

Melanoma Detection Using Delaunay Triangulation

A. Pennisi*, D. D. Bloisi*, D. Nardi*, A. R. Giampetruzzi†, C. Mondino‡ and A. Facchiano†

*Department of Computer, Control, and Management Engineering
Sapienza University of Rome, Italy
{pennisi, bloisi, nardi}@dis.uniroma1.it

† Istituto Dermatologico dell'Immacolata IDI-IRCCS, Rome, Italy
{a.facchiano, a.giampetruzzi}@idi.it

‡ Department of Dermatology
Cantonal Hospital of Bellinzona, Switzerland
chiara.mondino@eoc.ch

Abstract—The detection of malignant lesions in dermoscopic images by using automatic diagnostic tools can help in reducing mortality from melanoma. In this paper, we describe a fully-automatic algorithm for skin lesion segmentation in dermoscopic images. The proposed approach is highly accurate when dealing with benign lesions, while the detection accuracy significantly decreases when melanoma images are segmented. This particular behavior lead us to consider geometrical and color features extracted from the output of our algorithm for classifying melanoma images, achieving promising results.

Keywords-Melanoma detection ; Dermoscopy images ; Automatic segmentation ; Border detection

I. INTRODUCTION

Melanoma is an aggressive tumor and it can be lethal, if not diagnosed in time [1]. Early diagnosis is today the main way to prevent mortality from melanoma. Indeed, malignant melanoma has a cure rate of more than 95% if detected at an early stage [2]. Thus, in order to reduce the number of deaths from melanoma, it is fundamental to develop innovative strategies to help dermatologists making early diagnosis. Furthermore, the development of reliable computer-aided diagnosis (CAD) tools can help in reducing the number of presumptive diagnoses that have to be confirmed histologically on skin biopsy.

The analysis of dermoscopic images is a key process in the early diagnosis of melanoma. Dermoscopic images are generated by combining optical magnification with either cross-polarized lighting or liquid immersion, with a low angle-of-incidence lighting. Analysing the dermoscopic image it is possible to extract information about the principal characteristics of the lesion, including symmetry, edges, color variety, size, presence and distribution of visual features (e.g., mesh of the net, black points, globules and striae), that are essential for a correct presumptive diagnosis.

In this paper, we describe a fully-automatic lesion segmentation method, which can be used to generate a binary mask of the lesion area (see Fig. 1). The proposed method is designed to be sensitive with respect to images containing irregular borders, multiple shades of pigmentation,




Figure 1. Skin lesion segmentation. a) Dermoscopic image in input. b) Binary mask in output. c) The border of the mask overlaid on the image. Images a and b are from the PH² database [3].

and varying texture. This is demonstrated by quantitative experimental results, carried out on the publicly available PH² database [3]. In particular, we show that the accuracy of the segmentation is extremely high when dealing with benign lesions (common and atypical nevi), while the precision of the segmentation results significantly decreases when malignant lesions (melanoma) are processed. This behavior lead us to consider the use of the binary masks generated by our algorithm as input for a classification stage. Preliminary results are promising and allow for considering the described method as a suitable tool for the development of CAD support systems for melanoma detection.

The remainder of the paper is organized as follows. Related work is discussed in Section II, while our skin lesion segmentation method is presented in Section III. Experimental results are shown in Section IV. Finally, conclusions are drawn in Section V.

II. RELATED WORK

Existing approaches in the literature dealing with the problem of segmenting skin lesion images can be classified according to the three following main categories [4].

Thresholding methods. Approaches in this category aims at comparing visual feature values for single or group of pixels in the dermoscopic image with threshold values. For example, a pixel can be classified as a lesion point if it is darker than a given color threshold. The thresholding process produces a binary image as output, which can be further processed by using morphological operators to filter




Figure 2. The four steps in our method.

out outliers, to fill small holes, or to select the largest connected component. Different local (per-pixel and per-region) and global (per frame) thresholding methods have been proposed: adaptive thresholding [5], histogram thresholding [6], and clustering [7]. The main drawback of thresholding methods is that they can achieve good results only if there is a high contrast between the lesion area and the surrounding skin region, which is not always the case.

Edge and contour-based methods. Algorithms belonging to this class rely on the identification of the discontinuities (i.e., the edges) in the dermoscopic images to detect the lesion borders. For example, an active contour method based on gradient vector flow (GVF) snakes for contour extraction is described in [8], while in [9] two contour based methods, i.e., adaptive snake and active contour by level set, are applied to skin lesion images. Edge and contour-based methods usually fail if the transition between the lesion and the surrounding skin is smooth and also in presence of hair.

Region-based methods. This category includes methods that exploit multi-scale region growing, multi-resolution Markov random fields, and statistical region merging. The basic assumption is that the image in input contains two different regions: lesion and skin. A method called JSEG [10], based on color quantization and spatial segmentation, has been applied to skin lesion images in [11] and [12]. The method uses J-images, corresponding to measurements of local homogeneities at different scales, to find potential boundary locations. The final segmentation is obtained by growing regions from seed areas of the J-images. Statistical region merging (SRM) is used in [13]. SRM treats the image as an observed instance of an unknown theoretical image, whose statistical regions are to be reconstructed. Region-based algorithms suffer from the possibility of generating over-segmentation results. In particular, when the skin or the lesion region are textured or the interior of the lesion exhibit multi-colored areas, region-based methods can produce inaccurate results.

In this paper, we present a region-based, fully-automatic algorithm for skin lesion segmentation. As a difference with previous work, we compute two parallel processes of skin

detection and lesion segmentation and then merge the results, thus obtaining an accurate binary mask containing the lesion area. The details of our algorithm are given in the next section.

III. SKIN LESION SEGMENTATION

The great variety of lesion shapes, size and colors, the different skin types and textures, as well as the possible presence of hair and air/oil bubbles make segmentation a hard task. In order to deal with such difficulties, we propose an approach for skin lesion image segmentation structured in four steps (see Fig. 2). In step 1, noise removal is performed by morphological closing. Then, two segmentation processes (steps 2 and 3) are carried out in parallel, yielding two different segmented images. Step 2, i.e., skin detection, aims at determining the lesion area by first detecting the skin region and then filtering it out. In step 3, i.e., lesion segmentation, the lesion area is extracted by using Edge detection and Delaunay Triangulation. In step 4, the results obtained in steps 2 and 3 are merged, in order to obtain the final detected lesion area.

Referring to Fig. 2, the artifact removal and image equalization process produces an image E , which is a filtered (possible hair are removed) and equalized version of the dermoscopic image I in input. E represents the input for both the skin detection and the lesion segmentation functions. The former generates a so-called skin image S by using a color thresholding mechanism, while the latter exploits a set of adaptive parameters to compute the so-called lesion image L , that contains the different color regions in E . During the merging step, S and L are analysed for possibly fusing adjacent regions, finally obtaining the binary image B .

A. Artifact Removal and Image Equalization

The RGB dermoscopic image I in input is processed in order to remove hair. I is filtered through a morphological transformation, thus obtaining a new RGB image F (filtered image). In particular, F is the result of a closing operation with an 11×11 kernel having each element $e_{ij} = 1$. The morphological transformation aims at removing the outlier pixels that can be introduced in the image acquisition phase,




Figure 3. Noise removal: original input image and filtered one.




Figure 4. Image equalization process. a) Input image and the corresponding luminance spectrum before equalization. b) Equalized Image: the luminance spectrum is modified.

while maintaining the properties of the lesion region (see Fig. 3 for an example of hair removal).

The filtered image F , coming from the artifact removal phase, is processed to get an equalized image E . The equalization step, performed applying the OpenCV function `equalizeHist` on the Y channel, helps in highlighting the lesion borders and in obtaining a more accurate output, since the color difference between the lesion area and the surrounding skin are stressed (see Fig. 4). The image E is used as input for both the parallel processes of Skin Detection and Lesion Segmentation.

B. Skin Detection

After equalization, the subsequent step consists in identifying the pixels in the image belonging to the skin. A number of methods for skin segmentation in color images are available in the literature. The simplest methods define boundaries in the chosen color space for identifying skin clusters. The main advantage of such methods is that they do not require a training phase. However, it is difficult to define the boundaries that give good results by considering




Figure 5. Skin detection process.

a single color space [14]. For such a reason, we adopt a combination of multiple color spaces.

The main steps in the skin detection process are shown in Fig. 5. E is converted into the YCrCb color space and the skin region is detected by using a thresholding on the luminance and chrominance values, producing an image T .

Then, T is converted into the HSV color space and normalized. The resulting normalized image N is further filtered by applying a second thresholding, this time on the HSV values, thus obtaining an image S , that contains the lesion area only, with the skin region filtered out.

C. Lesion Segmentation

The process of extracting the contours of the lesion area is derived from the algorithm proposed in [15] and it is arranged in two phases: *Edge Detection* and *Delaunay Triangulation*. The former is used for splitting the image into local coherent regions, the latter for aggregating homogeneous regions in a global fashion at image level.

The RGB input equalized image E is filtered by a Gaussian blur filter with a kernel size $\sigma = 5$. Then, it is converted to grayscale and the Edge Detection procedure begins with a Canny edge extraction, that leads to the creation of a grayscale image C containing the intensity edges in E (see Fig. 6). The two parameters min and max in the Canny algorithm have been set to the values 0.03 and 2.0 respectively, in order to focus on short edges in the input image. The detected edges are then vectorized into connected line segments — generated as described in [16] — and used as input for the Delaunay Triangulation procedure, which computes a triangular tessellation of the image.

The triangular graph is segmented by using a *region association* procedure, which iteratively finds and associates the two regions with the lowest normalized boundary cost, by considering a predefined association threshold ω . In particular, each of the triangles in the graph is considered in turn, by calculating the average HSV color of all the pixels that lie within its circumcircle: If a pair of triangles have a similar HSV value, then they are fused into a single triangle (see the example in Fig. 7).

The C++ source code for the image segmentation procedure is available at: www.dis.uniroma1.it/~pennisi/FHIS-Image_Segmentation_Library.html




Figure 6. Image Segmentation process.




Figure 7. Delaunay Triangulation. a) Input image from the PH² database [3]. b) Detail of the Delaunay Triangulation. c) Resulting triangles after association.




Figure 8. The results of the skin detection and the lesion segmentation processes are merged to obtain the final binary image B . Images are from the PH² database [3].

D. Merging

In the final step, the results generated by the two parallel processes of Skin Detection and Lesion Segmentation are merged. The idea is that a correctly extracted lesion blob from a binary mask can be circumscribed by a circle with a diameter equal to the major axis of the detected blob. First, the merging procedure detects if one (or more lesion areas) is (are) present in each image S and L . To this end, S and L are converted into binary images (called S_B and L_B , respectively) by assigning the value 255 to the pixels having color values different from the HSV value $\langle 0, 0, 0 \rangle$. In presence of multiple lesion areas, only the biggest one is considered.

Then, the probabilistic Hough transform is applied to each one of the two images, in order to obtain the number of circles that can be inscribed or circumscribed to the skin lesion area. We adopt the OpenCV function *HoughCircles*, with the following experimentally selected parameters: Inverse ratio of resolution = 1.0; Minimum distance between detected centers = 90 pixels; Thresholds for the internal Canny detector equal to 255 and 10.

If the number of the detected circles in S_B and L_B is greater than zero for both of them, then the final binary image B is the result of the pixel-wise logical AND of S_B and L_B , otherwise B is the image where almost one circle can be detected. Fig. 8 shows three examples for the merging procedure. The first row of Fig. 8 illustrates an example where a circle can be detected both in S_B and




Figure 9. Segmentation results on PH² images IMD046 (common nevus, first row), IMD048 (atypical mole, second row) and IMD058 (melanoma, third row).

L_B , thus the final image B is the pixel-wise logical AND of the two images. It can happen that the skin detection process generates an S_B image where *HoughCircles*, by using the predefined parameters, cannot find any circle (see the second row of Fig. 8). However, since a circle can be detected in the L_B image, then $B = L_B$. The third row of Fig. 8 shows an example where the final image B corresponds to S_B , since no circles can be detected in L_B by applying *HoughCircles*. It is worth noting that, for all the 200 images in the PH² database, it was always possible to find at least a circle in one of the two images S_B and L_B by using the above listed parameters.

IV. EXPERIMENTAL RESULTS

A quantitative experimental validation has been conducted on a publicly available database of dermoscopic images, containing ground truth annotations.

A. Data Set Description

The PH² database [3] has been realized by the Universidade do Porto, Técnico Lisboa in collaboration with the Hospital Pedro Hispano in Matosinhos, Portugal. The database is composed of 200 RGB dermoscopic images, with a resolution of 768×574 pixels and a magnification of $20\times$, annotated with ground truth data. The 200 images are divided into benign lesions (80 common and 80 dysplastic nevi) and malignant lesions (40 melanomas), with a skin color that varies from white to cream white.

For each image, the ground truth data include the following information:

- A ground truth binary image, manually generated by expert dermatologists, containing the skin lesion area;
- Clinical and histological diagnosis;
- Dermoscopic criteria.

In particular, in the provided ground truth binary image, the pixels with value 1 belong to the segmented lesion, while pixels with value 0 correspond to the background.

Dermoscopic criteria include asymmetry, colors, pigment network, dots/globules, streaks, regression areas, and blue-whitish veil.

B. Segmentation Results

All the 200 images in the PH² database have been segmented using the same parameters. Three examples of application for our algorithm are shown in Fig. 9: in the first row a common nevus is shown, in the second row an atypical mole, and in the third row a melanoma. It is worth noting that, for the images in the first and second rows of Fig. 9, the binary images obtained are:

- 1) In very good accordance with respect to the corresponding ground truth images in PH²;
- 2) The results are not affected by the presence of hair.

A situation where our method provides a binary image containing an under-estimated lesion area is shown in the third row of Fig. 9: This is an interesting behavior of the algorithm when dealing with melanoma images, which is discussed in the rest of this section. The complete set of the 200 binary images can be downloaded at: www.dis.uniroma1.it/~pennisi/skin_lesion_segmentation/results.zip

Four different metrics have been selected to calculate the segmentation results: *Sensitivity*, *Specificity*, *Accuracy*, and *F-measure*. The definitions for the used metrics are given in the following equations, where TP is the number of true positive pixels, FP is the number of false positive pixels, TN is the number of true negative pixels, and FN is the number of false negative pixels. The chosen metrics are widely used in the literature to measure the performance of skin lesion segmentation methods [17].

$$Sensitivity = \frac{TP}{TP + FN} \quad (1)$$

$$Specificity = 1 - \frac{FP}{TN + FP} \quad (2)$$

$$Accuracy = \frac{TP + TN}{TP + FN + TN + FP} \quad (3)$$

$$F-measure = \frac{1}{n} \sum_{i=1}^n 2 \frac{Prec_i \times Rec_i}{Prec_i + Rec_i} \quad (4)$$

where n is the total number of images and:

$$Rec_i(P) = TP_i / (TP_i + FN_i)$$

$$Prec_i(P) = TP_i / (TP_i + FP_i)$$

$$Rec_i(N) = TN_i / (TN_i + FP_i)$$

$$Prec_i(N) = TN_i / (TN_i + FN_i)$$

$$Rec_i = (Rec_i(P) + Rec_i(N)) / 2$$

$$Prec_i = (Prec_i(P) + Prec_i(N)) / 2$$

Table I shows the segmentation results obtained by considering the complete PH² data set and three different subsets of the images. Indeed, since the dermoscopic images in PH² are

Table I
SKIN LESION SEGMENTATION RESULTS

Input Images	Sensitivity	Specificity	Accuracy	F-measure
All	0.8024	0.9722	0.8966	0.8257
Common nevi Only	0.8717	0.9760	0.9477	0.8690
Atypical Moles Only	0.8640	0.9733	0.9271	0.8689
Melanomas Only	0.5404	0.9597	0.6615	0.6524

labeled according to their medical diagnosis, it is possible to carry out a finer analysis, by considering separately the three diagnostic classes: common nevi, atypical moles, and melanomas.

It can be noted from Table I that, when processing the 80 images of common nevi only, the sensitivity of our method increases from 0.8024 to 0.8717, the accuracy raises from 0.8966 to 0.9477, and the F-measure becomes 0.8690 from 0.8257. This means that very good results are achieved in segmenting images of common nevi. The same behavior can be observed by considering the segmentation performance on the 80 images of atypical moles only. In particular, the sensitivity increases from 0.8024 to 0.8640, the accuracy raises from 0.8966 to 0.9271, and the F-measure achieves 0.8689 from 0.8257. Thus, even in the case of dysplastic nevi (i.e., atypical moles), that are benign lesions, our algorithm obtains very good segmentation results.

On the other hand, a strong decrease in the quality of the segmentation results can be observed, on the totality of the used metrics, when only images containing melanomas are processed (see the last row in Table I). In particular, our algorithm presents a large decrease in the average accuracy, that becomes rather low (i.e., 0.6615) when compared to the accuracy obtained on all the PH² images (i.e., 0.8024).

Summarizing, a very high accuracy can be obtained when dealing with benign lesions, i.e., common nevi and atypical moles, while less accurate results are generated when melanoma images are processed.

C. Classification Results

The analysis of the segmentation results generated by evaluating the three classes of nevi separately leads to the following considerations:

- 1) For benign lesions (i.e., common and atypical nevi), the average accuracy is rather high (0.9477 and 0.9271, respectively);
- 2) For malignant lesions (i.e., melanoma images), the accuracy significantly decreases (0.6615).

This means that only in the case of malignant lesions our algorithm gives less accurate results. A possible motivation for such a behavior can be found in the studies described in [18] and in [19]. In those studies emerge that, the presence of light brown structureless areas in atypical




Figure 10. Segmentation results on PH² melanoma images IMD088 (streaks, regression areas, blue-whitish veil), IMD284 (blue-whitish veil, second row), IMD405 (blue-whitish veil), IMD419 (blue-whitish veil), IMD424 (streaks, blue-whitish veil), and IMD425 (regression areas, blue-whitish veil).

melanocytic lesions maybe very useful in differentiating atypical nevi from melanomas. Homogeneous areas and light brown structureless regions were the most sensitive and specific epiluminescence microscopy features for thin melanomas [19]. According to [18], particular attention is needed to melanocytic lesions that, over time, reveal a loss of network in favor of structureless areas and exhibit new colors such as dark brown, black, gray, blue, red, and white.

Since the proposed algorithm is based on a color region merging procedure for computing the segmentation results, it is strongly sensitive to structureless areas and homogeneous regions with a color different from the surrounding one. From the above considerations, it follows that our method tends to under-segment malignant lesions. Multiple examples of under-segmentation results in case of melanoma images are shown in Fig. 10, where images containing streaks, regression areas, and blue-whitish veil are processed.

The particular behavior of under-segmenting the lesion area when dealing with melanoma images, has been analyzed in order to understand if the binary masks can be used as input for a binary classifier. To this end, three features have been considered to represent the geometric properties of the detected lesion region:

- *Convex Area*: Scalar that specifies the number of pixel of the convex hull that contains the binary image;
- *Filled Area*: Scalar specifying the number of lesion

pixels in the binary image with all holes filled in.

- *Solidity*: Scalar specifying the proportion of the pixels in the convex hull that are also in the region. It is computed as $Area/ConvexArea$.

In addition to the three above described geometrical features, three color histograms for each dermoscopic image are computed. The first histogram of 255 bins represents the normalized hue (H) values extracted from the original dermoscopic image I by using the binary image B as a mask. The second and the third histograms contain the values for V and S , respectively, calculated in the same way of the H values.

For classifying the binary masks, we decided to compare four different well-known classifiers:

- *Adaboost*;
- *Naive Bayes*;
- *K-Nearest Neighbors* (KNN);
- *Random Trees* (RT).

The classifiers have been trained by taking into account a feature set made of the three above listed geometrical properties and the histograms related to the H , S , and V values. We selected the above listed classifiers since they are preferred when the number of images in each class varies, adopting the implementation provided by Weka (<http://www.cs.waikato.ac.nz/ml/weka/index.html>), a collection of machine learning algorithm developed by the University of Waikato, New Zealand.

We tested the proposed classification method by using a leave-one-out approach, as described in [20]. The classifiers are trained by using all the images except one, which is used for testing. Then, the process is repeated by changing the test image. The metrics selected for calculating the goodness of the classification process have been: Sensitivity, Specificity, F-measure, and Precision. The first three metrics have been computed as described above, while Precision, that represents the fraction of retrieved instances that are relevant, has been calculated as follows:

$$Precision = \frac{TP}{TP + FP} \quad (5)$$

Classification results are reported in Tables II, IV, VI, and VIII. The results show that by analysing the binary masks generated by the proposed segmentation approach it is possible to classify the instances in the PH² database by reaching good results in terms of sensitivity and specificity.

The Adaboost classifier is able to obtain a sensitivity of about 93.5% and a specificity of about 85.2%, thanks to the characteristic of combining rough and moderately inaccurate rules of thumb. In fact, the classifier is based on the observation that finding many rough rules of thumb can be easier than finding a single one, thus obtaining a highly accurate classifier. For such a reason, the classifier is able to recognize 154 over 160 nevi and 33 over 40 melanoma images (see Table III).

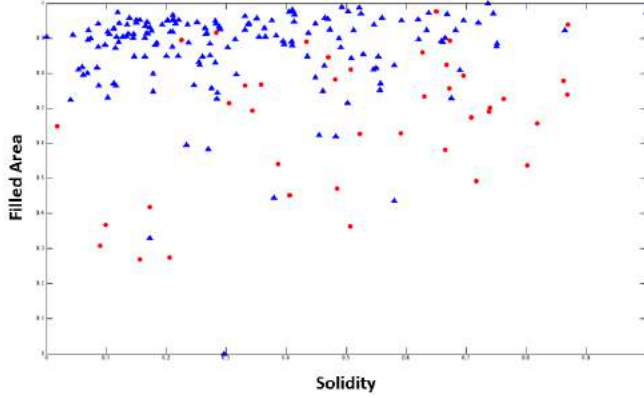


Figure 11. Filled Area plotted against Solidity. Melanoma images are represented as red circles and benign lesion images (common and atypical nevi) as blue triangles.

The Bayesian classification assumes that no dependencies exist among features and the training of such a classifier is based on the following main assumption: given a set of class variables, the value of a particular feature is independent of the value of any other feature. Therefore, a Bayesian classifier considers the contribution of each feature as being independent of the correlation probability between the single feature and the rest of the considered features. The Naive Bayes classifier is also preferred when the number of images in each class varies. Such reasons make the Naive Bayes classifier suitable for classification in our case: Indeed, the classifier can correctly classify 36 over 40 images of malignant lesions (see Table V), obtaining a sensitivity of 89.0% and a specificity of 89.7% (see Table IV).

The KNN and the RT classifiers obtain comparable results (see Table VI and Table VIII). However, both of them are based on the majority vote approach, which is influenced by the distribution of the features and thus, if the features are not well distributed as in our case — see for example Fig. 11, where filled area values are plotted against solidity values — then the classification cannot reach good results. Tables VII and IX show the confusion matrices related to the KNN and RT classifiers.

These preliminary classification results, obtained by considering only three geometrical features (i.e., convex area, filled area, and solidity) and three color histograms, are promising and allow to consider the use of our approach as a suitable tool for the development of CAD support systems for melanoma detection. Indeed, the classification results are rather accurate for all the four different considered classifiers.

V. CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, an automatic skin lesion image segmentation method, designed to deal with multiple types of lesion shapes, size and colors, and the presence of hair and air/oil bubbles, has been presented.

Table II
ADABOOST CLASSIFICATION RESULTS

Type	Sensitivity	Specificity	Precision	F-Measure
Nevus	0.963	0.825	0.963	0.960
Melanoma	0.825	0.962	0.846	0.835
Weighted Avg.	0.935	0.852	0.934	0.935

Table III
ADABOOST CONFUSION MATRIX

	Nevus	Melanoma
Nevus	154	6
Melanoma	7	33

Table IV
NAIVE BAYES CLASSIFICATION RESULTS

Type	Sensitivity	Specificity	Precision	F-Measure
Nevus	0.888	0.900	0.973	0.928
Melanoma	0.900	0.887	0.667	0.766
Weighted Avg.	0.890	0.897	0.911	0.896

Table V
NAIVE BAYES CONFUSION MATRIX

	Nevus	Melanoma
Nevus	142	18
Melanoma	4	36

Table VI
KNN CLASSIFICATION RESULTS

Type	Sensitivity	Specificity	Precision	F-Measure
Nevus	0.938	0.650	0.915	0.926
Melanoma	0.650	0.937	0.722	0.684
Weighted Avg.	0.880	0.707	0.876	0.878

Table VII
KNN CONFUSION MATRIX

	Nevus	Melanoma
Nevus	150	10
Melanoma	14	26

Table VIII
RANDOM TREE CLASSIFICATION RESULTS

Type	Sensitivity	Specificity	Precision	F-Measure
Nevus	0.913	0.675	0.918	0.915
Melanoma	0.675	0.912	0.659	0.667
Weighted Avg.	0.865	0.722	0.866	0.866

Table IX
RANDOM TREE CONFUSION MATRIX

	Nevus	Melanoma
Nevus	146	14
Melanoma	13	27

An experimental evaluation has been carried out on a publicly available database of dermoscopic images, in order to allow quantitative results on four different quality metrics.

The results demonstrate that the proposed method segments the lesion area in very good accordance with ground truth data only when images of benign lesions, i.e., common and dysplastic nevi, are considered, while the segmentation accuracy decreases considerably when the method is applied

to images of malignant lesions (i.e., melanomas). This behavior can be explained by the the presence in malignant lesion images of streaks, regression areas, and blue-whitish veil. Indeed, our algorithm is highly sensitive to structureless areas and homogeneous regions with a color different from the surrounding one and it generates a binary mask that presents a lesion area that is smaller than the actual one.

We exploited this particular sensitivity to images containing irregular borders (which is often the case of reticular pattern and atypical network in melanoma images) by considering geometrical features of the binary masks generated by our approach as input for four different binary classifiers, obtaining promising results. Although non suitable for diagnostic applications, the obtained classification results represent, in our opinion, a relevant starting point to further develop an automated analysis.

As future work, since the segmentation errors on melanoma images can easily be visually detected, we intend to study the possibility of integrating the described method into a computer-aided diagnosis (CAD) system.

REFERENCES

- [1] P. Roma, I. Savarese, A. Martino, D. Martino, and et al., "Slow-growing melanoma: Report of five cases," *Journal of Dermatological Case Reports*, vol. 1, no. 1, 2007.
- [2] D. Meckbach, J. Bauer, A. Pflugfelder, F. Meier, and et al., "Survival according to braf-v600 tumor mutations an analysis of 437 patients with primary melanoma," *Plos one*, vol. 9, no. 1, pp. 1–10, 2014.
- [3] T. Mendonca, P. M. Ferreira, J. S. Marques, A. R. Marcal, and J. Rozeira, "Ph² - a dermoscopic image database for research and benchmarking," *Conf Proc IEEE Eng Med Biol Soc*, vol. 2013, pp. 5437–5440, 2013, <http://www.fc.up.pt/addi/ph2>
- [4] F. Xie and A. C. Bovik, "Automatic segmentation of dermoscopy images using self-generating neural networks seeded by genetic algorithm," *Pattern Recognition*, vol. 46, no. 3, pp. 1012–1019, 2013.
- [5] G. Sforza, G. Castellano, S. A. Arika, LeAnder, and et al., "Using adaptive thresholding and skewness correction to detect gray areas in melanoma in situ images," *IEEE Transactions on Instrumentation and Measurement*, vol. 61, no. 7, pp. 1839–1847, 2012.
- [6] R. Garnavi, M. Aldeen, M. E. Celebi, G. Varigos, and S. Finch, "Border detection in dermoscopy images using hybrid thresholding on optimized color channels," *Computerized Medical Imaging and Graphics*, vol. 35, no. 2, pp. 105 – 115, 2011.
- [7] H. Zhou, M. Chen, L. Zou, R. Gass, L. Ferris, L. Drogowski, and J. Rehg, "Spatially constrained segmentation of dermoscopy images," in *5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2008, pp. 800–803.
- [8] B. Erkol, R. H. Moss, R. J. Stanley, W. V. Stoecker, and E. Hvatum, "Automatic lesion boundary detection in dermoscopy images using gradient vector flow snakes," *Skin Res Technol*, vol. 11, no. 1, pp. 17–26, 2005.
- [9] M. Silveira, J. C. Nascimento, J. S. Marques, and A. R. e. Marcal, "Comparison of segmentation methods for melanoma diagnosis in dermoscopy images," *J. Sel. Topics Signal Processing*, vol. 3, pp. 35–45, 2009.
- [10] D. Yining and B. S. Manjunath, "Unsupervised segmentation of color-texture regions in images and video," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 8, pp. 800–810, 2001.
- [11] M. E. Celebi, Y. A. Aslandogan, W. V. Stoecker, H. Iyatomi, H. Oka, and X. Chen, "Unsupervised border detection in dermoscopy images," *Skin Research and Technology*, vol. 13, no. 4, pp. 454–462, 2007.
- [12] P. Wighton, T. K. Lee, H. Lui, D. I. McLean, and M. S. Atkins, "Generalizing common tasks in automated skin lesion diagnosis," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 15, no. 4, pp. 622–629, 2011.
- [13] M. E. Celebi, H. A. Kingravi, H. Iyatomi, and Y. A. e. Aslandogan, "Border detection in dermoscopy images using statistical region merging," *Skin Research and Technology*, vol. 14, no. 3, pp. 347–353, 2008.
- [14] F. Gasparini and R. Schettini, "Skin segmentation using multiple thresholding," in *Internet Imaging VII, IS&T/SPIE Electronic Imaging*, vol. 6061. SPIE, 2006, pp. 60 610F–1, 60 610F–8.
- [15] C. J. Taylor and A. Cowley, "Parsing indoor scenes using rgb-d imagery," in *Proceedings of Robotics: Science and Systems*, 2012, pp. 401–408.
- [16] Q. Wu and Y. Yu, "Two-level image segmentation based on region and edge integration," in *Digital Image Computing: Techniques and Applications*, 2003, pp. 957–966.
- [17] A. Madooei, M. S. Drew, M. Sadeghi, and M. S. Atkins, "Automated pre-processing method for dermoscopic images and its application to pigmented skin lesion segmentation," *Color and Imaging Conference*, vol. 2012, no. 1, pp. 158–163, 2012.
- [18] V. Terushkin, S. W. Dusza, A. Scope, G. Argenziano, and et al., "Changes observed in slow-growing melanomas during long-term dermoscopic monitoring," *The British Journal of Dermatology*, vol. 166, no. 6, pp. 1213–1220, 2012.
- [19] G. Annessi, R. Bono, F. Sampogna, T. Faraggiana, and D. Abeni, "Sensitivity, specificity, and diagnostic accuracy of three dermoscopic algorithmic methods in the diagnosis of doubtful melanocytic lesions," *Journal of the American Academy of Dermatology*, vol. 56, no. 5, pp. 759–767, 2007.
- [20] C. Barata, M. Ruela, M. Francisco, T. Mendonca, and J. Marques, "Two systems for the detection of melanomas in dermoscopy images using texture and color features," *Systems Journal, IEEE*, pp. 1–15, 2013.