

Approximate lesion localization in dermoscopy images

M. Emre Celebi¹, Hitoshi Iyatomi², Gerald Schaefer³ and William V. Stoecker⁴

¹Department of Computer Science, Louisiana State University, Shreveport, LA, USA, ²Department of Electrical Informatics, Hosei University, Tokyo, Japan, ³School of Engineering and Applied Science, Aston University, Birmingham, UK, and ⁴Stoecker & Associates, Rolla, MO, USA

Background: Dermoscopy is one of the major imaging modalities used in the diagnosis of melanoma and other pigmented skin lesions. Because of the difficulty and subjectivity of human interpretation, automated analysis of dermoscopy images has become an important research area. Border detection is often the first step in this analysis. **Methods:** In this article, we present an approximate lesion localization method that serves as a preprocessing step for detecting borders in dermoscopy images. In this method, first the black frame around the image is removed using an iterative algorithm. The approximate location of the lesion is then determined using an ensemble of thresholding algorithms.

Results: The method is tested on a set of 428 dermoscopy images. The localization error is quantified by a metric that uses dermatologist-determined borders as the ground truth. **Conclusion:** The results demonstrate that the method presented here achieves both fast and accurate localization of lesions in dermoscopy images.

Key words: Skin Cancer – Dermoscopy – Border Detection – Thresholding

© 2009 John Wiley & Sons A/S
Accepted for publication 14 April 2009

MALIGNANT MELANOMA, the most deadly form of skin cancer, is one of the most rapidly increasing cancers in the world, with an estimated incidence of 62,480 and an estimated total of 8420 deaths in the United States in 2008 alone (1). Early diagnosis is particularly important because melanoma can be cured with a simple excision if detected early.

Dermoscopy, also known as epiluminescence microscopy, has become one of the most important tools in the diagnosis of melanoma and other pigmented skin lesions. This non-invasive skin imaging technique involves optical magnification, which makes subsurface structures more easily visible when compared with conventional clinical images (2). This in turn reduces screening errors and provides greater differentiation between difficult lesions such as pigmented Spitz nevi and small, clinically equivocal lesions (3). However, it has also been demonstrated that dermoscopy may actually lower the diagnostic accuracy in the hands of inexperienced dermatologists (4). Therefore, in order to minimize the diagnostic errors that result from the difficulty and subjectivity of visual interpretation, the development of computerized image analysis techniques is of paramount importance (5).

Automated border detection is often the first step in the automated or semi-automated analysis of dermoscopy images (6–10). It is crucial for the image analysis for two main reasons. First, the border structure provides important information for accurate diagnosis, as many clinical features, such as asymmetry, border irregularity, and abrupt border cut-off, are calculated directly from the border. Second, the extraction of other important clinical features such as atypical pigment networks, globules, and blue-white areas, critically depends on the accuracy of border detection. Automated border detection is a challenging task due to several reasons: (i) low contrast between the lesion and the surrounding skin, (ii) irregular and fuzzy lesion borders, (iii) artifacts such as black frames, skin lines, hairs, and air bubbles, (iv) variegated coloring inside the lesion.

A number of methods have been developed for preprocessing dermoscopy images. Most of these focused on the removal of artifacts such as hairs and bubbles. Of the studies dealing with hair removal, Lee et al. (11) and Schmid (6) approached the problem using mathematical morphology. Fleming et al. (5) applied curvilinear structure detection with various constraints

followed by gap filling. More recently, Zhou et al. (12) improved Fleming et al.'s approach using feature guided exemplar-based inpainting. A method for bubble removal was introduced in (5), where the authors utilized a morphological top-hat operator followed by a radial search procedure.

In this article, we present a method for approximate lesion localization in dermoscopy images. First, the black frame around the image is removed using an iterative algorithm. Then, the approximate location of the lesion is determined using an ensemble of thresholding algorithms.

Materials and Methods

Black frame removal

Dermoscopy images often contain black frames that are introduced during the digitization process. These need to be removed because they might interfere with the subsequent lesion localization procedure. In order to determine the darkness of a pixel with $(R, G, \text{ and } B)$ coordinates, the lightness component of the HSL color space (13) is utilized:

$$L = \frac{\max(R, G, B) + \min(R, G, B)}{2} \quad (1)$$

In particular, a pixel is considered to be black if its lightness value is < 20 . Using this criterion, the image is scanned row by row starting from the top. A particular row is labeled as part of the black frame if it contains 60% black pixels. The top-to-bottom scan terminates when a row that contains less than the threshold percentage of pixels is encountered. The same scanning procedure is repeated for the other three main directions. Figure 1 shows the result of this procedure on a sample image.

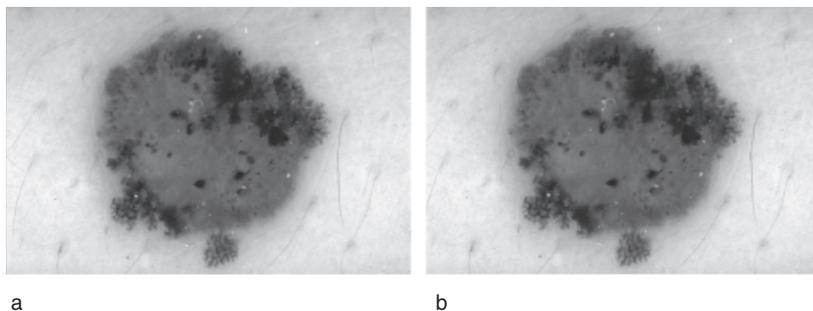


Fig. 1. Black frame removal: (a) Original image, (b) After frame removal.

Approximate lesion localization

Although dermoscopy images can be quite large, the actual lesion often occupies a relatively small area. Therefore, if we can determine the approximate location of the lesion, the border detection algorithm can focus on this region rather than the whole image. An accurate bounding box (the smallest axis-aligned rectangular box that encloses the lesion) might be useful for various reasons: (i) it provides an estimate of the lesion size (certain image segmentation algorithms such as region growing and morphological flooding can use the size of the region as a termination criterion), (ii) it might improve the border detection accuracy because the procedure is focused on a region that is guaranteed to contain the lesion, (iii) it speeds up the border detection because the procedure is performed on a region that is often smaller than the whole image, (iv) its surrounding might be utilized in the estimation of the background skin color, which is useful for various operations including the elimination of spurious regions that are discovered during the border detection procedure (10) and the extraction of dermoscopic features such as blotches (14) and blue-white areas (15).

In many dermoscopic images, the lesion can be roughly separated from the background skin using a grayscale thresholding method applied to the blue channel (8, 9). While there are a number of thresholding methods that perform well in general, the effectiveness of a method strongly depends on the statistical characteristics of the image (16). Figure 2 illustrates this phenomenon (the frame of this image is left intact for visualization purposes). Here, methods 2d, e, and g perform quite well. In contrast, methods 2c and h underestimate the optimal threshold, whereas method 2f overestimates the optimal threshold. Although method 2c is the most popular thresholding algorithm in the literature (17),

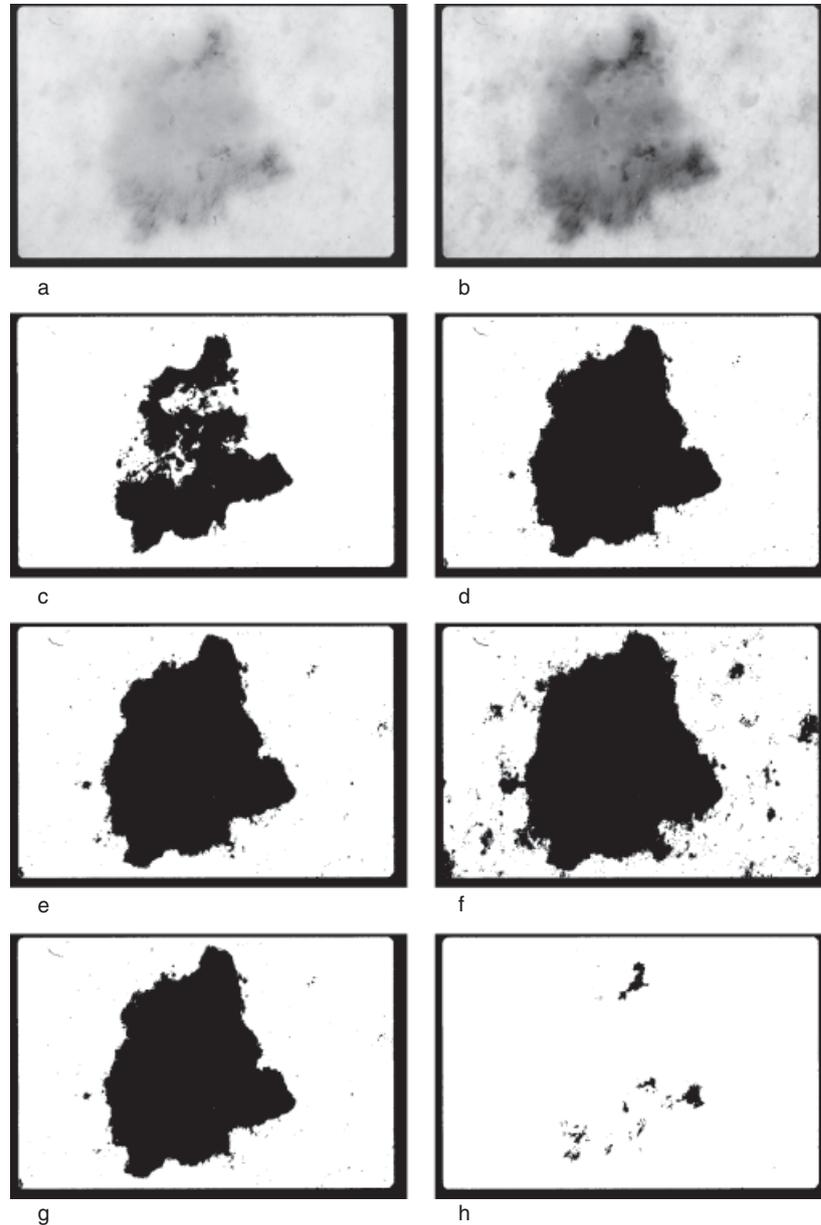


Fig. 2. Comparison of various thresholding methods: (a) Original image, (b) Blue channel, (c) Otsu's method (17) $T = 137$, (d) Kapur et al.'s method (18) $T = 178$, (e) Huang and Wang's method (19) $T = 183$, (f) Yen et al.'s method (20) $T = 200$, (g) Sahoo et al.'s method (21) $T = 179$, (h) Li and Tam's method (22) $T = 59$. ($T = \text{threshold}$).

for this particular image, it performs the second worst.

A possible approach to overcome this problem is to fuse the results provided by an ensemble of thresholding algorithms. In this way, it is possible to exploit the peculiarities of the participating thresholding algorithms synergistically, thus arriving at more robust final decisions than is possible with a single thresholding algorithm. We note that the goal of the fusion is not to outperform the individual thresholding algorithms, but to obtain accuracies comparable with that of the best thresholding algorithm independently of the image characteristics. In

this study, we used the threshold fusion method proposed by Melgani (16), which we describe briefly in the following.

Let $X = \{x_{mn}: m = 0, 1, \dots, M-1, n = 0, 1, \dots, N-1\}$ be the original scalar $M \times N$ image with L possible gray levels ($x_{mn} \in \{0, 1, \dots, L-1\}$) and $Y = \{y_{mn}: m = 0, 1, \dots, M-1, n = 0, 1, \dots, N-1\}$ be the binary output of the threshold fusion. Consider an ensemble of P thresholding algorithms. Let T_i and $A_i (i = 1, 2, \dots, P)$ be the threshold value and the output binary image associated with the i th algorithm of the ensemble, respectively. Within a Markov random field (MRF) framework the fusion problem can be formulated

as an energy minimization task. Accordingly, the local energy function U_{mn} to be minimized for the pixel (m, n) can be written as follows:

$$U_{mn} = \beta_{SP} U_{SP} [y_{mn}, Y^S(m, n)] + \sum_{i=1}^P \beta_i U_{II} [y_{mn}, A_i^S(m, n)] \quad (2)$$

where S is a predefined neighborhood system associated with pixel (m, n) , $U_{SP}(\cdot)$ and $U_{II}(\cdot)$ refer to the spatial and inter-image energy functions, respectively, whereas β_{SP} and β_i ($i = 1, 2, \dots, P$) represent the spatial and inter-image parameters, respectively. The spatial energy function can be expressed as

$$U_{SP} [y_{mn}, Y^S(m, n)] = - \sum_{y_{pq} \in Y^S(m, n)} I(y_{mn}, y_{pq}) \quad (3)$$

where $I(\cdot, \cdot)$ is the indicator function defined as

$$I(y_{mn}, y_{pq}) = \begin{cases} 1 & \text{if } y_{mn} = y_{pq} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

The inter-image energy function is defined as

$$U_{II} [y_{mn}, A^S(m, n)] = - \sum_{A_i(p, q) \in A_i^S(m, n)} \alpha^i(x_{pq}) I[y_{mn}, A_i(p, q)] \quad (5)$$

where $\alpha^i(\cdot)$ is a weight function given by

$$\alpha^i(x_{mn}) = 1 - \exp(-\gamma |x_{mn} - T_i|) \quad (6)$$

This function controls the effect of unreliable decisions at the pixel level that can be incurred by the thresholding algorithms. At the global (image) level decisions are weighed by the inter-image parameters β_i ($i = 1, 2, \dots, P$), which are computed as follows:

$$\beta_i = \exp(-\gamma |\bar{T} - T_i|) \quad (7)$$

where \bar{T} is the average threshold value

$$\bar{T} = \frac{1}{P} \sum_{i=1}^P T_i \quad (8)$$

The MRF fusion strategy proposed in (16) is as follows:

1. Apply each thresholding algorithm of the ensemble to the image X to generate the set of thresholded images A_i ($i = 1, 2, \dots, P$).
2. Initialize Y by minimizing for each pixel (m, n) the local energy function U_{mn} defined in Eq. (2) without the spatial energy term, i.e. by setting $\beta_{SP} = 0$.
3. Update Y by minimizing for each pixel (m, n) the local energy function U_{mn} defined in Eq. (2)

including the spatial energy term, i.e. by setting $\beta_{SP} \neq 0$.

4. Repeat step 3 K_{\max} times or until the number of different labels in Y computed over the last two iterations becomes very small.

In our preliminary experiments, we observed that, besides being computationally demanding, the iterative part (step 3) of the fusion algorithm makes only marginal contribution to the quality of the results. Therefore, in this study, we considered only the first two steps. The γ parameter was set to the recommended value of 0.1 (16). For computational reasons, α [Eq. (6)] and β [Eq. (7)] values were precalculated and the neighborhood system S was chosen as a 3×3 square.

The most important performance factor in the fusion algorithm seems to be the choice of the thresholding algorithms. We considered six popular thresholding algorithms to construct the ensemble: Otsu's (18), Kapur et al.'s (19), Huang and Wang's (20), Yen et al.'s (21), Sahoo et al.'s (22), and Li and Tam's (23) methods. In order to determine the best combination, we evaluated ensembles with three (20 ensembles), four (15 ensembles), five (six ensembles), and six (one ensemble) methods.

Figure 3 shows the output of four particular ensembles: Otsu–Kapur–Huang, Yen–Sahoo–Li, Otsu–Kapur–Huang–Yen, and Huang–Yen–Sahoo–Li. Note that each ensemble contains at least one method that either underestimates or overestimates the optimal threshold. It can be seen that each ensemble performs equally well, which demonstrates that failures in pathological cases might be prevented using a proper fusion strategy.

Figure 4a shows the result of the ensemble Otsu–Kapur–Huang–Sahoo. Here, the blue bounding box encloses the dermatologist determined border (see 'Results and Discussion'), whereas the red one encloses the binary output of the threshold fusion. It can be seen that the red box is completely contained within the blue box. This was observed in many cases because the automated thresholding methods tend to find the sharpest pigment change, whereas the dermatologists choose the outmost detectable pigment. We experimented with two different expansion methods to solve this problem. The first one involves expanding the automatic box by $P\%$ in four main directions. In other words, an automatic box of size $M_B \times N_B$ is expanded by $M_B \times P/100$ pixels in the west and east directions

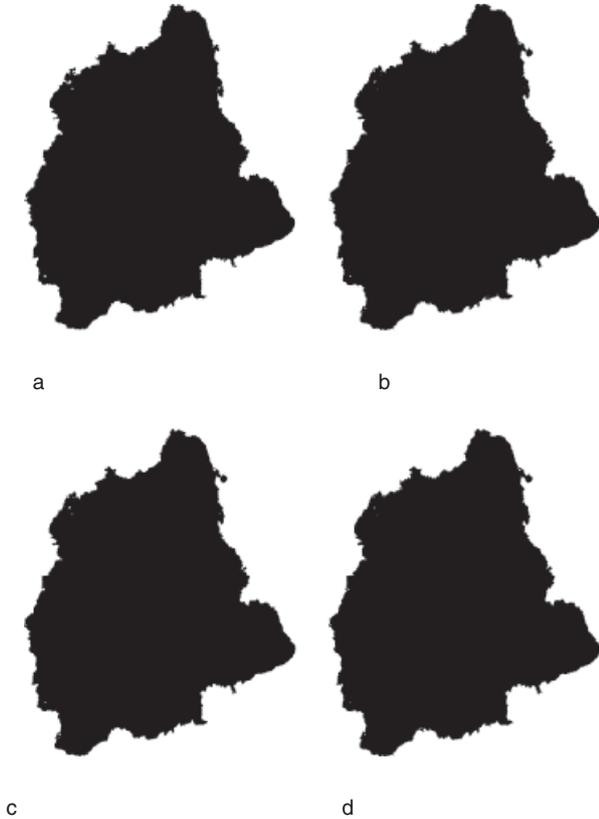


Fig. 3. Comparison of various threshold ensembles: (a) Otsu-Kapur-Huang, (b) Yen-Sahoo-Li, (c) Otsu-Sahoo-Li, (c) Otsu-Kapur-Huang-Yen, (d) Huang-Yen-Sahoo-Li.

and $N_B \times P/100$ pixels in the north and south directions. The second one involves incrementing the threshold values obtained by each algorithm in the ensemble by G gray levels. In the rest of this article, we will refer to these expansion methods as non-adaptive and adaptive, respectively. Figure 4a and b shows the results of these methods with the expanded box shown in green. In this particular example, the non-adaptive method performs better in bringing the automatic box closer to the manual one. In order to determine the optimal expansion amounts we evaluated $P \in \{2,4,6,8\}$ and $G \in \{4,6,8,10\}$.

Results and Discussion

The proposed method was tested on a set of 428 dermoscopy images obtained from the EDRA Interactive Atlas of Dermoscopy (2) and the Keio University Hospital. These were 24-bit RGB color images with dimensions ranging from 771×507 to 768×512 pixels. An experienced dermatologist (W.V.S.) determined the manual borders by selecting a number of points on the lesion border, which were then connected

by a second-order B-spline. The bounding box error was quantified using the grading system developed by Hance et al. (24)

$$\varepsilon = \frac{\text{Area}(\text{AutomaticBox} \oplus \text{ManualBox})}{\text{Area}(\text{ManualBox})} \times 100 \quad (9)$$

where *AutomaticBox* is the binary image obtained by filling the bounding box of the fusion output, *ManualBox* is the binary image obtained by filling the bounding box of the dermatologist-determined border, \oplus is the exclusive OR operation, which essentially determines the pixels for which the *AutomaticBox* and *ManualBox* disagree, and $\text{Area}(I)$ denotes the number of pixels in the binary image I .

We determined the optimal parameter combination for the presented approximate bounding box computation method as follows. First, the black frame removal procedure described in ‘Black frame removal’ is performed on each image in the data set. The lesion bounding box is then computed using the fusion method described in ‘Approximate lesion localization’ with one of the 42 ensembles. Finally, the approximate bounding box is expanded using either the non-adaptive method with $P \in \{2,4,6,8\}$ or the adaptive method with $G \in \{4,6,8,10\}$. Table 1 shows various statistics associated with the four most accurate ensembles for each expansion method. The last two columns refer to the mean and standard deviation values, respectively, for the percentage image size reduction, i.e. $\frac{\text{Area}(\text{AutomaticBox})}{MN} 100$, provided by the bounding box computation. The following observations are in order: (i) both expansion methods reduce the mean bounding box error, (ii) the lowest mean errors were obtained using the ensemble Otsu-Kapur-Huang-Sahoo, (iii) the non-adaptive expansion method was more effective than the adaptive one, (iv) the computation of the bounding box reduced the original image size by about 260%.

The adaptive method was less effective than the non-adaptive one probably because the former often expands the approximate box by unpredictable amounts: either too little (as in Fig. 4b) or too much depending on the shape of the histogram and the value of the G parameter. In contrast, the latter always expands the approximate box by an amount specified by the P parameter.

Table 2 shows the statistics for the individual thresholding methods. Note that, due to space limitations, we report only the results of the non-

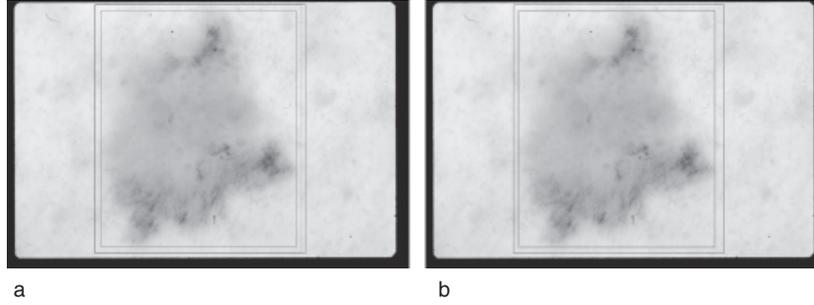


Fig. 4. Comparison of the bounding box expansion methods: (a) Non-adaptive expansion with $P = 3\%$, (b) Adaptive expansion with $G = 6$.

TABLE 1. Ensemble statistics

Ensembles	Expansion method	μ_{ε_i}	σ_{ε_i}	μ_{ε_x}	σ_{ε_x}	μ_s	σ_s
Otsu–Kapur–Huang–Sahoo	Non-adaptive ($P = 2$)	10.25	8.10	7.58	8.13	268.31	185.64
Otsu–Huang–Yen–Li	Non-adaptive ($P = 4$)	11.92	7.59	7.89	6.30	260.55	183.85
Otsu–Huang–Sahoo–Li	Non-adaptive ($P = 4$)	11.98	7.62	7.90	6.20	260.95	184.14
Otsu–Huang–Sahoo	Non-adaptive ($P = 2$)	11.14	7.17	7.91	6.71	273.84	195.69
Otsu–Kapur–Huang–Sahoo	Adaptive ($G = 6$)	10.25	8.10	9.27	7.68	276.92	192.14
Kapur–Huang–Sahoo–Li	Adaptive ($G = 8$)	10.98	7.66	9.43	7.69	279.03	194.42
Otsu–Kapur–Huang–Sahoo	Adaptive ($G = 4$)	10.25	8.10	9.44	7.56	279.98	194.26
Kapur–Huang–Sahoo–Li	Adaptive ($G = 6$)	10.98	7.66	9.67	7.58	282.09	196.58

μ , mean; σ , standard deviation; ε_i , initial box error; ε_x , expanded box error.

TABLE 2. Individual statistics

Thresholding method	Expansion method	μ_{ε_i}	σ_{ε_i}	μ_{ε_x}	σ_{ε_i}	μ_s	σ_s
Otsu	Non-adaptive ($P = 2$)	12.05	9.10	9.00	8.95	275.07	199.28
Kapur	Non-adaptive ($P = 2$)	12.87	16.86	12.68	17.56	261.95	197.94
Huang	Non-adaptive ($P = 2$)	20.31	67.97	17.17	69.76	269.59	190.09
Yen	Non-adaptive ($P = 2$)	14.98	27.12	15.74	27.74	255.61	250.53
Sahoo	Non-adaptive ($P = 2$)	13.43	24.60	13.37	25.19	254.43	184.36
Li	Non-adaptive ($P = 2$)	15.12	9.65	11.06	9.07	293.54	215.80
Otsu	Non-adaptive ($P = 4$)	12.05	9.10	9.10	9.14	256.86	182.82
Kapur	Non-adaptive ($P = 4$)	12.87	16.86	15.54	18.61	245.36	183.78
Huang	Non-adaptive ($P = 4$)	20.31	67.97	16.83	70.69	251.99	174.44
Yen	Non-adaptive ($P = 4$)	14.98	27.12	19.32	28.49	239.46	230.91
Sahoo	Non-adaptive ($P = 4$)	13.43	24.60	16.43	25.98	238.32	170.38
Li	Non-adaptive ($P = 4$)	15.12	9.65	9.41	7.99	273.93	198.61

μ , mean; σ , standard deviation; ε_i , initial box error; ε_x , expanded box error.

adaptive expansion method (as in the ensemble case, the adaptive method has inferior performance). It can be seen that, in most configurations, the individual methods obtain significantly higher mean errors than the best ensemble methods, i.e. the first four rows of Table 1. This is because, as explained in ‘Approximate lesion localization,’ the individual methods are more prone to catastrophic failures when given pathological input images. The high standard deviation values also support this explanation. Only the performance of Otsu (with $P = 2, 4$) and Li

et al.’s (with $P = 4$) methods is close to the performance of the ensembles. However, as mentioned in ‘Approximate lesion localization,’ the goal of fusion is not to outperform the individual thresholding algorithms, but to obtain accuracies comparable with that of the best thresholding algorithm independently of the image characteristics.

As mentioned in ‘Approximate lesion localization,’ an accurate bounding box can provide an estimate of the lesion size, i.e. Area (*Manual Border*). In order to verify this, we calculated the best

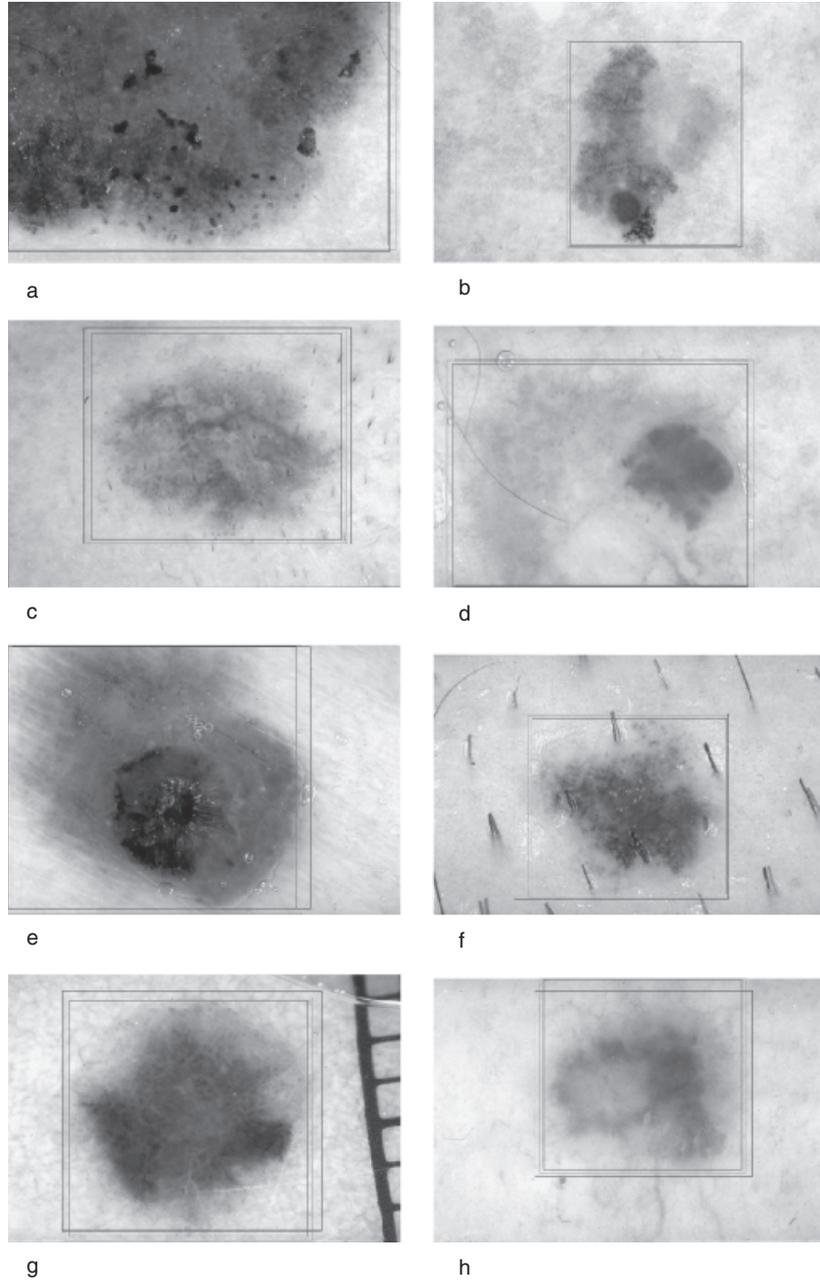


Fig. 5. Sample bounding box computation results: (a) $\epsilon_i = 3.26\%$, $\epsilon_x = 1.83\%$, (b) $\epsilon_i = 4.89\%$, $\epsilon_x = 3.62\%$, (c) $\epsilon_i = 10.87\%$, $\epsilon_x = 4.03\%$, (d) $\epsilon_i = 3.18\%$, $\epsilon_x = 4.82\%$, (e) $\epsilon_i = 6.75\%$, $\epsilon_x = 5.36\%$, (f) $\epsilon_i = 14.24\%$, $\epsilon_x = 9.21\%$, (g) $\epsilon_i = 14.05\%$, $\epsilon_x = 10.91\%$, (h) $\epsilon_i = 18.36\%$, $\epsilon_x = 13.62\%$ (ϵ_i : initial box error, ϵ_x : expanded box error).

fitting line for $Area(AutomaticBox)$ vs. $Area(ManualBorder)$ using the generalized least-squares method (25):

$$Area(ManualBorder) \approx Area(AutomaticBox) \times 0.861271 - 15627.226419 \quad (10)$$

where $ManualBorder$ is the binary image obtained by filling the dermatologist-determined border. The accuracy of this relation was calculated by plugging the area of the approximate bounding

box for each image into Eq. (10), and then comparing the result with the actual area of the lesion, which is calculated from the dermatologist-determined border. The percentage mean and standard deviation errors over the entire image set were 11.88 and 11.49, respectively. These results demonstrate that the lesion size can be estimated from the bounding box area with relatively high accuracy. An even better estimate can be made from the binary output of the threshold fusion. The best fitting line for $Area(FusionOutput)$ vs. $Area(ManualBorder)$ was

calculated as

$$\begin{aligned} \text{Area}(\text{ManualBorder}) &\approx \text{Area}(\text{FusionOutput}) \\ &\times 1.158209 - 4485.287871 \end{aligned} \quad (11)$$

where *FusionOutput* is the binary output of the threshold fusion. The percentage mean and standard deviation errors for this relation were 8.16 and 8.54, respectively.

Figure 5 shows sample bounding box computation results obtained using the ensemble Otsu–Kapur–Huang–Sahoo with $P = 2$. It can be seen that the presented method determines an accurate bounding box even for lesions with fuzzy borders. We note that while the expansion operation is useful in most cases, in some cases such as Fig. 5d, it might deteriorate the results slightly.

Conclusions

In this paper, an automated method for approximate lesion localization in dermoscopy images is presented. The method is comprised of three main phases: black frame removal, initial bounding box computation using an ensemble of thresholding algorithms, and expansion of the initial bounding box. The execution time of the method is about 0.15s for a typical image of size 768×512 pixels on an Intel Pentium D 2.66 GHz computer.

The presented method may not perform well on images with significant amount of hair or bubbles because these elements alter the histogram, which in turn results in biased threshold computations. For images with hair, a preprocessor such as DullRazor™ (11) might be helpful. Unfortunately, the development of a reliable bubble removal method remains an open problem.

Future work will be directed toward testing the utility of the presented method in a border detection study. The implementation of the threshold fusion method will be made publicly available as part of the Fourier image processing and analysis library, which can be downloaded from <http://sourceforge.net/projects/fourier-ipal>

Acknowledgments

This work was supported by grants from Louisiana Board of Regents (LEQSF2008-11-RD-A-12) and NIH (SBIR #2R44 CA-101639-02A2).

References

1. Jemal A, Siegel R, Ward E et al. Cancer statistics, 2008. *CA Cancer J Clin* 2008; 58: 71–96.
2. Argenziano G, Soyer HP, De Giorgi V et al. *Dermoscopy: a tutorial*. Milan, Italy: EDRA Medical Publishing & New Media, 2002.
3. Steiner K, Binder M, Schemper M et al. Statistical evaluation of epiluminescence dermoscopy criteria for melanocytic pigmented lesions. *J Am Acad Dermatol* 1993; 29: 581–588.
4. Binder M, Schwarz M, Winkler A et al. Epiluminescence microscopy. A useful tool for the diagnosis of pigmented skin lesions for formally trained dermatologists. *Arch Dermatol* 1995; 131: 286–291.
5. Fleming MG, Steger C, Zhang J et al. Techniques for a structural analysis of dermatoscopic imagery. *Comput Med Imaging Graph* 1998; 22: 375–389.
6. Schmid P. Segmentation of digitized dermatoscopic images by two-dimensional color clustering. *IEEE Trans Med Imaging* 1999; 18: 164–171.
7. Erkol B, Moss RH, Stanley RJ, Stoecker WV, Hvatum E. Automatic lesion boundary detection in dermoscopy images using gradient vector flow snakes. *Skin Res Technol* 2005; 11: 17–26.
8. Iyatomi H, Oka H, Saito M et al. Quantitative assessment of tumor extraction from dermoscopy images and evaluation of computer-based extraction methods for automatic melanoma diagnostic system. *Melanoma Res* 2006; 16: 183–190.
9. Celebi ME, Aslandogan YA, Stoecker WV et al. Unsupervised border detection in dermoscopy images. *Skin Res Technol* 2007; 13: 454–462.
10. Celebi ME, Kingravi HA, Iyatomi H et al. Border detection in dermoscopy images using statistical region merging. *Skin Res Technol* 2008; 14: 347–353.
11. Lee TK, Ng V, Gallagher R et al. Dullrazor: a software approach to hair removal from images. *Comput Biol Med* 1997; 27: 533–543.
12. Zhou H, Chen M, Gass R et al. Feature-preserving artifact removal from dermoscopy images. *Proc SPIE Medical Imaging 2008 Conf.*, 6914: 1B1–1B9.
13. Levkowitz H, Herman GT. GLHS: A generalized lightness, hue, and saturation color model. *CVGIP: Graphic Models Image Process* 1993; 55: 271–285.
14. Stoecker WV, Gupta K, Stanley RJ et al. Detection of asymmetric blotches in dermoscopy images of malignant melanoma using relative color. *Skin Res Technol* 2005; 11: 179–184.
15. Celebi ME, Iyatomi H, Stoecker WV et al. Automatic detection of blue-white veil and related structures in dermoscopy images. *Comput Med Imaging Graph* 2008; 32: 670–677.
16. Melgani F. Robust image binarization with ensembles of thresholding algorithms. *J Electronic Imaging* 2006; 15: 023010.
17. Sezgin M, Sankur B. Survey over image thresholding techniques and quantitative performance evaluation. *J Electronic Imaging* 2004; 13: 146–165.
18. Otsu N. A threshold selection method from gray level histograms. *IEEE Trans Syst Man Cybern* 1979; 9: 62–66.
19. Kapur JN, Sahoo PK, Wong AKC. A new method for gray-level picture thresholding using the entropy of the histogram. *Graphic Models Image Process* 1985; 29: 273–285.
20. Huang L-K, Wang M-JJ. Image thresholding by minimizing the measures of fuzziness. *Pattern Recognition* 1995; 28: 41–51.

21. Yen JC, Chang FJ, Chang S. A new criterion for automatic multilevel thresholding. *IEEE Trans Image Process* 1995; 4: 370–378.
22. Sahoo PK, Wilkins C, Yeager J. Threshold selection using renyi's entropy. *Pattern Recognition* 1997; 30: 71–84.
23. Li CH, Tam PKS. An iterative algorithm for minimum cross entropy thresholding. *Pattern Recognition Lett* 1998; 18: 771–776.
24. Hance GA, Umbaugh SE, Moss RH, Stoecker WV. Un-supervised color image segmentation with application to skin tumor borders. *IEEE Eng Med Biol* 1996; 15: 104–111.
25. Li HC. A generalized problem of least squares. *Am Math Month* 1984; 91: 135–137.

Address:
M. Emre Celebi
Department of Computer Science
Louisiana State University
Technology Center 206
One University Place
Shreveport, LA 71115
USA
e-mail: ecelebi@lsus.edu