentation. *IEEE Transactions on Medical Imaging*, 23(2):256–264.
- Osareh, A., Mirmehdi, M., Thomas, B., and Markham, R. (2001). Automatic recognition of exudative maculopathy using fuzzy c-means clustering and neural networks. In Claridge, E. and Bamber, J., editors, *Medical Image Understanding and Analysis*, pages 49–52. BMVA Press.
- Osher, S. and Sethian, J. A. (1988). Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79:12–49.
- Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639.
- Porta, M. and Bandello, F. (2002). Diabetic retinopathy a clinical update. *Diabetologia*, 45(12):1617–1634.
- Sapiro, G. (2001). *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press.
- Sinthanayothin, C., Boyce, J., Williamson, T., Cook, H., Mensah, E., and Lal, S. and Usher, D. (2002). Automated detection of diabetic retinopathy on digital fundus images. *Diabetic Medicine*, 19:105–112.
- Srolovitz, D. (2005). Level set method simulation of surface and defect evolution in materials.
- Suri, J., Liu, K., Singh, S., Laxminarayan, S., Zeng, X., and Reden, L. (2002). Shape recovery algorithms using level sets in 2-d/3-d medical imagery: A state-of-the-art review. *IEEE Transactions On Information Technology In Biomedicine*, 6:8–28.
- Wang, H., Hsu, W., Goh, K., and Lee, M. (2000). An effective approach to detect lesions in color retinal images. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 181–186.
- Wang, L., Bhalerao, A., and Wilson, R. (2004). Robust modelling of local image structures and its application to medical imagery. In *ICPR04*, pages III: 534–537.
- Ward, N., Tomlinson, S., and Taylor, C. J. (1989). Image analysis of fundus photographs: the detection and measurement of exudates associated with diabetic retinopathy. *Ophthalmology*, 96(1):80–86.