Automatic Skin Lesion Segmentation via Iterative Stochastic Region Merging

Alexander Wong, *Member, IEEE*, Jacob Scharcanski, *Senior Member, IEEE*, and Paul Fieguth, *Member, IEEE*

***Abstract*—An automatic method for segmenting skin lesions in conventional macroscopic images is presented. The images are acquired with conventional cameras, without the use of a dermoscope. Automatic segmentation of skin lesions from macroscopic images is a very challenging problem due to factors such as illumination variations, irregular structural and color variations, the presence of hair, as well as the occurrence of multiple unhealthy skin regions. To address these factors, a novel iterative stochastic region-merging approach is employed to segment the regions corresponding to skin lesions from the macroscopic images, where stochastic region merging is initialized first on a pixel level, and subsequently on a region level until convergence. A region merging likelihood function based on the regional statistics is introduced to determine the merger of regions in a stochastic manner. Experimental results show that the proposed system achieves overall segmentation error of under 10% for skin lesions in macroscopic images, which is lower than that achieved by existing methods.**

***Index Terms*—lesion, skin cancer, stochastic, region merging, iterative.**

I. INTRODUCTION

According to the Skin Cancer Foundation, skin cancer is the most frequent form of cancer in the United States [1]. Of greater concern is the fact that there is a continuous increase in skin cancer incidence over time in many areas of the world such as Singapore [2], Slovakia [3], and the Netherlands [4], with no signs of leveling off [4].

To date, the most effective way to reduce the mortality related to skin cancer is through early skin examination to identify possible signs of skin discoloration and abnormalities. The early detection of such abnormalities is particularly important, since these abnormalities may evolve to skin cancer [5]. In skin examinations it is important to differentiate malignant from nonmalignant lesions. This enables malignant melanoma skin cancer to be diagnosed at an earlier and much more treatable stage, in turn helping to avoid unnecessary surgery and improving patient morbidity and mortality. Unfortunately, the individual skin examination of a large population (where each individual may be a potential patient) can be expensive for most health care systems, given the limited number of specialists trained for performing proper skin diagnosis, and their location in specific health care providers at larger centers

Alexander Wong and Paul Fieguth are with Dept. of Systems Design Engineering, U. of Waterloo, Waterloo, Canada N2L 3G1. {a28wong,pfieguth}@uwaterloo.ca Jacob Scharcanski is with Instituto de Informatica and Programa de Pós-Graducacão em Engenharia Eletrica, Universidade Federal do Rio Grande do Sul, Caixa Postal 15064, 91501-970, Porto Alegre, RS, Brasil. He is also is an Adjunct Associate Professor with Systems Design Engineering, U. of Waterloo, Waterloo, Canada N2L 3G1. jacobs@inf.ufrgs.br

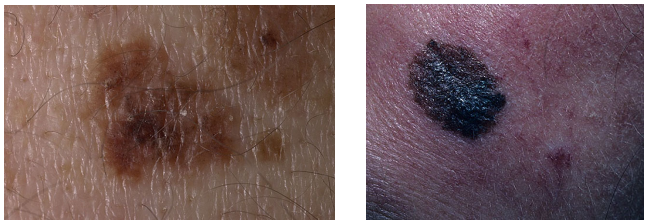

Fig. 1. Examples of macroscopic images of skin cancer regions. Challenges in segmenting cancer regions in an automated manner include illumination variations, irregular structural and color variations, the presence of hair, as well as the formation of multiple cancerous regions.

to optimize the associated infrastructure costs (i.e., hospitals and clinics). Usually, a dermatologist performs a thorough patient skin examination using a dermoscope[1].

A particularly promising development that has the potential for improving skin cancer diagnosis is the remote clinical diagnosis in telemedicine, where macroscopic images of the suspected skin area are transferred to the clinician over the Internet for remote examination. In this way, suspicious skin lesions are pre-screened using images acquired with conventional cameras, and if the image analysis suggest a lesion that needs special attention, the patient is immediately referred to a dermatologist [6]. Therefore, allowing for patients who would otherwise have no access to skilled dermatologists to receive proper clinical diagnosis for suspicious skin discoloration and abnormalities. Given the potential for this enabling technology, there has been great interest in developing computer-aided systems to assist in the rapid clinical analysis and diagnosis of dermatological skin lesions (e.g. melanomas) [7], [8] Such computer-aided systems have now been tested remotely using the Internet and has been shown to be effective and have great potential for widespread use. Motivated by the potential to allow for remote clinical diagnosis in telemedicine, where macroscopic images are used for initial diagnosis instead of dermoscopy images, our focus is on computer-aided clinical diagnosis of skin lesions from conventional macroscopic images as opposed to dermoscopy images.

An essential step in the computer-aided clinical diagnosis of skin lesions is the automatic segmentation of the lesions from macroscopic images. The segmentation is very challenging due

[1]A dermoscope, or dermatoscope, is a tool widely by dermatologists to help with examining the skin. Usually, a dermoscope is a hand held optical device, much like a magnifying glass, with a light source attached (generally specific for dermatological use): (a) They help to differentiate melanocytic lesions from nonmelanocytic and malignant lesions from nonmalignant; (b) Dermoscopy can be used to justify a case for surgical excision; and (c) The examination with a dermoscope helps to provide a careful examination of the patient's skin, and can reassure the patient that he/she has received a thorough skin examination.

to factors such as illumination variations, irregular structural and color variations, the presence of hair, as well as the existence of multiple lesions in the skin, some of which can be seen in Fig. 1.

Several methods have been proposed to tackle this important task. A popular class of approaches to skin lesion segmentation is thresholding [8], [9], [10], [11], [12], which has been shown to be effective for situations where the lesions have consistent characteristics and the surrounding skin regions are homogeneous in nature, as well as well handling situations characterized by multiple regions. However, such approaches face difficulties in situations characterized by structural, illumination, and color variations, where no clear threshold can be found that separates the lesion regions from the surrounding skin regions, and resulting in poor segmentation accuracy. Another popular class of approaches is based on active contours [25], [14], [15], where a curve is evolved towards the boundaries of the skin lesion regions. While more robust than thresholding at handling noise, artifacts, and variations in illumination and color, active contour approaches can face difficulties in situations characterized by weak contrast between the skin lesions and the surrounding skin regions.

A new class of approaches to skin lesion segmentation that has been recently investigated are methods based on Markov Random Fields (MRF) [20] and other statistics [16], [17], which take into consideration local spatial interactions and densities. Of particular interest is the employment of the statistical region merging algorithm [18] for skin lesion segmentation as proposed by Celebi et al. [17], which has been shown to work well with dermoscopy images. While computationally efficient and effective for dermoscopy images since the dermoscope is focused on a small area (typically containing a single lesion under ideal lighting conditions), such an approach has difficulties dealing with macroscopic images characterized by noise and artifacts, structural, illumination and color variations, multiple lesions, and weak boundary separation. This leads to oversegmentation and undersegmentation results under the degradations commonly found in macroscopic images. A survey on skin lesion segmentation can be found in [19].

The research goal of this paper is to address the aforementioned issues to obtain accurate skin lesion segmentation performance, specifically from macroscopic images in an automated manner. The proposed iterative stochastic region merging method is robust to noise and artifacts, multiple lesion regions, structural variations, illumination and color variations, and weak boundary separation between skin lesions and surrounding skin, which are the key challenges to skin lesion segmentation in macroscopic images. To the best of the authors' knowledge, such an iterative stochastic region merging approach to skin lesion segmentation has not been previously proposed. The proposed method differs greatly from current methods based on the statistical region merging algorithm [18] such as [17]. First, while current methods make use of the merging criteria proposed in [18] to make a deterministic merging decision, the proposed merging decision extends upon this by introducing a new merging likelihood function that allows for a stochastic merging decision. More importantly, unlike current methods that attempt to obtain the segmentation results within a single pass, the proposed method introduces a new multi-pass strategy that continually refines the segmentation results to improve segmentation accuracy. Such differences play an important part in the improved segmentation performance of the proposed method for skin lesion segmentation, which will be illustrated in Section III-A.

This paper is organized as follows. First, the proposed method is described in Section II. Second, the experimental results using real macroscopic images of various types of dermatological lesions are presented in Section III. Finally, conclusions are drawn and future work is discussed in Section IV.

II. Method

The proposed iterative stochastic region merging method for skin lesion segmentation can be summarized as follows. Each pixel in the image is assigned to a unique region, and these regions are subsequently merged with other regions in a stochastic manner, based on a region merging likelihood function. The regions formed during this initial phase are then merged using the same stochastic region merging process to yield the final segmentation results. The theory behind the proposed automatic segmentation system will be described in the following sections.

A. Problem Formulation

Let S be the discrete lattice upon which the image of the skin lesions is defined, and $s \in S$ be a site in the lattice, which in the case of the image refers to a pixel location. Let R_s be a random variable taking on a region label $\{1, \ldots, k\}$, denoting a unique lesion region or the surrounding skin region, and F_s be a random variable taking on the color associated with s. To keep the method illumination and color invariant, there is no prior model or prior region associated with healthy skin.

Given an image f, let $f = \{f_s | s \in S\}$ be the observed color image, and $r = \{r_s | s \in S\}$ an instantiation of the region label field. The problem of skin lesion segmentation can therefore be treated as an inverse problem, where we wish to determine r given f, formulated as the following Maximum a posteriori (MAP) estimation problem:

$$\hat{r} = \arg\max_{r} \{P(r\,|f)\}, \tag{1}$$

where $P(r|f)$ is the posterior. Based on Bayes' theorem, the problem posed in (1) is equivalent to

$$\hat{r} = \arg\max_{r} \{P(f\,|r)P(r)\}, \tag{2}$$

where $P(f|r)$ is the color-to-label assignment likelihood, $P(r)$ is the region label prior, and $P(f) = 1$. The solution to the problem defined in (2) is difficult to solve directly due to complex imaging challenges, such as noise and artifacts, multiple lesion regions, structural variations, illumination and color variations, and weak boundary separation between unhealthy skin regions and surrounding skin regions. These are the key challenges to skin lesion segmentation in macroscopic images.

B. First Stage - Initialization

To better handle illumination and color variations, as well as noise and artifacts, one effective approach is to employ a Markov Random Field (MRF) model [20] for the prior $P(r)$ in (2), where the image is modeled in terms of local spatial interactions of its sites. More specifically, the probability distribution at a site s is dependent only on its neighboring sites. This MRF spatial context model is realized in the first stage of the proposed algorithm in the following manner.

Given a $M \times N$ image f, since there are no initial knowledge or assumptions about the image (e.g., no knowledge about the number of regions), we first assume that each site s represents a unique region (i.e. a pixel), and we assign each site s a unique region label $R = \{1, \ldots, M \times N\}$. An arbitrary site labeling order can be used.

We then construct a region adjacency graph [20], where each vertex represents a region R, and the graph edges E connect each vertex (representing an individual region R) to its adjacent graph vertices (i.e., to the regions adjacent to R on the discrete lattice S) (see Fig. 2). At this stage, we set up the region adjacency graph such that:

- each site s represents a unique region and is represented by a single vertex,
- each site s has eight neighbors on S that it is directly connected to, and so each vertex has eight adjacent vertices on the graph.

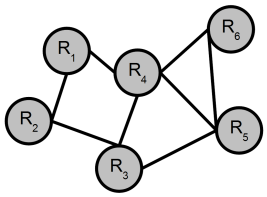

Fig. 2. An example of a region adjacency graph, where each vertex represents a region R, and the graph edges E connect each vertex (representing an individual region R) to its adjacent graph vertices (i.e., to the regions adjacent to R on the discrete lattice S).

C. Second Stage - Stochastic Region Merging

Based on this MRF model, we employ a stochastic region merging strategy to solve the problem posed in (2), where a region R_a is merged with an adjacent region R_b, thus taking on the region label R_b, with a merging probability of $\alpha(R_a, R_b)$. A thorough explanation of the MAP problem being solved via region merging and an MRF model can be found in [21]. To compute the merging probability $\alpha(R_a, R_b)$, we introduce a new stochastic region merging likelihood function, which accounts for $P(f|r)$ in (2). The proposed stochastic region merging likelihood function extends upon the statistical region merging theory proposed by Nock and Nielsen [18],

$$\alpha(R_a, R_b) = \exp\left[-\frac{\left(E\left[R_a\right] - E\left[R_b\right]\right)^2}{\Lambda\left(R_a, R_b\right)}\right], \qquad (3)$$

where $E[.]$ is the expected value of elements in the region, and Λ is a statistical region merging penalty function defined as :

$$\Lambda\left(R_a, R_b\right) = \frac{{D_f}^2}{2Q}\left[\frac{\ln\left(\Phi(f)^2\right)}{\Phi(R_a)} + \frac{\ln\left(\Phi(f)^2\right)}{\Phi(R_b)}\right], \qquad (4)$$

where $\Phi(.)$ represents the number of elements contained in the set provided as the argument (e.g., $\Phi(f)$ is the number of pixels in the image f), D_f represents the dynamic range of f (e.g., 256 for an 8-bit image), and Q is a regularization term. In the situation where we are dealing with color images, Eq. (3) is computed for each color channel, and the final merging likelihood is the product of the individual likelihoods of each channel. Based on the computed merging probability $\alpha(R_a, R_b)$, we decide whether R_a is merged with R_b in the following probabilistic manner. We generate a random number u from a uniform distribution between 0 and 1. If the random number u is less than $\alpha(R_a, R_b)$, then R_a is merged with R_b. Otherwise, the regions are not merged.

Several observations can be made about the proposed stochastic region merging likelihood function. First, the merging probability α increases as the size of regions decrease. Therefore, small regions are more likely to merge than large regions, which is useful for the handling of noise and artifacts, which are typically characterized as small regions and should merge with larger regions to avoid oversegmentation. Second, α decreases as the difference in regional expectations $(E\left[R_a\right] - E\left[R_b\right])$ increases, thus reducing the probability that regions with very different characteristics from merging and hence promoting intra-region homogeneity. Third, α decreases as Q increases, thus increasing the probability of merger when Q is low and decreasing the probability of merger when Q is high. What this effectively does is allow for more flexible control over the segmentation scale (e.g., increasing Q promotes oversegmentation). Unlike the original deterministic statistical region merging theory proposed by Nock and Nielsen [18], the introduction of the stochastic aspect into the region merging process allows for increased segmentation robustness that is less prone to being trapped in local optima. For example, in the deterministic approach, region merging occurs only if the regional expectation difference (i.e. $E\left[R_a\right] - E\left[R_b\right]$) is less than a hard statistical threshold. Therefore, even if the regional expectation difference is greater than the threshold by even a very minute quantity due to factors such as noise and artifacts, structural, illumination and color variations, and weak boundary separation, region merging will not occur despite their relative similarity. On the other hand, in the proposed method, the merging probability of such a case would be relatively high and so region merging will likely occur. We will show in the experimental results an illustrative comparison between the deterministic approach and the proposed stochastic approach to justify our proposed approach.

To determine the order of merger evaluation, all adjacent region pairs are placed in a priority queue based on ascending regional expectation differences $(E\left[R_a\right] - E\left[R_b\right])$ (lower regional expectation differences have higher merger evaluation priority). As the stochastic region merging progresses, the region adjacency graph is updated to reflect the region merging decisions at each step.

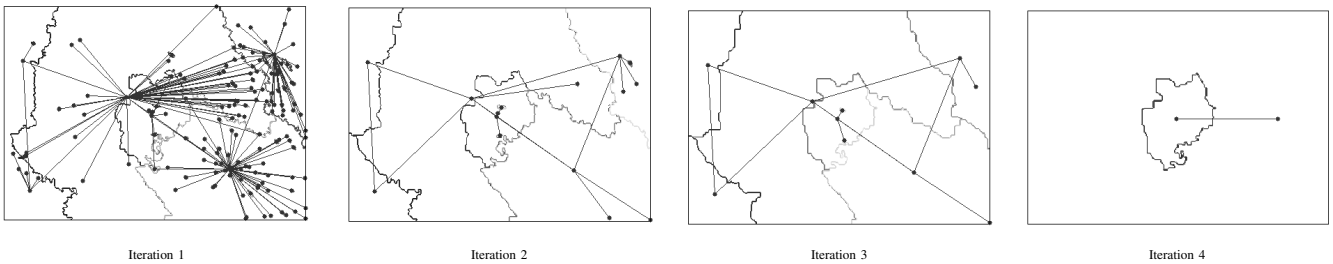

Fig. 3. An illustration of the iteration process for an example image. The nodes and edges of the region adjacency graph are overlayed on the segmentation results at each iteration. In these cases, each region is considered a site in the lattice. Each site (i.e., graph vertex) is depicted in the centroid location of each spatial region. At each iteration, the set of regions determined in the previous iteration are refined using a stochastic region merging process. This is repeated until convergence.

Algorithm 1 Iterative Stochastic Region Merging

1: Assign each site s in the image f a unique region label (Note: order of assignment does not matter).
2: Construct an initial region adjacency graph, where each vertex represents a site s with eight adjacent vertices each (i.e., its eight neighbors).
3: **repeat**
4: Place all adjacent region pairs into a priority queue based on ascending regional expectation differences.
5: **repeat**
6: Remove region pair R_a and R_b from priority queue.
7: Merge region pair R_a and R_b with a probability $\alpha(R_a, R_b)$ based on the proposed region merging likelihood function in (3).
8: If merging occurs, update the adjacency graph.
9: **until** Priority queue is empty
10: Decrease Q by half.
11: **until** Convergence

This process is continued until all adjacent region pairs in the priority queue have been evaluated to yield an initial segmented result, thus constituting one iteration in the proposed method.

D. Stage Three - Iterative Refinement

After the priority queue has been emptied, the adjacent region pairs from the initial segmentation results are inserted into the priority queue in ascending regional expectation difference order, and the stochastic region merging process is repeated, with Q decreased by half for each iteration, to yield refined segmented results. The number of regions is decreased at each step, and this iterative process is continued until convergence (no further changes to the region adjacency graph, or only two regions left), to yield the final set of segmentation results with k different regions, each assigned a unique region label $\{1, \ldots, k\}$. An example of the iterative process is shown in Fig. 3. The largest region, measured by area in pixels, in the segmentation results is assumed to correspond to the surrounding skin region, while the remaining regions correspond to the skin lesion. The pseudo-code for the proposed iterative stochastic region merging algorithm is presented in Algorithm 1.

This proposed iterative process departs significantly from the original deterministic statistical region merging theory proposed by Nock and Nielsen [18], which is not iterative and has a tendency to oversegment with high values of Q, and undersegment with low values of Q. Given the iterative nature of the proposed algorithm, Q can be set to an arbitrarily high initial value (e.g., $Q = 100$ in our tests, with $Q \geq 5$ for all iterations to avoid undersegmentation based on our experiments).

III. Experimental Results and Discussion

To evaluate the effectiveness of the proposed method, segmentation was performed on 60 test macroscopic images containing 40 malignant cases (20 melanoma cases and 20 carcinoma cases), and 20 benign cases (20 naevi cases). This large variety of dermatological lesions allows us to better study the practical performance of the proposed method. These test macroscopic images range in size from 720×480 pixels to 473×720 pixels, and are representative of real-world scenarios, where the images are characterized by noise and artifacts, structural, illumination, and color variations, as well as unclear separation between unhealthy skin regions and surrounding skin regions, making them noticeably more challenging to segment than dermascopy images. The proposed method was implemented in MATLAB, and tested on an Intel Pentium 4.3 GHz machine with 1 GB of RAM, with Q set to an initial value of 100.

A. Comparison with Deterministic Statistical Region Merging

Given that the proposed method extends upon the original deterministic statistical region merging algorithm proposed by Nock and Nielsen [18] that integrates a MRF model [20], and other methods have utilized this approach for skin lesion segmentation [17], it is worthwhile to provide a comparison between the proposed method and the deterministic algorithm. The computational complexity of the proposed algorithm with respect to the original deterministic statistical region merging algorithm proposed by Nock and Nielsen [18] is difficult to determine from an analytical perspective, given that it depends on the complexity of the underlying image and on the amount of segments proposed at each iteration. However, to put things into perspective, the deterministic statistical region merging algorithm proposed by Nock and Nielsen [18] takes approximately 2 seconds to segment one image, while the

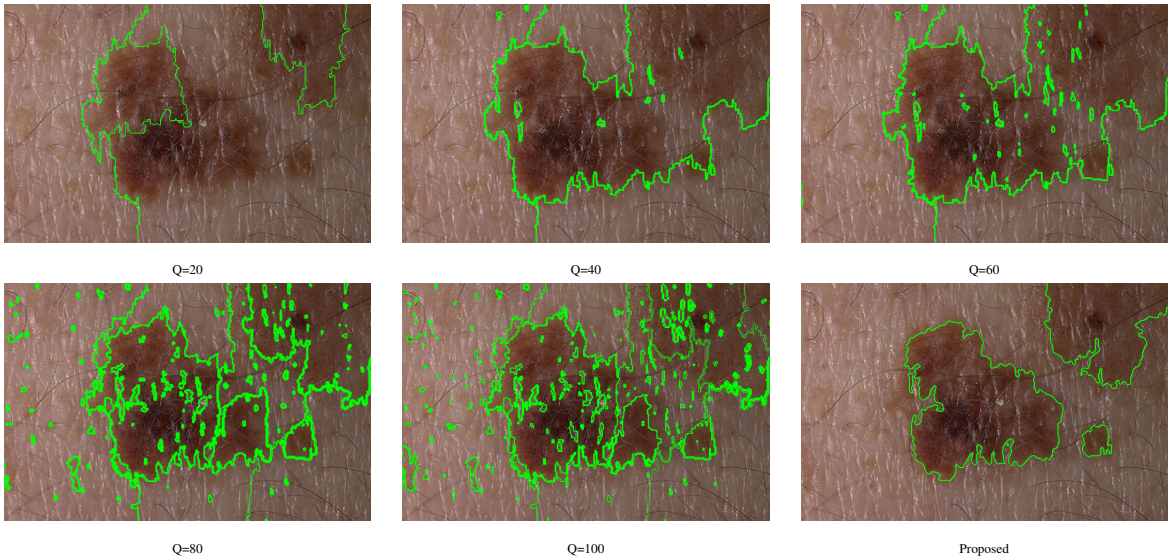

Fig. 4. Segmentation results using the original deterministic statistical region merging approach [18] with various values of Q, and the results of the proposed stochastic region merging approach. It can be observed that, regardless of the value of Q selected, the deterministic approach is unable to achieve correct segmentation results, while the proposed stochastic approach is able to give accurate segmentation results.

proposed method takes approximately 6 seconds to segment one image. However, in terms of segmentation accuracy, the original deterministic statistical region merging algorithm will give inaccurate results in situations characterized by noise and artifacts, structural, illumination and color variations, multiple lesions, and weak boundary separation, while the proposed algorithm is robust to these conditions and will provide significantly more accurate results. An illustration of the segmentation results using the original deterministic statistical region merging algorithm with various values of Q, and the results of the proposed stochastic region merging approach is shown in Fig. 4. It can be observed that, regardless of the value of Q selected, the deterministic approach is unable to achieve correct segmentation results, while the proposed stochastic approach is able to give accurate segmentation results. In fact, based on exhaustive tests for the Q range of $[20, 200]$, there are no values of Q that provides reasonable segmentation results for the set of test images.

B. Comparison with State-of-the-art Methods

For comparison purposes, we used the automatic skin lesion segmentation (SLS) approach proposed by Xu et al. [9], the multi-directional gradient vector flow (MGVF) approach proposed by Tang [15]. We also tested a state-of-the-art color image segmentation method recently proposed by Li et al. [22], which segments images using a region-based variational approach based on level set active contours (LSAC). The LSAC method was chosen as it represents the state-of-the-art in region-based active contour methods, as well as the SLS and MGVF methods, which were shown to provide strong segmentation accuracy for various types of skin cancer regions, thus acting as good indicators of the level of segmentation accuracy that can be achieved by current techniques in this area. It should be noted that the parameters of the LSAC, SLS, and MGVF methods have been set based on that presented in their respective literatures, given that the goal is for fully automatic skin lesion segmentation without manual intervention. The SLS and MGVF methods have been implemented based on their respective works in MATLAB, while the LSAC method used is that provided by the authors of [22].

Given the ground truth measurements by a trained expert (denoted by GT), the obtained segmentation results (denoted by SR) using the proposed system and the tested methods was computed over all test images and compared using three different metrics:

1) The segmentation error (SE) as defined by [17]

$$SE = \frac{A(SR \bigoplus GT)}{A(GT)} \times 100\%, \tag{5}$$

where $\bigoplus$ denotes XOR operation and $A(.)$ denotes the area.

2) The true detection rate (TDR) as defined by [25]

$$TDR = \frac{A(SR \bigcap GT)}{A(GT)} \times 100\%. \tag{6}$$

3) The false positive rate (FPR) as defined by [25]

$$FPR = \frac{A(SR \bigcap \bar{GT})}{A(GT)} \times 100\%. \tag{7}$$

The SE metric provides a good indication of overall segmentation performance, but does not give a good indication on specific characteristics of the segmentation methods. To complement this overall metric, the TDR metric provides a good indication of undersegmentation error while the FPR metric provides a good indication of oversegmentation error. As such, the use of these three metrics provide an overall picture of the different aspects of the segmentation algorithms. In order to evaluate the repeatability of the results given the stochastic nature of the algorithm, the overall segmentation error presented for the proposed approach is based on a total of 20 repeated trials for each image. Furthermore, note that, as discussed in [9], given that ground-truth measurements are made by trained experts, there exists minor variations between the ground-truth measurements made by different experts. Therefore, the ground truth measurements by a second

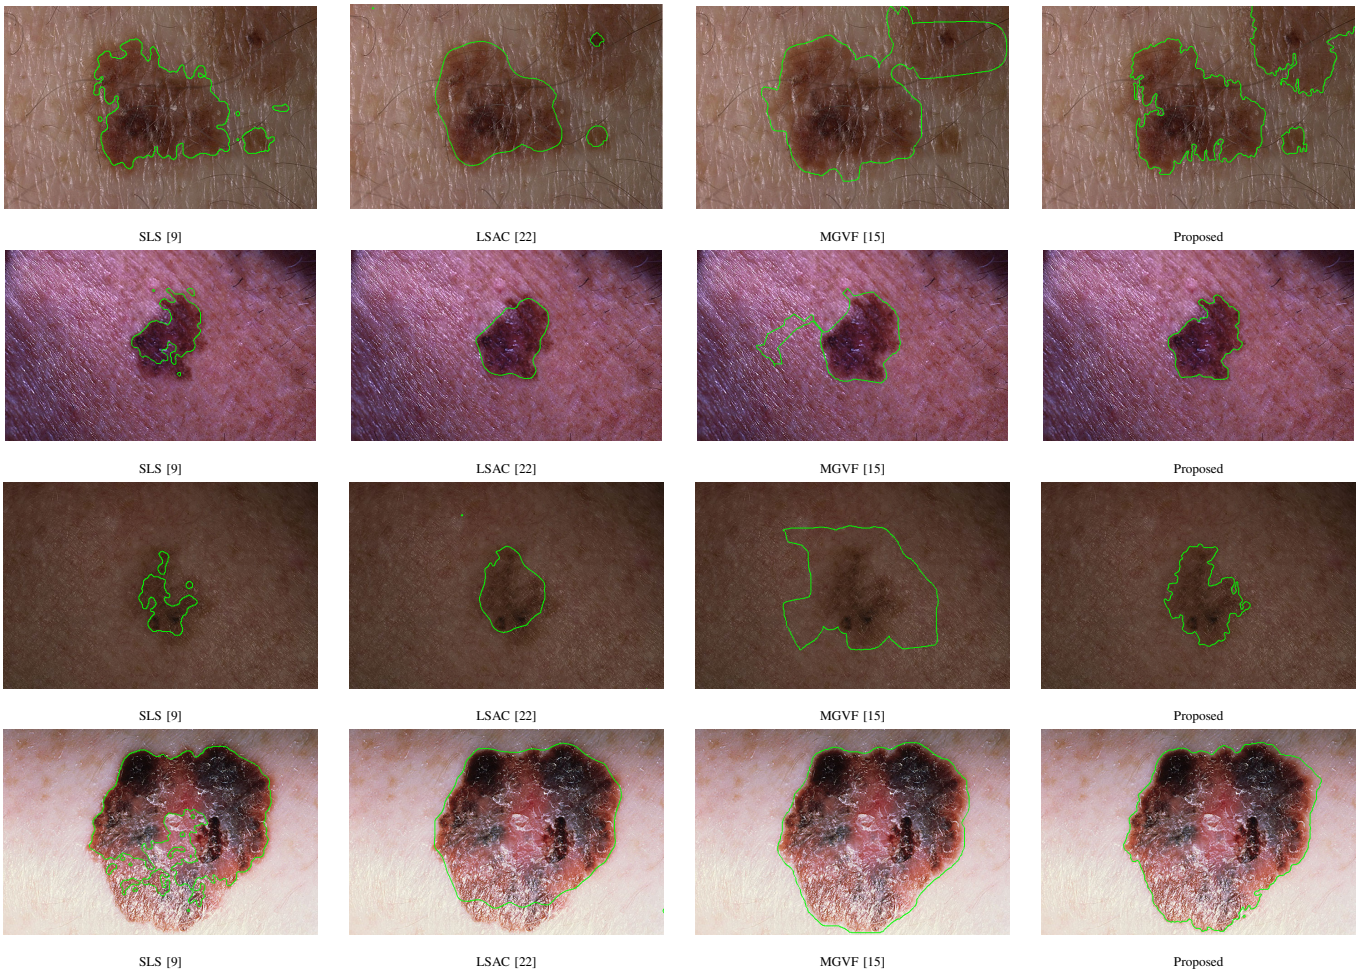

Fig. 5. Sample skin lesion segmentation results for Tests 1, 2, 3 and 4 (first to fourth rows, respectively).

trained expert was used to investigate inter-observer variation. All ground-truth measurements were produced using Adobe Photoshop. It was found that the inter-observer variation is within 3%, and as such allows for reasonable quantitative assessments. All the methods produced errors in their segmentations of the test images, and we verified statistically if these methods produced the same segmentation errors on average. This hypothesis was rejected using analysis of variance (i.e., the two-way ANOVA test) considering a significance level of 0.05 (i.e., $F = 29.50$ and $Prob(F > 29.50) < 0.0001$), suggesting that at least one of the methods produces an average segmentation error that is statistically different from the error averages for all tested methods. To identify the methods that have similar performances, we grouped the methods with statistically similar average segmentation errors using the Tukey's Honestly Significant Difference (HSD) test [23]. Considering a significance level of 0.05, the following three groups of segmentation methods were identified : (a) LSAC [22] and SLS [9] form group 1; (b) SLS [9] and MGVF [15] form group 2; and (c) the proposed method stands alone in group 3.

The segmentation performance metrics on the 60 real macroscopic images are shown in Table I, and shed several important insights on the performance and particularities of the proposed method. Firstly, the proposed method achieved the lowest SE, which indicates that the proposed method has the strongest overall performance when compared to the SLS, MGVF, and LSAC methods. Secondly, based on the TDR metric, it can be observed that the proposed method has a slight tendency to under-segment when compared to MGVF, but provides significantly better TDR than LSAC and SLS. Third, based on the FPR metric, it was observed that the proposed method resulted in relatively few false positives as with LSAC and SLS, while MGVF exhibited significantly higher number of false positives. Finally, the proposed method performs equally on both malignant and benign cases.

TABLE I
SEGMENTATION PERFORMANCE METRICS ON 60 REAL MACROSCOPIC IMAGES, BASED ON MEASUREMENTS MADE BY THE FIRST EXPERT.

Method	SE (%)	TDR (%)	FPR (%)
LSAC [22]	35.24	70.49	1.83
SLS [9]	27.98	75.28	2.94
MGVF [15]	26.74	95.92	15.78
Proposed[1]	9.16	93.06	3.03

[1]The results presented for the proposed approach is based on a total of 20 repeated trials for each image to evaluate the repeatability of the results given the stochastic nature of the algorithm.

Sample segmentation results from five of the tested macroscopic images are shown in Fig. 5. Visually, the proposed

method was able to segment the lesion regions with a higher level of accuracy for almost all tests when compared to the SLS, MGVF, LSAC methods. However, when compared to the proposed method in Test 4, MGVF provides a slightly more accurate segmentation at the top-right hand corner of the skin lesion. This suggests that our proposed method potentially can be more effective in segmenting skin lesion regions from macroscopic images than comparable state-of-the-art segmentation methods available in the literature. The effectiveness of the proposed method compared to the other tested methods may be attributed to two factors. First, unlike the other tested approaches, the proposed method introduces a stochastic aspect into the region merging process, which allows for increased robustness against being stuck in local optima, which is particularly important in complex situations occurring in dermatological image analysis, characterized by noise and artifacts. In fact, the key advantage of the proposed region-based approach over thresholding methods is its robustness to noise as well as illumination and color variations. Second, the iterative nature of the proposed method, where regions are merged from a pixel basis to region basis, lends itself better to situations characterized by strong illumination variations and weak boundary separation than approaches such as MGVF and LSAC. In fact, the key advantage of the proposed region-based approach over active contour based approaches is its robustness to illumination variations and weak boundary separation, which would result in oversegmentation or undersegmentation results for active contour based methods.

IV. Conclusions and Future Work

In this paper, we described an automatic method for segmenting skin lesions in macroscopic images based on iterative stochastic region merging. The images are acquired with conventional cameras, without using a dermoscope, which facilitates the remote image acquisition in telemedicine. The proposed method can be part of a system designed to help in the quantitative analysis of skin lesions for remote clinical pre-screening in teledermatology. Our experimental results indicate that our proposed segmentation method potentially can provide more accurate skin lesion segmentations from macroscopic images than comparable methods existing in the literature. In the next stage of our research work, we intend to test our method in larger datasets, including different kinds of pigmented skin lesions. Also, we plan to investigate how to incorporate new features to our segmentation approach to further improve its segmentation accuracy, such as using textural information in region merging, and techniques to compensate for illumination variations.

Acknowledgment

The authors would like to thank the Natural Sciences and Engineering Research Council (NSERC) of Canada, and Conselho Nacional de Desenvolvimento Cientifico e Tecnologico (CNPq) of Brazil, for providing funds for this project.

References

[1] Skin Cancer Foundation, “Skin Cancer Facts”, http://www.skincancer.org/Skin-Cancer-Facts/, Accessed: July 4, 2010.

[2] D. Koh, H. Wang, J. Lee, K. Chia, H. Lee, and C. Goh, “Basal Cell Carcinoma, Squamous-cell Carcinoma and Melanoma of the Skin: Analysis of the Singapore Cancer Registry Data 1968-97,” *Br. J. Dermatol.*, vol. 148, no. 6, pp. 1161-1166, 2003.

[3] I. Plesko, G. Severi, A. Obsitnkova, and P. Boyle, “Trends in the Incidence of Non-melanoma Skin Cancer in Slovakia, 1978-1995”, *Neoplasma*, vol. 47, pp. 137-142, 2000.

[4] E. de Vries, M. Louwman, M. Bastiaens, F. de Gruijl, and J. Coebergh, “Rapid and continuous increases in incidence rates of basal cell carcinoma in the southeast Netherlands since 1973,” *J. Invest. Dermatol.*, vol. 123, pp. 634-638, 2004.

[5] P. Boyle and B. Levin, *World Cancer Report 2008*, IARC Press, pp. 412–417, 2008.

[6] P. G. Cavalcanti and J. Scharcanski, *Automated Prescreening of Pigmented Skin Lesions Using Standard Cameras*, Preprint, Instituto de Informatica, UFRGS, Porto Alegre, Brazil, 2010.

[7] H. Oka, M. Hashimoto, H. Iyatomi, and M. Tanaka, “Internet-based program for automatic discrimination of dermoscopic images between melanoma and Clark nevi,” *British Journal of Dermatology*, vol. 150, no. 5, pp. 1041, 2004.

[8] H. Iyatomi, H. Oka, M. Emre Celebi, M. Hashimoto, M. Hagiwara, M. Tanaka, and K. Ogawa, “An improved Internet-based melanoma screening system with dermatologist-like tumor area extraction algorithm,” *Computerized Medical Imaging and Graphics*, vol. 32, pp. 566-579, 2008.

[9] L. Xu, M. Jackowski, A. Goshtasby, D. Roseman, S. Bines, C. Yu, A. Dhawan, and A. Huntley, “Segmentation of skin cancer images”, *Image and Vision Computing*, vol. 17, pp. 65-74, 1999.

[10] M. Celebi, Y. Aslandogan, and W. Stoecker, H. Iyatomi, H. Oka, and X. Chen, “Unsupervised border detection in dermoscopy images”, *Skin Research and Technology*, vol. 13, pp. 454-462, 2007.

[11] M. Hintz-Madsen, L. Hansen, J. Larsen, and K. Drzewiecki, “A probabilistic neural network framework for the detection of malignant melanoma”, *Artificial neural networks in cancer diagnosis. Prognosis and Patient Management*, pp. 141-183, 2001.

[12] M. Yuksel and M. Borlu, ”Accurate segmentation of dermoscopic images by image thresholding based on type-2 fuzzy logic”, *IEEE Transactions on Fuzzy Systems*, vol. 17, pp. 976-981, 2009.

[13] M. Silveira, J. Nascimento, and J. Marques, A. Marcal, T. Mendonca, S. Yamauchi, J. Maeda, and J. Rozeira, “Comparison of segmentation methods for melanoma diagnosis in dermoscopy images”, *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, pp. 35-45, 2009.

[14] B. Erkol, R. Moss, R. Stanley, W. Stoecker, E. Hvatum, “Automatic lesion boundary detection in dermoscopy images using gradient vector flow snakes”, *Skin Research and Technology*, vol. 11, pp. 17-26, 2005.

[15] J. Tang, “A multi-direction GVF snake for the segmentation of skin cancer images”, *Pattern Recognition*, vol. 42, no. 6, pp. 1172-1179, 2009.

[16] H. Zhou, G. Schaefer, A. Sadka, and M. E. Celebi, ”Anisotropic mean shift based fuzzy c-means segmentation of dermoscopy images”, *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, pp. 26-34, 2009.

[17] M. Celebi, H. Kingravi, H. Iyatomi, Y. Aslandogan, W. Stoecker, R. Moss, J. Malters, J. Grichnik, A. Marghoob, H. Rabinovitz, and S. Menzies, ”Border detection in dermoscopy images using statistical region merging”, *Skin Research and Technology*, vol. 14, pp. 347-353, 2008.

[18] R. Nock and F. Nielsen, “Statistical Region Merging”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 11, pp. 1452-1458, 2004.

[19] M. Celebi, H. Iyatomi, G. Schaefer, and W. Stoecker, ”Lesion border detection in dermoscopy images”, *Computerized Medical Imaging and Graphics*, vol. 33, pp. 148-153, 2009.

[20] S. Z. Li, Markov Random Field Modeling in Image Analysis. Springer, 2001.

[21] Q. Yu and D.A. Clausi, “IRGS: Image Segmentation Using Edge Penalties and Region Growing,” *IEEE TPAMI*, vol. 30, no. 12, pp. 2126-2139, 2008.

[22] C. Li, C. Kao, J. C. Gore, and Z. Ding, “Minimization of region-scalable fitting energy for image segmentation”, *IEEE Trans. Image Processing*, vol. 17, no. 10, pp. 1940-1949, 2008.

[23] R. Sokal and F. Rohlf, Biometry: the principles and practice of statistics in biological research. San Francisco: W. H. Freeman, 1994.

[24] A. Hammoude, “Computer-Assisted Endocardial Border Identification from a Sequence of Two-Dimensional Echocardiographic Images,” *Ph.D. dissertation*, University of Washington, Seattle, WA, 1988.

[25] M. Silveira, J. Nascimento, J. Marques, A. Marcal, T. Mendonca, S. Yamauchi, J. Maeda, and J. Rozeira, “Comparison of Segmentation Methods for Melanoma Diagnosis in Dermoscopy Images,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 3, pp. 35-45, 2009.