

Segmentation of skin cancer images

L. Xu^a, M. Jackowski^a, A. Goshtasby^{a,*}, D. Roseman^b, S. Bines^b, C. Yu^c, A. Dhawan^d, A. Huntley^e

^aDepartment of Computer Science and Engineering, Wright State University, Dayton, OH 45435, USA

^bDepartment of General Surgery, Rush-Presbyterian-St. Luke's Medical Center, Chicago, IL, USA

^cDepartment of Electrical Engineering and Computer Science, University of Illinois, Chicago, IL, USA

^dDepartment of Bioengineering, University of Toledo, OH, USA

^eDepartment of Dermatology, University of California, Davis, CA, USA

Received 2 June 1997; received in revised form 8 December 1997; accepted 29 December 1997

Abstract

An automatic method for segmentation of images of skin cancer and other pigmented lesions is presented. This method first reduces a color image into an intensity image and approximately segments the image by intensity thresholding. Then, it refines the segmentation using image edges. Double thresholding is used to focus on an image area where a lesion boundary potentially exists. Image edges are then used to localize the boundary in that area. A closed elastic curve is fitted to the initial boundary, and is locally shrunk or expanded to approximate edges in its neighborhood in the area of focus. Segmentation results from 20 randomly selected images show an average error that is about the same as that obtained by four experts manually segmenting the images. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Image segmentation; Double thresholding; Boundary extraction; Elastic curve; Melanoma; Pigmented lesion

1. Introduction

Skin cancers are the most common form of cancers in humans [1]. The American Cancer Society estimates that more than 700 000 new skin cancers are diagnosed annually in the United States alone [2]. Skin cancers can be classified into melanoma and non-melanoma. Although melanomas are much less common than non-melanomas, they account for most of the mortality from skin cancers [2]. Detection of malignant melanoma in its early stages considerably reduces morbidity and mortality. Early detection also saves hundreds of millions of dollars that otherwise would be spent on the treatment of advanced diseases [3].

If cutaneous melanoma is detected in its early stages and removed, there is a very high likelihood that the patient will survive [4,5]. Clinical features of pigmented lesions suggestive of melanoma are what are known as the ABCDs of melanoma [3]: asymmetry, border irregularity, color variegation, and diameter greater than 6 mm. Image analysis techniques for measuring these features have been developed [6]. Measurement of image features for diagnosis of melanoma requires that first the lesions be detected and localized in an image. It is essential that lesion boundaries are determined accurately so that measurements, e.g.

maximum diameter, asymmetry, irregularity of the boundary, and color characteristics can be accurately computed. For delineating lesion boundaries, various image segmentation methods have been developed. These methods use color and texture information in an image to find the lesion boundaries.

To segment a skin image into lesions, Umbaugh et al. [7] transformed the RGB color space into a spherical color space with coordinates defined by $L = \sqrt{R^2 + G^2 + B^2}$, $A = \cos^{-1}(B/L)$ and $B = \cos^{-1}(R/L \sin(A))$. By quantizing the AB space into four colors, they were then able to partition a color image into different regions and isolate lesions from the background. In a separate study, Umbaugh et al. [8] developed a principal-components transform in a user-selected color space to segment skin cancer images. Green et al. [9] segmented a color image by first finding the average color of a small area of a lesion and the average color of a small area of the background interactively. Then, by mapping the image colors to the vector connecting the two average colors, they obtained a histogram. The color corresponding to the valley between the two peaks in the histogram was then used as the threshold value to segment the image. Dhawan and Sicsu [10] used image gray values and textures separately to segment skin images. They then combined the results to obtain the lesion boundaries. Hance et al. [11] compared the accuracy of six different color

* Corresponding author.

segmentation techniques, and found that when two or more of the techniques are combined, an accuracy that is considerably higher than the accuracy of any one of the individual techniques will be obtained.

Image segmentation is perhaps the most studied area in computer vision, with numerous methods reported [12,13]. A segmentation method is usually designed taking into consideration the properties of a particular class of images. In this paper, we develop a three-step segmentation method using the properties of skin cancer images. The steps of our method are as follows:

1. Preprocessing: a color image is first transformed into an intensity image in such a way that the intensity at a pixel shows the color distance of that pixel with the color of the background. The color of the background is taken to be the median color of pixels in small windows in the four corners of the image.
2. Initial segmentation: a threshold value is determined from the average intensity of high gradient pixels in the obtained intensity image. This threshold value is used to find approximate lesion boundaries.
3. Region refinement: a region boundary is refined using edge information in the image. This involves initializing a closed elastic curve at the approximate boundary, and shrinking and expanding it to fit to the edges in its neighborhood.

In the following, the details of the steps of our method are provided.

2. Preprocessing

The first step in our image segmentation method can be considered a preprocessing operation that transforms a color image into an intensity image. This operation is motivated by two observations:

1. Skin lesions come in a variety of colors; therefore, absolute colors are not very useful in segmenting images. However, changes in color from a lesion to its

background (its surrounding healthy skin) are similarly observed in all images; therefore, changes in color can be used to effectively segment images.

2. When segmenting a skin image, significant color variations may exist within a lesion or in the background. Such variations should be suppressed since our interest is in color changes from the background to a lesion or from a lesion to the background.

Observation 1 suggests that we should use changes in color rather than absolute colors to segment images. Therefore, we transform pixel colors that are vector quantities into intensities that are scalars and represent color differences. Observation 2 states that, among the color changes, only those belonging to a lesion boundary are important in image segmentation, and color changes inside a lesion or in the background should be ignored.

We transform our images that are in RGB color coordinates into images that are in CIELAB or CIE 1976 $L^*a^*b^*$ color coordinates [14]. CIELAB is a color space standardized by the CIE (Commission Internationale de l'Éclairage) in 1976 to measure color differences. This is a uniform color space defined in such a way that Euclidean distance between two colors (defined as ΔE) is proportional to their visual difference. Color in the CIELAB space can be described with less redundancy than in the RGB space. RGB color coordinates can be transformed into $L^*a^*b^*$ color coordinates using the following formulae [14]:

$$L^* = 116(Y/Y_n)^{1/3} - 16, (Y/Y_n) > 0.008856$$

$$L^* = 903.3(Y/Y_n), (Y/Y_n) \leq 0.008856$$

$$a^* = 500[f(X/X_n) - f(Y/Y_n)]$$

$$b^* = 200[f(Y/Y_n) - f(Z/Z_n)]$$

where $f(t) = t^{1/3}$ when $t > 0.008856$ and $f(t) = 7.787t + 16/116$ when $t \leq 0.008856$. X_n , Y_n and Z_n are the coordinates of the CIELAB reference white, which are usually chosen to be 0.9642, 1.0 and 0.8249, respectively.

If we require that the images be taken such that lesions do




Fig. 1. (a) A color image showing an atypical lesion. (b) Image obtained after mapping colors into intensities in such a way that the intensity at a pixel is proportional to the CIELAB color distance of the pixel to the color of the background. (c) Gradient magnitudes of (b) obtained by the Sobel operator. Larger gradients are shown brighter.

not fall on image corners, we can then use colors in the four corners of an image to estimate the color of the background. We take small windows, typically 10×10 pixels in size, from the four corners of an image and determine the median L^* , a^* and b^* of the pixels. We use this median color as an estimate to the color of the background. We use median color rather than average color because image averaging uses the hair colors as well as the skin colors to estimate the color of the background. Since the number of hair pixels is usually much smaller than the number of skin pixels in an image, when the median color is used, the color of a pixel belonging to the hair will not be used and the color of a pixel belonging to the skin will be used to estimate the color of the background.

If ΔL^* , Δa^* , Δb^* , correspondingly, show the differences of L^* , a^* , b^* color components of an image pixel and its background, the color distance between the pixel and background can be computed from [15]:

$$\Delta E = \sqrt{\Delta L^{*2} + \Delta a^{*2} + \Delta b^{*2}}$$

If the intensities assigned to pixels are proportional to color distances of the pixels to the color of the background, we will obtain an image that has high values in lesions and small values in the background. An image generated in this manner will, therefore, show lesions as bright spots. This is demonstrated in Fig. 1. Fig. 1(a) shows an image of a skin lesion, and (b) shows the image representing the distances of pixel colors to the color of the background. This mapping not only transforms a color image into an intensity image, it assigns intensities to the pixels in such a way that a lesion can be distinguished from the background.

The image of Fig. 1(b) is obtained from the image of Fig. 1(a) after implementing observation 1. Note that higher intensity pixels in Fig. 1(b) correspond to pixels in Fig. 1(a) whose colors are more dissimilar to the color of the background. Therefore, they are more likely to belong to a lesion.

Image gradients have been used in the past to determine region boundaries [16]. However, as Fig. 1(c) shows, detection of lesion boundaries using pure image gradients is a difficult task. We need to include observation 2 in the preprocessing operation so that edges on lesion boundaries are distinguished from edges inside a lesion or in the background.

To implement observation 2, we will need a function that provides the property shown in Fig. 2(a). For a wide range of intensities in the background, this function produces very similar intensities. Therefore, the function reduces image gradients in the input corresponding to details in the background. Similarly, it reduces image gradients belonging to a lesion. For intensities falling on lesion boundaries, however, we see that gradients are increased. Therefore, if we map image intensities according to the function depicted in Fig. 2(a), we will increase gradients on lesion boundaries, while decreasing gradients inside a lesion or in the

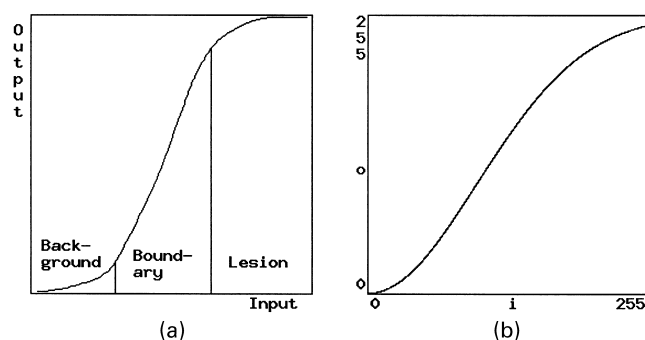


Fig. 2. (a) A desirable function for mapping color distances to image intensities. (b) Approximation of function (a) by $o = 1/\sqrt{2\pi\sigma}(1 - \exp(i^2/2\sigma^2))$. i and o show the input and output image intensities, respectively.

background. Mapping intensities in this manner facilitates detection of lesion boundaries.

We use a Gaussian function to achieve the property demonstrated in Fig. 2(a). If we let $G(x) = 1/\sqrt{2\pi\sigma} \exp(-x^2/2\sigma^2)$, we see in Fig. 2(b) that $F(x) = 1/\sqrt{2\pi\sigma} - G(x)$ has a shape similar to that in Fig. 2(a). The standard deviation of this Gaussian (σ) depends on intensity variations in the background. An image containing large intensity variations in the background requires a larger standard deviation than an image containing small intensity variations. Since windows in the four corners of an image are used to characterize the background, we use the standard deviation of intensities in the four windows as the standard deviation in function F . The more detailed the background, the larger the standard deviation needed to suppress intensity variations in the background.

Fig. 3(a) shows mapping of the intensities of Fig. 1(b) according to the function shown in Fig. 2(b). As can be observed, most details within the lesion and some details in the background have been suppressed, while variations from the lesion to the background and from the background to the lesion have been enhanced. We segment the image of Fig. 3(a) instead of the image of Fig. 1(a) to isolate lesion boundaries. Note that the preprocessing operation not only reduces a color image into an intensity image, it enhances the boundary of a lesion while suppressing details inside and outside a lesion.




Fig. 3. (a) Transforming intensities of Fig. 1(b) according to the function of Fig. 2(b). (b) Smoothing of (a) with a 2D Gaussian kernel of standard deviation 2 pixels.

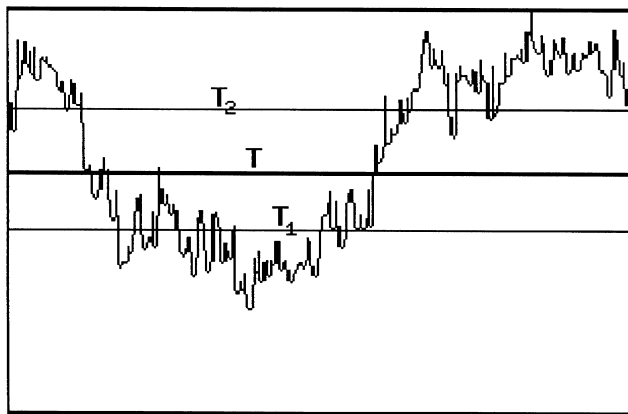


Fig. 4. Intensities along an image scanline, and the relation between the initial threshold value T and the double threshold values T_1 and T_2 .

3. Initial segmentation

To reduce the effect of image noise and intensity variations due to skin's repetitive texture and hair, an image is first low-pass filtered before being segmented. Fig. 3(b) shows the image of Fig. 3(a) after being smoothed with a 2D Gaussian kernel of standard deviation 2 pixels. As can be observed, although smoothing reduce details in the image, the smoothed image still contains information about the lesion, which is brighter than the background. The objective in the initial segmentation is to determine the approximate position and shape of a lesion, and then use double thresholding to narrow in on an image area where the optimal lesion boundary exists. Since the optimal threshold value at one boundary point may differ from that at another boundary point, the objective in double thresholding is to select a range of threshold values that includes the optimal threshold value at every boundary point.

Double thresholding also reduces the number of noisy regions obtained as a result of intensity thresholding. Consider an image scanline, shown in Fig. 4. The horizontal axis shows pixels in the scanline, while the vertical axis shows the intensities of the pixels. As can be observed, if the threshold value is not selected properly, noisy regions

from intensity variations in the lesion or background will be obtained. If there were no details from hair or skin texture, or if there were no intensity variations inside a lesion, a single threshold value T would have been sufficient to isolate a lesion from its background. However, since image variations from hair and skin texture usually exist in an image, a single threshold value may detect noisy regions from hair and skin texture. If such regions are close to each other, they may merge and create larger regions. By increasing the threshold value, we will observe that the number and size of noisy regions in the background will decrease. If we decrease the threshold value, we will see that the number and size of noisy regions in the lesion will decrease. The use of two threshold values will, therefore, make it possible to obtain rather noise-free regions for both the lesion and background. Double thresholding will produce a segmentation that either will be free of noisy regions, or will contain fewer and smaller noisy regions than when a single threshold value is used. Double thresholding, however, requires the use of an initial threshold value.

To determine an initial threshold value automatically, we observe that gradients of pixels on lesion boundaries are generally higher than gradients of pixels inside or outside lesions. We will, therefore, use the average intensity of the top $p\%$ highest gradient pixels in the image to compute the threshold value. p is typically a small number, e.g. 5. Because noise and details from skin texture and hair could also result in high gradients, this process may detect details from noise, skin texture and hair. Such regions, however, are often small and can be removed. In Fig. 4, the average intensity of the top $p\%$ highest gradient pixels in an image is denoted by T .

The choice of p depends on the given image. For an image containing a large lesion, this percentile should be taken larger than for an image containing a small lesion. Since we do not know the size of a lesion ahead of time, we cannot select the best percentile for every image. The threshold values computed under variations of this percentile tend to change very little though. If we compute the average intensities of the 3, 6, 10 and 15% highest gradient pixels in Fig. 3(b), we find threshold values equal to 124,




Fig. 5. (a)–(d) Thresholding the image of Fig. 3(b) at intensities equal to the average intensities of the 3, 6, 10 and 15% highest gradient pixels in the image, respectively.

125, 131 and 148, respectively. The results of segmenting Fig. 3(b) with these threshold values are shown in Fig. 5(a)–(d). As can be observed, because of the considerable intensity differences between the lesion and background obtained after the intensity transformation, the regions obtained at various threshold values are about the same shape and size. Any error made in this initial segmentation will be corrected later when the regions are refined.

For a noise-free image, the threshold value T obtained in this manner will represent the average intensity of pixels with locally maximum gradients. Due to noise and image details from skin texture and hair, this threshold value may not detect the optimal lesion boundary, and points in the detected boundary may be displaced from their true position. If we assume that the threshold value T approximately determines the boundary of a lesion, and an optimal boundary point is in the neighborhood of the estimated one, we can then refine the obtained boundary using image edges.

Assuming that a point on the lesion boundary obtained by intensity thresholding needs to be displaced by d pixels to fall on the optimal boundary pixel, parameter d depends on the gradient at the pixel. If the boundary pixel has a high gradient, a small d would be sufficient, while if the pixel has a low gradient, a larger d would be necessary. This is because when the gradient is low, by changing the threshold value slightly, the boundary point will move more than when the gradient is high. Suppose that the average movement cannot exceed d pixels. We will then determine threshold values T_1 and T_2 as follows: If $h[i]$ denotes the number of pixels in an image with intensity i , for $i = 0, \dots, 255$, and if T is the threshold value determined by the average intensity of the $p\%$ highest gradient pixels in the image, we will determine intensities T_1 and T_2 in such a way that

$$d \times n = \sum_{i=T_1}^{T-1} h[i] = \sum_{i=T+1}^{T_2} h[i]$$

where n is the number of pixels in the region boundaries obtained by thresholding the image at intensity T . If there are fewer than $d \times n$ pixels in the image with intensities smaller than T , we will let $T_1 = 0$, while if fewer than $d \times n$ pixels exist in the image with intensities larger than T , we will let $T_2 = 255$. Note that d is the average correction anticipated of boundary pixels obtained by thresholding the image at intensity T . Parameter d will be one of the input parameters of the proposed method.

Fig. 6 demonstrates the use of double thresholds T_1 and T_2 in defining a mask inside which the optimal lesion boundary is expected to exist. Fig. 6(a) and (b) shows regions obtained by thresholding the image of Fig. 3(b) at $T_1 = 131$ and $T_2 = 195$. By Exclusive-Oring Fig. 6(a) and (b), we obtain the image of (c), which defines a mask inside which the optimal lesion boundary should be found. We will use this mask as our search area to find the optimal lesion boundary using the image edges. Parameter d shows the average of




Fig. 6. (a), (b) Thresholding the image of Fig. 3(b) with threshold values $T_1 = 131$ and $T_2 = 195$, respectively. T_1 and T_2 were obtained by setting parameter $d = 10$. (c) Mask obtained by Exclusive-Oring images (a) and (b). (d), (e) Same as (c), except that d was set to 5 and 15, respectively.

displacements needed at the approximate boundary points to obtain the optimal boundary points. If this parameter is too small, some parts of the optimal boundary may be missed, while if it is too large, detection of the optimal boundary may become difficult, and a suboptimal boundary may be found as a result. This parameter is image-dependent and has to be selected using experimental results. Fig. 6(d) and (e) shows the masks obtained when $d = 5$ and $d = 15$, respectively. We will investigate the dependency of the segmentation result on this parameter later in this paper.

4. Region refinement

The double thresholding process described in the preceding section determines an image area where a lesion boundary obtained from a range of threshold values will exist. Since the best threshold value in one local area may be different from the best threshold value in another local area, this range of threshold values is expected to include the best threshold value for all boundary pixels. We will assume that an optimal threshold value produces a boundary pixel that has a locally maximum gradient magnitude. Therefore, we will move an initial boundary pixel to the pixel in its neighborhood having a locally maximum gradient magnitude.

Region boundaries in an image are best described by pixels with locally maximum gradient magnitudes [17–20]. Pixels with locally maximum gradient magnitudes can be determined without any user interaction; therefore, the process is automatic. Locally maximum gradients in an image, however, not only represent lesion boundaries, they also represent small details inside and outside a lesion. In addition, an obtained boundary may merge with another boundary due to image noise and produce a false lesion boundary. Therefore, although detected edges from locally maximum gradient magnitudes show sharp changes in

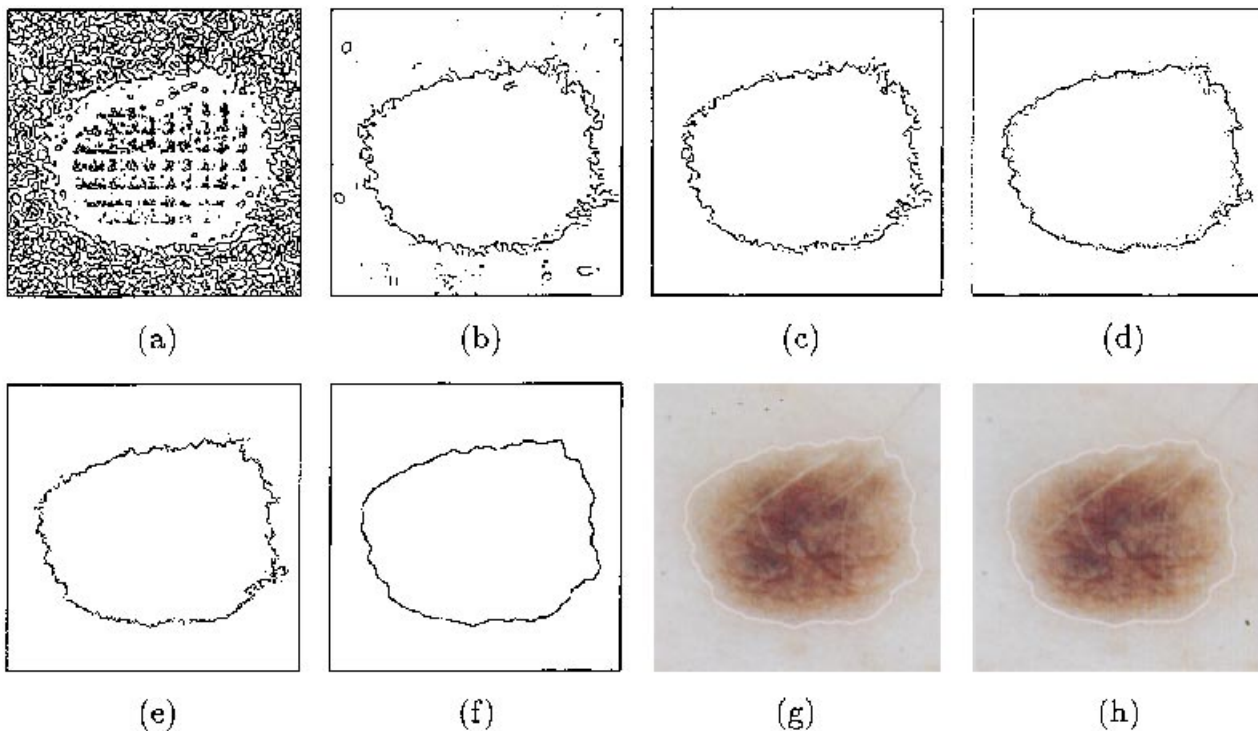


Fig. 7. (a) Authentic edges [18] of Fig. 3(b). (b) Edges of (a) corresponding to the top 10% highest gradients in the image. (c) Edges of (a) falling in the mask of Fig. 6(c). (d) Initial boundary shown in Fig. 5(c) when overlaid with (c). (e) Edges in (d) obtained by expanding or shrinking the initial contour. (f) A rational Gaussian curve fitting the points in (e). (g) Overlaying the curve obtained in (f) with the original image shown in Fig. 1(a). This is the final segmentation result when using $p = 8\%$, $d = 10$ pixels and $\sigma = 2$ pixels. Increasing σ to 3 pixels, we obtain the segmentation result shown in (h).

image intensities, the obtained edges by themselves do not always represent true region boundaries.

The locally maximum gradient pixels of Fig. 3(b) are shown in Fig. 7(a). Although many of the noisy contours can be eliminated by removing the small magnitude edges [19], the obtained result may not represent a closed lesion boundary, as shown in Fig. 7(b). Since we know the approximate position of a lesion boundary from double thresholding, we can filter out the edge points that do not contribute to the lesion boundary and eliminate them. Fig. 6(c) shows the image area obtained by two threshold values, $T_1 = 131$ and $T_2 = 195$. We use Fig. 6(c) as a mask to select from among all locally maximum gradient pixels those that potentially contribute to the lesion boundary. Fig. 7(c) shows such edges. Although many points on the lesion boundary have been detected, we see that the boundary is disconnected and that some points not belonging to the boundary have also been detected. By combining double thresholding and edge detection, we have been able to focus on a small fraction of image edges that lie on or near the initial lesion boundary.

From obtained edges in Fig. 7(c), we must determine those that belong to the lesion boundary. Since the obtained edges may not form a closed contour, the contour segments must be connected to construct a closed boundary. The approach we choose uses a closed elastic curve [21]. When the curve is initialized at the boundary obtained by segmenting the image with threshold T , the curve will

expand or shrink locally to fit to the edges and form a lesion boundary. Fig. 7(d) shows the lesion boundary obtained by intensity thresholding [Fig. 5(c)] when initialized at the edges of Fig. 7(c). For each point in the initial contour, the edge point closest to it in the direction normal to the initial contour is determined. If for a point on the initial contour an edge point in the direction normal to the contour is not found, the closest edge point is selected. The edges obtained in this manner are shown in Fig. 7(e). These edges will be the only ones used to obtain the final curve. The entire set of edges in the mask area is not used in curve fitting, because in some areas in the mask a large number of edges may be obtained, pulling the curve unnecessarily towards themselves. When a closed elastic curve [21] is fitted to points in Fig. 7(e), we obtain the boundary contour shown in Fig. 7(f). Fig. 7(g) shows the curve of Fig. 7(f) when overlaid with the image of Fig. 1(a). This is the final segmentation result. The elasticity of the curve may be varied to obtain boundaries with different smoothnesses. Fig. 7(h) shows a smoother lesion boundary obtained by decreasing the elasticity of the curve.

The elastic curve model used here is based on the rational Gaussian formulation [21]. Given a set of points $\{\mathbf{V}_i; i = 1, \dots, n\}$ along a contour, the rational Gaussian curve that fits the points is defined by

$$\mathbf{P}(u) = \sum_{i=1}^n \mathbf{V}_i g_i(u), \quad u \in [0, 1]$$

where $g_i(u)$ is the i th basis function of the curve given by

$$g_i(u) = \frac{G_i(u)}{\sum_{j=1}^n G_j(u)}$$

and $G_i(u)$ is a Gaussian:

$$G_i(u) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(u-u_i)^2}{2\sigma^2}}$$

u_i is the i th node of a curve determined from

$$u_i = \frac{|\mathbf{V}_i - \mathbf{V}_{i+1}|}{\sum_{j=1}^n |\mathbf{V}_j - \mathbf{V}_{j+1}|}, \quad i \in [1, n]$$

where $\mathbf{V}_{n+1} = \mathbf{V}_1$ and $|\mathbf{V}_i - \mathbf{V}_{i+1}|$ are the Euclidean distances between points \mathbf{V}_i and \mathbf{V}_{i+1} . σ is the standard deviation of Gaussians in a curve and shows the smoothness (inverse elasticity) of the curve. The larger the standard deviation, the smoother the curve. The standard deviation used in curve fitting should be equal to or greater than the standard deviation used in image smoothing. When an image is smoothed, both the boundary obtained from intensity thresholding and the contours obtained from image edges will be smoothed. Reducing the standard deviation in curve fitting to a value that is smaller than the standard




Fig. 8. Twenty images of skin lesions. Images 1–6, 7–13 and 14–20 show atypical, benign and malignant lesions, respectively. These images are digitized clinical photographs from our database of skin lesions.

deviation used in image smoothing will not make the boundary contour more detailed. To obtain a more detailed boundary contour, the standard deviation used in both image smoothing and curve fitting should be reduced.

The curve obtained in this manner will pass near locally maximum gradients and will connect edge fragments

Table 1
Segmentation errors obtained by varying input parameters σ , p and d

σ	p	d	Expert 1	Expert 2	Expert 3	Expert 4	
2	3	3	0.0694	0.0819	0.0964	0.0989	
		5	0.0687	0.0812	0.0966	0.0982	
		7	0.0656	0.0814	0.0977	0.0982	
	6	3	0.0694	0.0815	0.0966	0.0984	
		5	0.0733	0.0835	0.1002	0.0989	
		7	0.0655	0.0810	0.0979	0.0976	
	10	3	0.0694	0.0821	0.0974	0.0986	
		5	0.0700	0.0796	0.0970	0.0979	
		7	0.0655	0.0817	0.0988	0.0978	
	3	3	3	0.0727	0.0845	0.0978	0.1021
			5	0.0717	0.0833	0.0981	0.1008
			7	0.0714	0.0830	0.0981	0.1004
6		3	0.0721	0.0840	0.0974	0.1014	
		5	0.0712	0.0829	0.0978	0.1001	
		7	0.0707	0.0824	0.0977	0.0996	
10		3	0.0717	0.0837	0.0980	0.1105	
		5	0.0708	0.0825	0.0982	0.0992	
		7	0.0703	0.0821	0.0981	0.0988	
4		3	3	0.0744	0.0864	0.0990	0.1041
			5	0.0732	0.0852	0.0990	0.1028
			7	0.0729	0.0849	0.0989	0.1025
	6	3	0.0738	0.0856	0.0985	0.1030	
		5	0.0725	0.0844	0.0986	0.1018	
		7	0.0721	0.0841	0.0984	0.1014	
	10	3	0.0729	0.0846	0.0984	0.1017	
		5	0.0716	0.0835	0.0986	0.1004	
		7	0.0714	0.0832	0.0986	0.1002	

A table entry under a particular expert was obtained by overlaying an automatically segmented lesion with the lesion manually traced by that expert. The ratio of the sum of the lesion areas that did not overlap and the sum of the lesion areas were then computed, and the average of such ratios for the 20 images found. Experts 1 and 2 were surgeons, Expert 3 was a dermatologist, and Expert 4 was a bioengineer.

obtained in the lesion boundary. This refinement step is completely automatic and does not require any user interaction. However, the user is required to provide values for σ : the standard deviation of the Gaussian used in image smoothing and curve fitting, $\%p$: the percentile highest $\%$ gradient pixels used in determining the initial threshold value T , and d : the average displacement anticipated of the boundary points obtained from threshold T to obtain the optimal boundary points (d is used to determine T_1 and T_2). These parameters are image dependent, and have to be selected based on the properties of the class of images used in the segmentation. In the following, the sensitivity of the proposed segmentation to its parameters is investigated.

5. Sensitivity of the segmentation to input parameters

To determine the sensitivity of the proposed segmentation

method to its input parameters, we use 20 representative skin lesions, as shown in Fig. 8. Images 1–6 are atypical lesions, images 7–13 are benign lesions and images 14–20 are melanoma lesions. The boundaries of some of the lesions are not well defined, and some images contain hair and details from skin texture. Fig. 9 shows the lesion boundaries manually traced by four experts. Fig. 9(a) and (b) were traced by two surgeons, and Fig. 9(c) and (d) were traced by a dermatologist and bioengineer, respectively. It is interesting to note that the experts have segmented the images differently, and most differences exist in segmenting the atypical lesions [image 2 of Fig. 9(a) and (c), image 4 of Fig. 9(c) and (d), and image 5 of Fig. 9(c) and (d)]. This raises the question as to where the actual boundaries are, and it appears that different experts use different rules to segment an image. Some set the boundary at the junction of the lesion with normal skin, while others set it at the junction of darker pigmentation with lighter pigmentation, even though the lighter areas were clearly not normal




Fig. 9. Lesion boundaries of images of Fig. 8 manually traced by (a), (b) two surgeons, (c) a dermatologist, and (d) a bioengineer.


skin. The observer variation in placing boundaries may well reflect the fact that for some lesions the real boundary is not perfectly vertical and may therefore have some width. We found that variation of lesion boundaries were the smallest among benign lesions where lesion boundaries were the least fuzzy.

Segmenting these images with our segmentation program with parameters $\sigma = 2, 3, 4$, $p = 3, 6, 10$ and $d = 5, 10, 15$, and comparing our results with those obtained manually by the experts, we get the results shown in Table 1. To obtain a table entry, first the boundary automatically determined for the i th lesion is overlaid with the corresponding boundary manually traced by expert j . Then the sum of the two areas that did not overlap, a_{ij} , and the sum of the lesion areas, b_{ij} , are determined. Next, the ratio of the two areas $r_{ij} = a_{ij}/b_{ij}$ is computed. r_{ij} will be called the error of the automatic method when segmenting the i th lesion and using the result from the j th expert as the standard. This error measure was used, because when segmentation results from our program and from an expert are exactly the same, an error of 0.0 will be obtained, while when the two segmentation results do not have any overlap, an error of 1.0 will be obtained. This error measure is always between 0.0 and 1.0, and is independent of the size of a lesion. The average of these ratios for the 20 lesions ($i = 1, \dots, 20$) was determined for each expert and entered into Table 1. The errors in Table 1 show that our segmentation is stable under variations of its parameters and that changing the input parameters slightly will not change the segmentation result drastically. Segmentation errors measured with respect to different experts are considerably different. Errors obtained using results from the two surgeons were smaller than those obtained from the dermatologist or bioengineer. The overall best and worst segmentation results obtained by our program according to Table 1 are shown in Fig. 10. These images show that even with two extreme sets of input parameters, the segmentation results are not very different and the method is stable under the variation of its parameters.


We also determined the variability among experts and tabulated the results in Table 2. First, the i th lesion boundary segmented by experts j and k was overlaid, and the ratio of the sum of areas non-overlapping and sum of lesion areas was determined. Then the average of the ratios for all lesions ($i = 1, \dots, 20$) was determined. The obtained average ratio was then entered into entry (j, k) of the table as the average of the difference between experts j and k when segmenting the 20 images. This shows that, although the variation between the two surgeons is small, the overall average variation between the experts is slightly higher than the average variation between our program and the four experts.

6. Conclusions

Image segmentation is the first step in many image analysis problems. To analyze skin lesions, it is necessary



(a)



(b)

Fig. 10. (a) Best automatic segmentation result. (b) Worst automatic segmentation result.

to accurately locate and isolate the lesions. In this paper, an automatic method for segmentation of skin cancer images was presented. The method starts with an initial segmentation and uses edge information in the neighborhood of the initial segmentation to refine the results. An elastic curve model is used to represent the final segmentation.

Although the method is devised for segmentation of color images, early on in processing, a color image is transformed into an intensity image where the intensity at a pixel shows the color distance of that pixel to the background. Intensities in the image obtained in this manner are then transformed according to a function shown in Fig. 2 to suppress details in the background and in a lesion while enhancing details across lesion boundaries. Transformation of a color image into an intensity image and mapping of image intensities to enhance lesion boundaries are considered to be the main contributions of this work.

Table 2
Variations among experts when segmenting the images of Fig. 8

kj	Expert 1	Expert 2	Expert 3	Expert 4
Expert 1	0	0.0502	0.1023	0.0870
Expert 2	0.0502	0	0.1186	0.0863
Expert 3	0.1023	0.1186	0	0.0976
Expert 4	0.0870	0.0863	0.0976	0

The entry corresponding to experts j and k is obtained by computing the sum of non-overlapping areas over the sum of areas manually traced by experts j and k for each lesion, and finding the average of such ratios for the 20 images. Experts 1 and 2 were surgeons, Expert 3 was a dermatologist, and Expert 4 was a bioengineer.

The proposed segmentation has three parameters. One parameter specifies the standard deviation of the Gaussian smoother. An image with considerable skin texture, hair and specularly requires a larger smoothing parameter than an image with very little skin texture, hair and specularly. Another parameter is linked to image gradients and determines the threshold value needed to find the initial lesion boundaries. The third parameter depends on the sharpness of color changes across lesion boundaries and is provided in terms of the maximum displacement needed to reposition a boundary pixel obtained in the initial segmentation to its final position. The algorithm is designed such that slight variations in the input parameters will not change the final segmentation result drastically. Experimental results show that the segmentation result remains stable under a wide range of input parameters.

The proposed segmentation requires less than 1 min on a Sun Sparc V computer to segment an image of 256×256 pixels. The output from this program is being used to quantify the size, irregularity, asymmetry and color variation of skin lesions.

Acknowledgements

The authors would like to thank the reviewers for their insightful comments and suggestions. The support of the National Science Foundation under grants IRI-9529045 and IRI-9509253 is also greatly appreciated. The images used in this study were provided by the Department of General Surgery, Rush-Presbyterian-St. Luke's Medical Center, Chicago, and the Department of Dermatology, University of California, Davis.

References

- [1] A.W. Kopf, T.G. Salopek, J. Slade, A.A. Marghood, R.S. Bart, Techniques of cutaneous examination for the detection of skin cancer, *Cancer Supplement* 75 (2) (1994) 684–690.
- [2] D.E. Elder, Skin cancer: Melanoma and other specific nonmelanoma skin cancers, *Cancer Supplement* 75 (1) (1994) 245–256.
- [3] NIH Consensus Conference, Diagnosis and treatment of early melanoma, *JAMA* 268 (10) (1992) 1314–1319.
- [4] A.J. Sober, T.B. Fitzpatrick, M.C. Mihm, Early recognition of cutaneous melanoma, *JAMA* 242 (1979) 2795–2799.
- [5] M.M. Wick, A.J. Sober, T.B. Fitzpatrick, Clinical characteristics of early cutaneous melanoma, *Cancer* 45 (1980) 2684–2686.
- [6] W.V. Stoecker, W.W. Li, R.H. Moss, Automatic detection of asymmetry in skin tumors, *Computerized Medical Imaging and Graphics* 16 (3) (1992) 191–197.
- [7] S.E. Umbaugh, R.H. Moss, W.V. Stoecker, Automatic color segmentation of images with applications in detection of variegated coloring in skin tumors, *IEEE Engng Med. Biol.* 8 (1989) 43–52.
- [8] S.E. Umbaugh, R.H. Moss, W.V. Stoecker, An automatic color segmentation algorithm with application to identification of skin tumor borders, *Computerized Medical Imaging and Graphics* 16 (3) (1992) 227–235.
- [9] A. Green, N. Martin, J. Pfitzner, M. O'Rourke, N. Knight, Computer image analysis in the diagnosis of melanoma, *J. American Academy of Dermatology* 31 (6) (1994) 958–964.
- [10] A.P. Dhawan, A. Sicsu, Segmentation of images of skin lesions using color and texture information of surface pigmentation, *Computerized Medical Imaging and Graphics* 16 (3) (1992) 163–177.
- [11] G.A. Hance, S.E. Umbaugh, R.H. Moss, W.V. Stoecker, Unsupervised color image segmentation with application to skin tumor borders, *IEEE Engineering in Medicine and Biology* 15 (1) (1996) 104–111.
- [12] R.M. Haralick, L.G. Shapiro, Image segmentation techniques, *Computer Vision Graphics, and Image Processing* 29 (1) (1985) 100–132.
- [13] R.K. Sahoo, S. Soltani, A.K.C. Wong, A survey of thresholding techniques, *Computer Vision, Graphics, and Image Processing* 41 (1988) 233–260.
- [14] J.K. Kasson, W. Plouffe, An analysis of selected computer interchange color spaces, *ACM Trans. Graphics* 11 (4) (1992) 373–405.
- [15] B. Hill, Th. Roger, F.W. Vorhagen, Comparative analysis of the quantization of color spaces on the basis of the CIELAB color-difference formula, *ACM Trans. Graphics* 16 (2) (1997) 109–154.
- [16] V.S. Nalwa, T.O. Binford, On detecting edges, *IEEE Trans. Pattern Anal. Machine Intell.* 8 (6) (1986) 600–714.
- [17] D. Marr, E. Hildreth, Theory of edge detection, *Proc. Royal Society of London B-207* (1980) 187–217.
- [18] J.J. Clark, Authenticating edges produced by zero crossing algorithms, *IEEE Trans. Pattern Anal. Machine Intell.* 11 (1) (1989) 43–57.
- [19] M.M. Fleck, Some defects in finite difference edge finders, *IEEE Trans. Pattern Anal. Machine Intell.* 14 (3) (1992) 337–346.
- [20] J. Canny, A computational approach to edge detection, *IEEE Trans. Pattern Anal. Machine Intell.* 8 (6) (1986) 679–698.
- [21] A. Goshtasby, Geometric modeling using rational Gaussian curves and surfaces, *Computer-Aided Design* 27 (5) (1995) 363–375.