UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM COMPUTAÇÃO

PABLO GAUTERIO CAVALCANTI

Automated Prescreening of Melanocytic Skin Lesions Using Standard Camera Images

Thesis presented in partial fulfillment
of the requirements for the degree of
Doctor of Computer Science

Prof. Dr. Jacob Scharcanski
Advisor

Porto Alegre, May 2013

CIP – CATALOGING-IN-PUBLICATION

Gauterio Cavalcanti, Pablo

Automated Prescreening of Melanocytic Skin Lesions Using Standard Camera Images / Pablo Gauterio Cavalcanti. – Porto Alegre: PPGC da UFRGS, 2013.

108 f.: il.

Thesis (Ph.D.) – Universidade Federal do Rio Grande do Sul. Programa de Pós-Graduação em Computação, Porto Alegre, BR–RS, 2013. Advisor: Jacob Scharcanski.

1. Melanocytic skin lesions classification. 2. Image processing. 3. Melanoma. I. Scharcanski, Jacob. II. Título.

“A noble spirit embiggens the smallest man.”
— JEBEDIAH SPRINGFIELD

CONTENTS

LIST OF ABBREVIATIONS AND ACRONYMS

ABCD	Asymmetry, Border irregularity, Color variation and Diameter
ABCD	(Dermoscopy) Asymmetry, Border irregularity, Color variation and Differential structures
CART	Classification and Regression Tree
CFS	Correlation-based Feature Selector
EM	Expectation-Maximization
HSV	Hue Saturation Value
ICA	Independent Component Analysis
KNN	K-Nearest Neighbors
LMT	Logistic Model Tree
MPSL	Macroscopic Pigmented Skin Lesion
PCA	Principal Component Analysis
RGB	Red Green Blue
UVA	Ultraviolet A

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Melanoma is a type of malignant pigmented skin lesion, and currently is among the most dangerous existing cancers. However, differentiating malignant and benign cases is a hard task even for experienced specialists, and a computer-aided diagnosis system can be an useful tool. Usually, the system starts by pre-processing the image, i.e. removing undesired artifacts such as hair, freckles or shading effects. Next, the system performs a segmentation step to identify the lesion boundaries. Finally, based on the image area identified as lesion, several features are computed and a classification is provided.

In this Thesis, presented as a collection of published papers, we detail approaches to automatically execute all these steps, resulting in a pre-diagnosis for a pigmented skin lesion based only in a standard camera image (i.e. a simple color photograph). We tested our methods on publicly available datasets and achieved better segmentation and classification results than methods previously proposed in the literature.

Keywords: Melanocytic skin lesions classification, Image processing, Melanoma.

Análise Automática de Lesões de Pele Melanocíticas Utilizando Imagens de Câmeras Convencionais

RESUMO

Melanoma é um tipo maligno de lesão de pele pigmentada, e atualmente está entre os tipos de câncer existentes mais perigosos. Entretanto, diferenciar casos malignos de benignos é uma tarefa difícil mesmo para experientes especialistas, e um sistema de diagnóstico auxiliado por computador pode ser uma ferramenta bastante útil. Normalmente, este sistema inicia por um pré-processamento da imagem, isto é, remoção de artefatos indesejados, como pelos, sardas ou efeitos de sombreamento. A seguir, o sistema executa uma etapa de segmentação, identificando as bordas da lesão. Por fim, baseando-se na área da imagem identificada como lesão, diversas feições são computadas e uma classificação é gerada.

Neste tese, apresentada na forma de uma coleção de artigos publicados, nós apresentamos técnicas para automaticamente executar todos estes passos, resultando em um pré-diagnóstico para uma lesão de pele pigmentada baseado apenas em uma imagem convencional (uma simples fotografia). Nós testamos nossos métodos em bases de imagens públicas e atingimos melhores resultados de segmentação e classificação que os demais métodos presentes na literatura.

Palavras-chave: classificação de lesões de pele melanocíticas, processamento de imagens, melanoma.

1 INTRODUCTION

Melanocytes, cells located in the bottom layer of the skin's epidermis, are responsible for the production of melanin, the pigment responsible for skin colour. Neoplastic transformation of melanocytes results in a tumor known as a melanocytic lesion, usually visible as a pigmented growth on the skin. The benign type is called melanocytic nevus (also referred to as a "mole") and tends to appear in anyone's skin without being dangerous. The malignant type is called a melanoma and it is among the most dangerous forms of cancer.

The early diagnosis of melanomas is essential for patient prognosis since most malignant skin lesion can be treated successfully in their initial stages. However, there is a benign melanocytic lesion know as atypical nevus or dysplastic nevus that shares at least some, and sometimes all, of the clinical characteristics of a malignant melanoma. Consequently, discriminating benign from malignant lesions is often challenging, sometimes in the extreme (RAO et al., 1997; FIKRLE; PIZINGER, 2007).

Currently, the most reliable method for identifying malignant melanomas is histopathology analysis - requiring surgical excision. While this is not a significant problem in most cases, it is problem for patients with the atypical-mole syndrome, presenting many melanocytic nevi (i.e. 100 nevi or more (RAO et al., 1997)). In these cases histopathology is a time consuming, expensive, and uncompfortable option.

According to World Health Organization (2011), approximately 132,000 melanoma cases occur globally each year. In Brazil, according to the Ministry of Health, the incidence was 6 cases per 100 thousand inhabitants in 2008 and 2009 (DATASUS, 2011). Table 1.1 presents the average numbers of cases by regions in Brazil. The South region has the highest incidence and according to (BAKOS et al., 2002) and (BONFÁ et al., 2011), a combination of three factors is responsible:

- The population of the South region is predominantly light-skinned. Light-skinned people are at greater risk of developing melanomas;
- Lower latitudes present higher levels of UVA (Ultraviolet A), known risk factor for melanoma;
- Cultural practices such sun-bathing and working, without the use of sunscreen or other protection.

Inhabitants in other regions of the world with low latitudes and high percentages of light-skinned people, such as Scandinavia and Oceania, also have higher probabilities of developing with malignant melanomas (PEREZ-GOMEZ et al., 2004; CUST et al., 2009). The highest incidence is in Queensland, Australia (BAXTER et al., 2008).

Table 1.1: Malignant melanoma incidence in the regions of Brazil (cases for each 100 thousand inhabitants).

Region	Males	Females
North	0.68	0.38
Northeast	0.86	0.89
Southeast	3.76	3.88
South	7.31	6.58
Central-West	1.88	1.84

1.1 Diagnosis of Melanocytic Skin Lesions

Initial diagnosis of melanomas is most commonly done by visual appraisal of a melanocytic lesion (WURM; SOYER, 2010). Since dysplastic nevi and melanomas can occur on any skin area, it is essential that a full body examination is done on every patient, regardless of age (KOPF, 2007). The appraisal is performed based on the ABCD rule, a simple mnemonic that was created to alert both the general community and health care professionals of the main features of malignant melanomas (FRIEDMAN; RIGEL; KOPF, 1985) :

- A = Asymmetry. Independently where the lesion is bisected, one half of the lesion does not match the other in terms of shape and/or lesion content.
- B = Border irregularity. The perimeter of the lesion is uneven, undulating, ragged, notched, or blurred.
- C = Colour. The lesion presents a mixture of brown, black, red, white and blue, producing a mottled appearance.
- D = Diameter > 6 mm. This is not an inviolate rule, and currently a significant portion of melanomas are diagnosed by experts when these cancers are $<$ 6 mm in diameter.

Beyond the ABCD characteristics, the E letter may be added when the patient has a history of melanocytic lesions. The E stands for Evolving, since melanomas usually melanomas usually change in size, shape, colour, symmetry, etc., while common melanocytic nevi evolve slowly and reach a final stage of growth usually within the first few decades of life.

There are other visual methods, such as the Menzies Scoring Method or the 7-point Checklist (JOHR, 2002), that specialists use to identify malignant lesions. Both of these these methods help the physician to identify the characteristic features of melanomas.

Based on visual examination, other methods can be applied to provide a more definitive diagnosis. Dermoscopy (also known as epiluminescence microscopy, dermatoscopy, and amplified surface microscopy) is frequently to magnify submascroscopic lesion structures. The ABCD rule of dermoscopy (NACHBAR et al., 1994), changes the D letter to Differential structures since the diameter is, at this stage in the diagnosing process, not a reliable parameter. Dermatologists look for the submacroscopic morphologic and vascular structures that are only visible using a dermoscope and are typical of malignant melanomas.

Dermoscopy can increase the diagnostic sensitivity by 10-27% compared to simple visual appraisal according to MAYER (1997). However, even with the assistance of dermoscopy, differentiating malignant and benign lesions is difficult. Some specialists claim that in the early stages of malignant lesions, dermoscopy does not improve the diagnosis accuracy (SKVARA et al., 2005).

Other methods, based on emerging imaging technologies, are being developed. While some of these methods show promise, the diagnostic reliability of most of them has yet to be demonstrated. These methods include: multispectral analysis (GUTKOWICZ-KRUSIN; RABINOVITZ, 2007), reflectance spectroscopy MARCHESINI; BONO; CARRARA (2009); ZONIOS et al. (2010); MATTHEWS et al. (2011), confocal laser scanning microscopy and optical coherence tomography (SMOLLE, 2007).

1.2 Automated Melanocytic Skin Lesions Screening Systems

The work presented in this thesis is based on the proposal that computational analysis of melanocytic skin lesions images could approximate, or even exceed, the accuracy of that done by a dermatologist using dermoscopic analysis. Such a system must, *a priori*, be based on what characteristics dermatologists use diagnose melanomas, and how the characteristics are used in the diagnostic process. Such a system could, *a posteriori*, exceed the accuracy of the human diagnostician by enhancing detection and discrimination beyond the limits of the human visual system. Having proposed this, this, or any other, system must be considered as a tool that contributes to diagnostic accuracy and not as a replacement for the diagnostic process.

An automated screening system must provide the same, or better, sensitivity (signal detection) and specificity (signal descrimination) as the best human diagnostician using dermoscopy, and, consequently, a high degree of accuracy (RAO et al., 1997). This does not mean that a screening system needs to achieve 100% accuracy. Even experienced dermatologists do not obtain 100% of accuracy, independent of the method in use (ARGENZIANO; ZALAUDEK; SOYER, 2004).

Many approaches for classifying lesion images as malignant or benign have been proposed in the last decades. Although most of these approaches were proposed as computer-aided diagnosis systems, the resultant diagnoses were not entirely reliable. These systems are more appropriate as screening or triage tools, i.e. as an aid in determining the probability malignancy, and subsequent order and priority of treatment.

Screening systems for melanocytic skin lesions usually include the following steps:

- Preprocessing: this initial stage is responsible by removing image artifacts that may complicate the subsequent steps. For instance, the presence of hair and/or shading effects may decrease the discrimination of the lesion border, and consequently interfere with lesion classification;

- Segmentation: partitioning of the image into healthy skin and lesion areas, determining the lesion boundaries and identifying the regions of interest;

- Feature extraction: quantification of important lesion characteristics used for classification;

- Classification: an algorithm that uses the identified features computed to classify the lesion as malignant or benign.

The performance of these systems is usually based on three statistical measures: (1) Sensitivity, i.e. the percentage of malignant (melanomas) cases correctly classified; (2) Specificity, i.e. the percentage of benign (nevi) cases correctly classified; and, (3) Accuracy, i.e. the overall percentage of cases correctly classified.

Efforts to develop a screening system for melanocytic skin lesions started in the 90s. SCHINDEWOLF et al. (1994) were the first to report on an automated analysis system using digital cameras, which were not very common at the time. The use of a classification and regression tree (CART) (BREIMAN et al., 1984) resulted in 91% sensitivity, 83% specificity and, 85% accuracy. In subsequent years, the use of dermoscopy became more widespread, and most of the screening systems proposed in the literature rely on dermoscopy images.

Several approaches have already been proposed for the segmentation step, usually including preprocessing. These approaches are based on different types of algorithms, such as: region-based, usually region growing (IYATOMI et al., 2008) (CELEBI et al., 2008); clustering techniques (GOMEZ et al., 2008); thresholding algorithms (GANSTER et al., 2001) (GARNAVI et al., 2011); iterative procedures, such as active-contours or snakes (SILVEIRA et al., 2009); and combinations of the above (ZHOU et al., 2011; WANG et al., 2011).

Complete systems for the analysis of melanocytic skin lesions have been proposed. However, many of these systems, described in the medical literature, do not provide details about the algorithms used. For example SEIDENARI; PELLACANI; PEPE (1998) compared the results of examination by two physicians with the results from an automated analysis system using 90 dermoscopy images. The non-experienced physician achieved 74% sensitivity and 75% specificity. The experienced physician achieved 81% sensitivity and 95% specificity, and the automated system achieved 93% sensitivity and 95% specificity. However, no processing details were presented and, as such, the publication could not be considered as a reference in later works. FIKRLE; PIZINGER (2007) report 91.3% sensitivity and 80.8% specificity using a system called microDERM. Again, processing details were not provided.

Approches proposed in the informatics literature normally detail all of the processing steps used for the classification of dermoscopic lesion images. IYATOMI et al. (2008) were the first to present an online system (http://dermoscopy.k.hosei.ac.jp) to which registered users can send an image and receive an classification. After a region-based segmentation step, their system computes 428 features (80 relying on symmetry, 32 on border definition, 140 on colour, and 176 on texture). Next, a feature selection step reduces the data dimensionality, and a Neural Network is used for the classification. According to their experiments using 1258 dermoscopy images, their system achieves 85.9% sensitivity and 86% specificity.

CELEBI et al. (2007) proposed a system based on the extraction of 437 features (11 based on lesion shape, 354 on colour and 72 on texture). After feature normalization and selection, a Support Vector Machine classifier produced results with 92.34% sensitivity and 93.33% specificity.

Most of the proposed approaches are based on the ABCD rule of dermoscopy. However, WADHAWAN et al. (2011) choose the 7-point Checklist as reference for the computed features. Using a linear Support Vector Machine, their experiment with 347 dermoscopy images resulted in 87.27% sensitivity and 71.31% specificity.

The Internet provides an opportunity to acquire images locally and submit them for remote classification, by a dermatologist or a screening system. This dimension of tele-

medicine would benefit both patient and physician: providing remote access to expert knowledge, reducing the inconvenience of travel, reducing the wait times for exam procedures and providing results in a short time (MASSONE et al., 2008; WHITED, 2006).

SCHINDEWOLF et al. (1994) experimented with high resolution images acquired using commercially available digital cameras, including cameras present in tablets and smartphones. The use of such devices for imaging would certainly facilitate access to automated screening systems. However, preprocessing and segmentation of melanocytic skin lesions in standard camera images has not received much attention in the literature.

Most of the proposed segmentation techniques convert the original colour image to a monochromatic image, and use a thresholding algorithm to identify the lesion area (MANOUSAKI et al., 2006; RUIZ et al., 2008; TABATABAIE; ESTEKI; TOOSSI, 2009; ALCON et al., 2009). Recently, the use of active-contours segmentation algorithms has been proposed in some approaches (PAROLIN; HERZER; JUNG, 2010; TANG, 2009).

Feature extraction is usually based on the ABCD rule of dermoscopy. However, systems that use standard camera images cannot measure lesion diameter because image acquisition standards do not exist, yet. Consequently, systems usually try to identify the differential structures using texture descriptors. MANOUSAKI et al. (2006) proposed to extract 43 features to characterize the lesions: 9 related to its geometry, 28 to its colours, 2 to its borders, and 4 based on texture. Using logistic regression, their experiments with 132 images obtained 89.4% accuracy, 60.9% sensitivity and 95.4% specificity.

TABATABAIE; ESTEKI; TOOSSI (2009) proposed using Independent Component Analysis (ICA) to quantify lesion texture information. Based on image samples (sub-images of 16×16 pixels), they extracted the 100 most frequent independent components from images of benign and malignant lesions. Next, they performed feature selection based on the t-test separability criterion, and classified the images using Support Vector Machines. Using only those ICA-based texture data, their experiments achieved 73.7% accuracy. Next, they included colour features (average and standard deviation of each RGB channel), and increased the accuracy to 88.7%.

Another important screening system for melanocytic skin lesions using standard camera images has been proposed by ALCON et al. (2009). They proposed a set of 55 lesion features: 12 based on lesion shape, 2 on border discrimination, 21 on colour, and 20 on texture. After feature selection to identify only the 5 most significant features, a Logistic Model Tree provide lesion classification. Their experiments with 152 images achieved 86.84% accuracy. Then, including the Adaboost meta-classifier (FREUND; SCHAPIRE, 1997), their system achieved 89.72% accuracy. The accuracies obtained by those proposed approaches are usually higher than the diagnostic accuracy obtained by trained physicians, which range between 31% and 85% according to the literature (WHITED, 2006).

With this in mind, the objective of the current work was to automate and integrate the four processing stages and to develop a an automated melanocytic lesion classification system, using conventional digital camera images, that provides the same or better sensitivity and specificity than that provided by experienced dermatologists. Further, that the system be usable locally or remotely, via telemedicine systems, by persons with little or no diagnostic experience.

2 THESIS OBJECTIVES

This Thesis is was developed and published as a set of papers, and this document is organized as a collection of papers. Each one of the Sections of Chapter 3 is, in fact, a paper that has already been published. Every paper is self-contained and consequently can be read independently. The order of presentation is chronological, so the Sections show papers in the same order they have been published.

The objective of these papers is to present a framework for the automated prescreening of melanocytic skin lesions, trying to differentiate malignant from benign cases using only standard camera images (i.e., simple photographs). As described in Section 1.2, an automated prescreening system for melanocytic skin lesions requires sequential steps that preprocess, segment, extract features and classify the images. Accordingly, the approaches presented next were proposed to address at least one of these steps.

The goals and contributions of each one of these published papers are:

- Section 3.1: we present a preprocessing procedure to attenuate shading effects in human skin color images. Specially in melanocytic skin lesion images, if used directly in the segmentation process, shading and lesion regions could be confused. So, we proposed in this paper an automatic method to prevent this possibility. We also present experiments attenuating shading effects in face images.

- Section 3.2: this paper presents a complete automated prescreening system. It starts with the preprocessing method presented in Section 3.1. Then, a new multi-channel representation is created to better differentiate healthy skin and lesion areas, and this representation is used to generate a segmentation mask. Next, we extract a set of 52 features based on the ABCD rule of dermoscopy. Using the Nearest Neighbor Classifier, and a combination of this classifier with a Decision Tree, this system obtained highly accurate classification results.

- Section 3.3: the image segmentation is an important task in a screening system, and in this paper we presented a parameter-free segmentation method that in our experiments resulted in lower average segmentation error in comparison to state-of-the-art methods proposed in literature. The method starts using Independent Component Analysis to locate skin lesions in the image, and this location information is further refined by an active-contours segmentation method.

- Section 3.4: we propose in this paper novel lesion features based on the melanin information of lesion areas. According to recent studies, the melanin quantity in malignant melanomas may vary in comparison to benign lesions. Our approach tries to determine this variation using only standard camera images. We included

those features in a second classification stage to the system proposed in Section 3.2, and our preliminary experiments indicate that this two-stage approach can achieve accuracy levels near to 100% with adequate training.

After these Sections detailing the four most important papers published during the Doctorate, we also included Section 3.5 detailing other publications of the author of this Thesis.

It is important to observe that the paper in Section 3.3 about image segmentation is presented between two papers about complete screening systems (Sections 3.2 and 3.4), including feature extraction and classification steps. Ideally, a image segmentation technique should have been presented before these steps. However, the order of presentation is chronological, as already mentioned, and we decided to include the paper in Section 3.3 in this Thesis since this method achieved improved segmentation results in comparison to the method used in Section 3.2. Meantime, after proposing this segmentation method presented in Section 3.3, we performed some experiments comparing this technique to other techniques proposed for melanocytic skin lesions segmentation and how the segmentation results influence on the final classification result. These experiments are described in a published book chapter, and the obtained results indicate that the features usually extracted based on the ABCD rule are not significantly influenced by the segmentation step (for more details about these experiments, we included the referred book chapter as Appendix A of this Thesis). Consequently, we decided to use in the paper presented in Section 3.4 the same segmentation method used in Section 3.2, since this technique is less complex and usually also very efficient. Moreover, in this way the potential of the second classification stage based on melanin features presented in Section 3.4 can be easier analyzed.

3 COLLECTION OF PUBLISHED PAPERS IN JOURNALS AND CONFERENCES PRESENTING THE DEVELOPMENT OF THIS THESIS

3.1 Shading Attenuation in Human Skin Color Images

Cavalcanti, P. G. ; Scharcanski, J. ; Lopes, C. B. O. . Shading Attenuation in Human Skin Color Images. In: 6th International Symposium on Visual Computing (ISVC), 2010, Las Vegas. Lecture Notes in Computer Science: Advances in Visual Computing, 2010. v. 6453. p. 190-198.

Abstract

This paper presents a new automatic method to significantly attenuate the color degradation due to shading in color images of the human skin. Shading is caused by illumination variation across the scene due to changes in local surface orientation, lighting conditions, and other factors. Our approach is to estimate the illumination variation by modeling it with a quadric function, and then relight the skin pixels with a simple operation. Therefore, the subsequent color skin image processing and analysis is simplified in several applications. We illustrate our approach in two typical color imaging problems involving human skin, namely: (a) pigmented skin lesion segmentation, and (b) face detection. Our preliminary experimental results show that our shading attenuation approach helps reducing the complexity of the color image analysis problem in these applications.

3.1.1 Introduction

Interpret the shading of objects is a important task in computer vision. This is specially true when dealing with human skin images, because the color of structures can be significantly distorted by shading effects. The occurrence of shading depends mainly on the color of the object and the light illuminating them. However, roughness of the surface, the angles between the surface and both the light sources and the camera, and the distance of the surface from both the light sources and the camera, can also significantly influence the way the scene is processed (SHAPIRO; STOCKMAN, 2001). Specifically, human skin images are impacted by these factors and the analysis of these images can become difficult if the uneven illumination is not correctly understood and corrected.

In teledermatology, for example, often a standard camera color image containing a skin lesion is transmitted to a specialist, or analyzed by a pre-diagnosis system, without special attention to the illumination conditions (WHITED, 2006)(MASSONE et al., 2008). However, these conditions can affect the quality of the visualization, and impact on the physician diagnosis, or limit the efficiency of the pre-screening system. Pigmented

skin lesions typically have low diagnosis accuracy if the illumination condition is insufficient. These lesions usually are darker than healthy skin, and automatic approaches to segment such lesions tend to confuse shading areas with lesion areas. As a consequence, the early detection of malignant cases is more difficult without removing shading effects from the images. Considering that melanoma is the most dangerous type of pigmented skin lesion, and that this disease results in about 10000 deaths in 40000 to 50000 diagnosed cases per year (only considering the United States of America (Melanoma Research Project, 2010)), any contribution to improve the quality of these images can be an important step to increase the efficiency of pre-diagnosis systems, and to help to detect cases in their early-stages.

Another important human skin color imaging application that is severely affected by shading effects is face detection. In this case, color images containing human skin are used in head pose estimation or in face recognition systems, and shading effects may occlude some important features of the face (e.g., eyes, nose, head geometry). Usually, it is not feasible to control the illumination condition during image acquisition, and an automatic preprocessing step to mitigate these effects is an important contribution to these systems efficiency, as will be illustrated later.

In this paper, we propose an new automatic approach to attenuate the shading effects in human skin images. In Section 3.1.2, we describe the algorithm that executes this operation. In Section 3.1.3, some preliminary experimental results of our method are shown, focusing on the benefits of this operation for the color image analysis of pigmented skin lesions and face images. Finally, in Section 3.1.4 we present our conclusions.

3.1.2 Our Proposed Shading Attenuation Method

Our method for shading effect attenuation improves on the approach proposed by Soille (SOILLE, 1999). He proposed to correct the uneven illumination in monochromatic images with a simple operation:

$$R(x, y) = I(x, y) \ / \ M(x, y), \tag{3.1}$$

where, R is the resultant image, I is the original image, $M = I \bullet s$ is the morphological closing of I by the structuring element s, and (x, y) represents a pixel in these images. The main idea behind Soille method is to use the closing operator to estimate the local illumination, and then correct the illumination variation by normalizing the original image I by the local illumination estimate M. The division in Eq. 3.1 relights unevenly illuminated areas, without affecting the original image characteristics. Unfortunately, it is often difficult to determine an efficient structuring element for a given image, specially for human skin images that have so many distinct features, such as hair, freckles, face structures, etc. In this way, the results tends to be unsatisfactory for this type of images, as can be seen in Figs. 3.1(b)-(c).

Our method modifies the Soille approach by providing a better local illumination estimate M. In order to provide this local illumination estimate, we start by converting the input image from the original RGB color space to the HSV color space, and retain the Value channel V. This channel presents a higher visibility of the shading effects, as proposed originally by Soille.

We propose an approach inspired on the computation of shape from shading (SHAPIRO; STOCKMAN, 2001). The human body is assumed to be constituted by curved surfaces (e.g. arms, back, faces, etc.) and, in the same way humans see, digital images present a

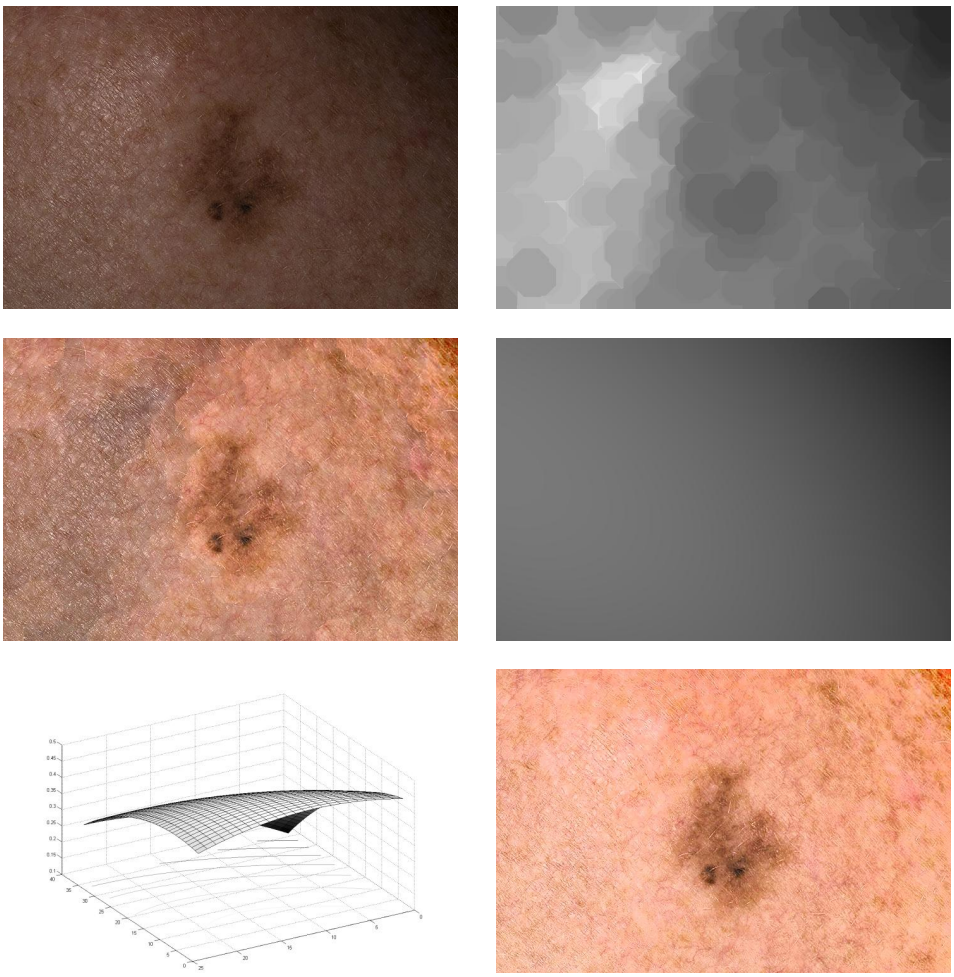

Figure 3.1: Shading attenuation in a pigmented skin lesion image : (a) Input image; (b) Morphological closing of Value channel by a disk (radius = 30 pixels); (c) Unsatisfactory shading attenuation after replacing the Value channel by $R(x, y)$, as suggested by Soille (SOILLE, 1999); (d) Local illumination based on the obtained quadric function; (e) 3D plot of the obtained quadric function; (f) Shading attenuation by using our approach.

smoothly darkening surface as one that is turning away from the view direction. However, instead of using this illumination variation to model the surface shape, we use this information to relight the image itself.

Let S be a set of known skin pixels (more details in Section 3.1.3). We use this pixel set to adjust the following quadric function $z(x, y)$:

$$z(x, y) = P_1x^2 + P_2y^2 + P_3xy + P_4x + P_5y + P_6, \tag{3.2}$$

where the six quadric function parameters P_i $(i = 1, ..., 6)$ are chosen to minimize the error ϵ:

$$\epsilon = \sum_{j=1}^{N} \left[V(S_{j,x}, S_{j,y}) - z(S_{j,x}, S_{j,y})\right]^2, \tag{3.3}$$

where, N is the number of pixels in the set S, and $S_{j,x}$ and $S_{j,y}$ are the x and y coordinates of the jth element of the set S, respectively.

Calculating the quadric function $z(x, y)$ for each image spatial location (x, y), we have an estimate z of the local illumination intensity in the image V. Replacing $M(x, y)$ by $z(x, y)$, and $I(x, y)$ by $V(x, y)$ in Eq. 3.1, we obtain the image $R(x, y)$ normalized with respect to the local illumination estimate $z(x, y)$. The final step is to replace the original Value channel by this new Value channel, and convert the image from the HSV color space to the original RGB color space. As a consequence of this image relighting, the shading effects are significantly attenuated in the color image. Figs. 3.1(d)-(e) illustrate the results obtained with our shading attenuation method.

3.1.3 Experimental Results and Discussion

As mentioned before, our method is initialized by a set of pixels S known to be associated with skin areas. In this section, we discuss how to select this set of pixels S in two typical applications of human skin color image analysis, namely, the segmentation of pigmented skin lesions and of faces in color images. Our goal is to show that our shading attenuation approach helps in the analysis of these images, making the processing steps simpler.

3.1.3.1 Pigmented Skin Lesion Segmentation in Color Images

In this application, the focus is in the image skin area that contains the lesion. As consequence, during image acquisition, the lesion is captured in the central portion of the image, and is surrounded by healthy skin. Therefore, we assume the four image corners to contain healthy skin. This assumption is common in dermatological imaging, and also has been made by other researchers in this field (CELEBI et al., 2008) (MELLI; GRANA; CUCCHIARA, 2006). Therefore, we use 20×20 pixel sets around each image corner, and determine S as the union of these 1600 pixels (i.e. the four pixel sets).

Many methods have been proposed for analyzing pigmented skin lesions in dermoscopic images (MAGLOGIANNIS; DOUKAS, 2009). However, dermoscopes are tools used by experts, and there are practical situations where a non-specialist wishes to have a qualified opinion about a suspect pigmented skin lesion, but only standard camera imaging is available on site (i.e., telemedicine applications). In the following discussion we focus in this situation that justifies the use of telemedicine and standard camera imaging. To illustrate the effectiveness of our method, we compare the segmentation results in pigmented skin lesions with and without the application of our method. Usually, pigmented

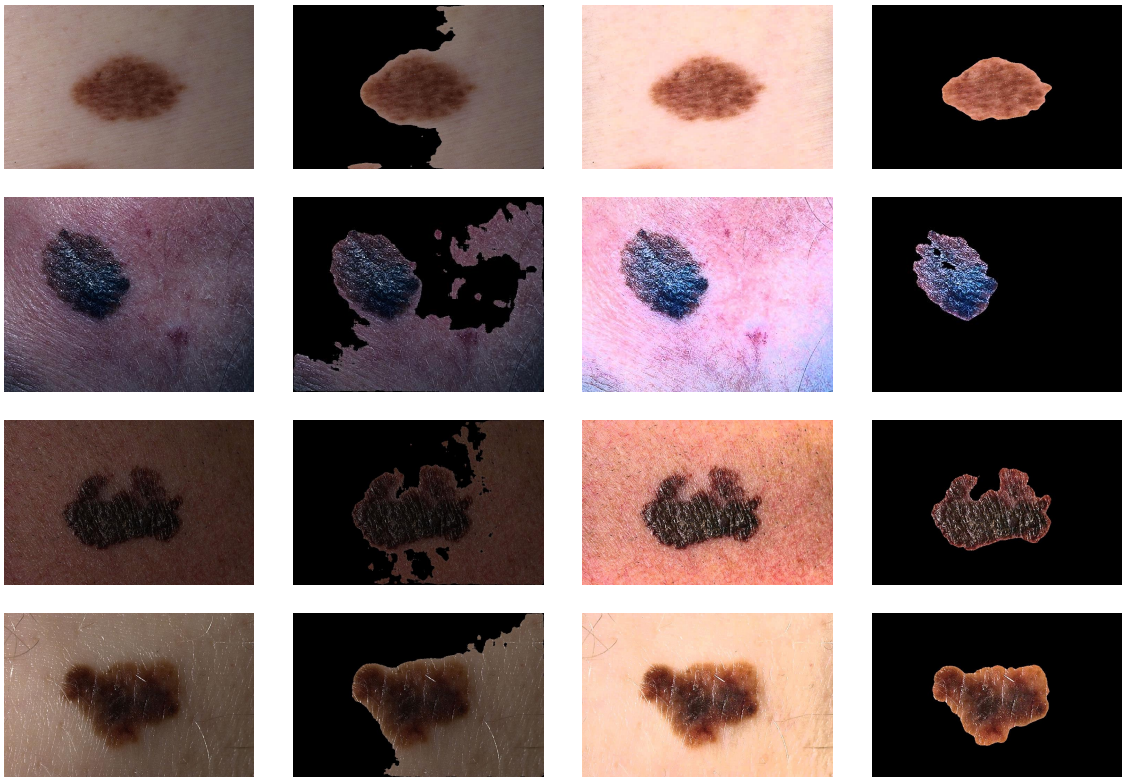

Figure 3.2: Examples of pigmented skin lesion segmentation. In the first and second columns, the original images and their respective segmentation results. The third and fourth columns show the resulting images after the application of our shading attenuation method, and the respective segmentation results.

skin lesions correspond to local darker skin discolorations. The segmentation method used is a well known thresholding procedure based on Otsu's method (OTSU, 1979). This algorithm assumes two pixel classes, usually background and foreground pixels (specifically in our case, healthy and unhealthy skin pixels), and searches exhaustively for the threshold th that maximizes the inter-class variance $\sigma_b^2(th)$:

$$\sigma_b^2(th) = \omega_1(th)\omega_2(th)\left[\mu_1(th) - \mu_2(th)\right]^2, \tag{3.4}$$

where, ω_i are the a priori probabilities of the two classes separated by the threshold th, and μ_i are the class means. To segment the input color images, we determine a threshold th for each one of the RGB channels, and establish a pixel as lesion if at least two of its RGB values are lower than the computed thresholds. At the end, we eliminate possible small segmented regions filtering the thresholding result with a 15×15 median filter.

In Fig. 3.2, we present some pigmented skin image segmentation examples. These pigmented skin lesions images are publicly available in the Dermnet dataset (DERMNET, 2010). Although our segmentation method is very simple, the application of the shading attenuation method increases its efficacy. In this way, the feature extraction and the classification procedure (typically the next steps in pre-diagnosis systems) have higher probability to produce accurate results.

Our method may fail in some situations, as illustrated in Fig. 3.3. The situations illustrated in Fig. 3.3 are: (a) our method is adequate to model and attenuate the global illumination variation (which changes slowly), but tends to have limited effect on local

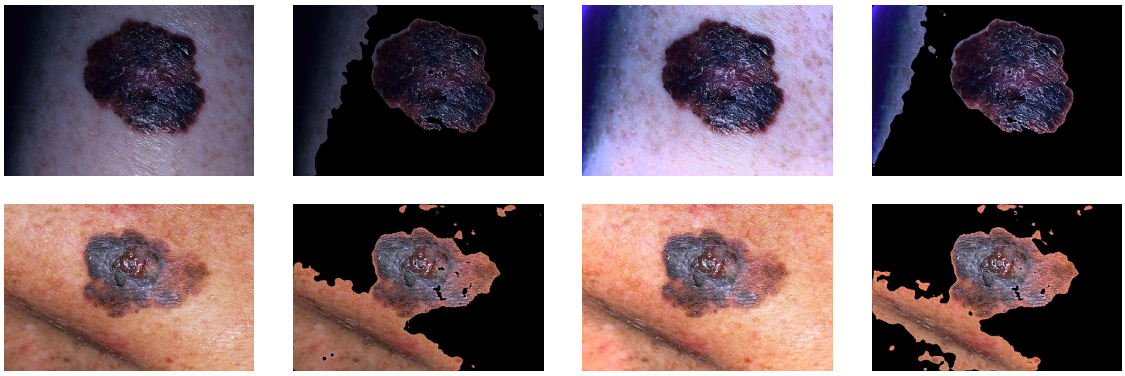

Figure 3.3: Illustrations of cases where our shading attenuation method tends to fail, such as cast shadows (first line) and surface shapes not well modeled by quadric functions (second line). The first and second columns show the original images and their respective segmentation results. The third and fourth columns show the resulting images after the application of our shading attenuation method, and the respective segmentation results.

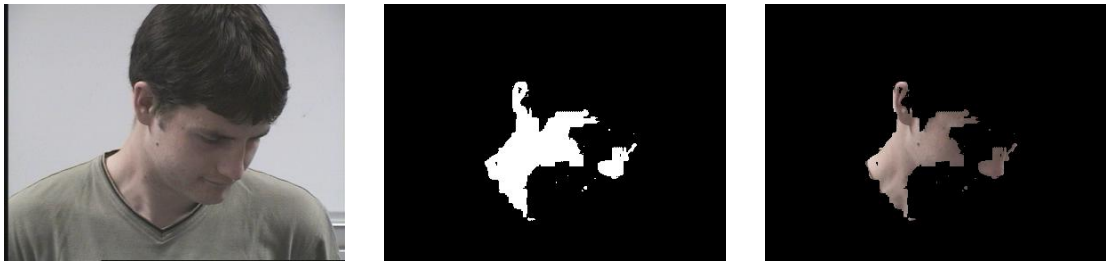

(a) (b) (c)

Figure 3.4: Illustration of skin pixels localization using Eq. 3.5 : (a) Input image; (b) Binary mask; and (c) adjacent pixels identified as human skin.

cast shadows; and (b) our approach tends to fail on surface shapes that are not locally smooth, since the quadric function is not able to capture the local illumination variation in this case. In such cases, the segmentation method may confuse healthy and unhealthy skin areas. Possibly, better results could be achieved in such cases by acquiring the images in a way that surface shapes are smoother and illumination varies slowly across the scene.

3.1.3.2 Face Segmentation in Color Images

A face can be found in virtually any image location. In this case, the selection of the initialization pixel set S it is not as trivial as in the pigmented skin lesion segmentation problem. Therefore, we obtain the initialization pixel set S based on previously known color information (VASSILI; SAZONOV; ANDREEVA, 2003). A pixel is considered to be associated to a skin region in an RGB face image if :

$$\begin{aligned} R > 95 \ \wedge \ G > 40 \ \wedge \ B > 20 \ \wedge \\ \max(R,G,B) - \min(R,G,B) > 15 \ \wedge \\ |R-G| > 15 \ \wedge \ R > G \ \wedge \ R > B, \end{aligned} \tag{3.5}$$

where, $\wedge$ denotes the logical operator *and*.

In Fig. 3.4, we present an example of the initialization skin pixels set S obtained with Eq. 3.5. Although this criterion to determine pixels associated to skin color is used often (VASSILI; SAZONOV; ANDREEVA, 2003), it can be very imprecise in practical situations, specially when there is image shading. However, its use here is justifiable since all we need is a set of adjacent image pixels with skin color (i.e. likely to be located in skin regions) to initialize our error minimization operation (see Eqs. 3.2 and 3.3), and erroneous pixels should not influence significantly the final result.

Once S has been determined, the shading effects in the face image can be attenuated. To demonstrate the efficacy of our method in this application, we show the face segmentations with, and without, shading attenuation using a known Bayes Classifier for the pixels based on their corrected colors (VASSILI; SAZONOV; ANDREEVA, 2003). A pixel is considered skin if:

$$\frac{P(c|skin)}{P(c|\neg skin)} > \theta, \tag{3.6}$$

$$where \;\; \theta = \kappa \times \frac{1 - P(skin)}{P(skin)}. \tag{3.7}$$

In Eq. 3.6, the a priori probability $P(skin)$ is set to 0.5, since we use the same number of samples for each class (i.e. 12800 skin pixels and 12800 non-skin pixels). The constant κ also is set to 0.5, increasing the chance of a pixel be classified as skin, and $P(c|skin)$ and $P(c|\neg skin)$ are modeled by Gaussian joint probability density functions, defined as:

$$P = \frac{1}{2\pi|\sum|^{1/2}} \times e^{-\frac{1}{2}(c-\mu)^T \sum^{-1}(c-\mu)}, \tag{3.8}$$

where, c is the color vector of the tested pixel, and μ and $\sum$ are the distribution parameters (i.e., the mean vector and covariance matrix, respectively) estimated based on the training set of each class (skin and non-skin).

Figs. 3.5 and 3.6 illustrate some face segmentation examples. These face images are publicly available in the Pointing'04 dataset (GOURIER; HALL; CROWLEY, 2004). The images in Fig. 3.5 show four different persons, with different physical characteristics and different poses (i.e. angles between their view direction and the light source), resulting in different shading effects. Clearly, the skin pixels, and consequently the face, is better segmented after we apply our shading attenuation method in all these different situations. In Fig. 3.6, we present four examples of the same person, just varying her head pose (the angle between her view direction and the light source). It shall be observed that even when the face is evenly illuminated, the face is better segmented after using our shading attenuation method. However, inaccuracies may occur near facial features partially occluded by cast shadows (e.g. near the nose and the chin). Based on these results, it should be expected that algorithms that extract facial features (e.g., eyes, mouth and nose) would perform their tasks more effectively, which helps in typical color image analysis problems such as head pose estimation or face recognition.

3.1.4 Conclusions

This paper presented a method for attenuating the shading effects in human skin images. Our preliminary experimental results indicate that the proposed method is applicable in at least two typical color image analysis problems where human skin imaging is of central importance. In the case of pigmented skin lesion segmentation, our shading

Figure 3.5: Face segmentation examples. In the first and second columns are shown the original images, and their respective segmentation results. In the third and fourth columns, are shown images after the application of our shading attenuation method, and their respective segmentation results.

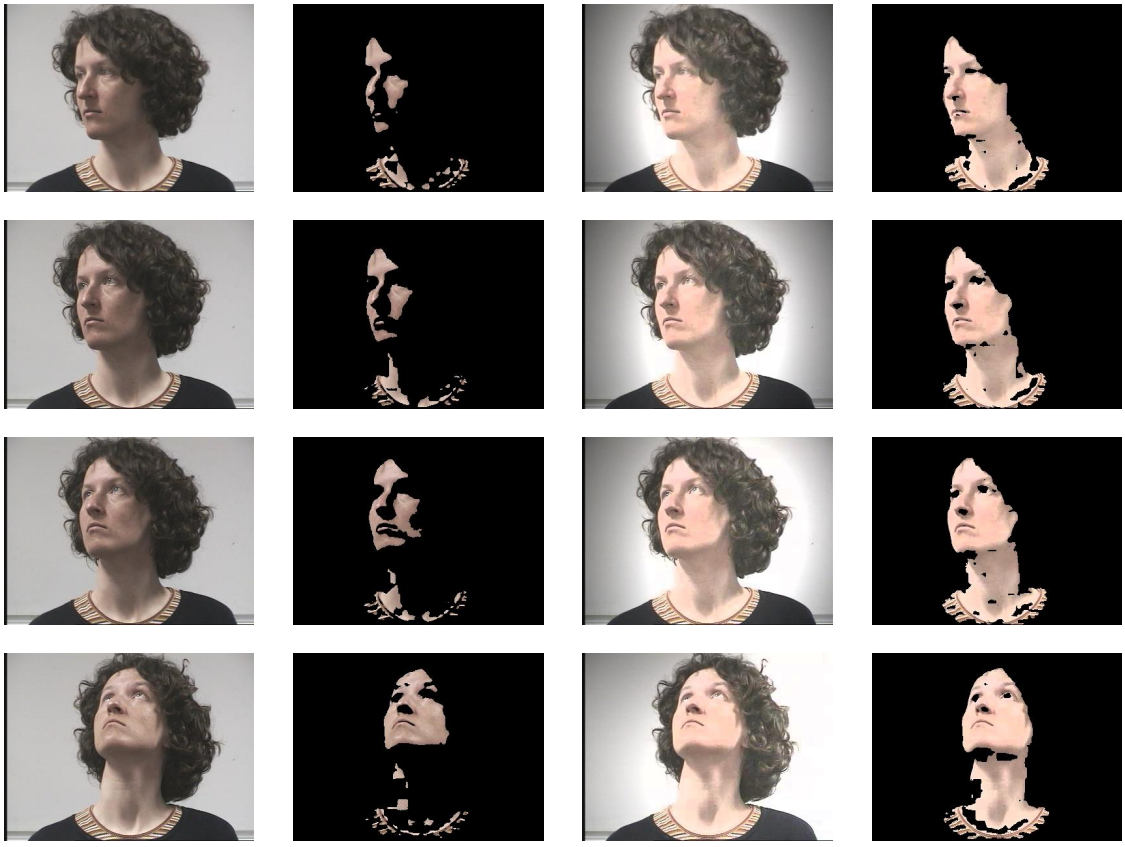

Figure 3.6: Face segmentation examples for the same person varying its head pose. In the first and second columns are shown the original images, and their respective segmentation results. In the third and fourth columns, are shown images after the application of our shading attenuation method, and their respective segmentation results.

attenuation method helps improving the lesion detection, and, hopefully, contributes for the early identification of skin cancer cases. We also studied the application of our shading attenuation method as a tool to increase the robustness of face segmentation, and our experiments suggest that potentially it can contribute to improve the efficiency of head pose estimation and facial recognition systems. We plan to further develop our approach using more complex quadric functions, and do a more extensive testing of our shading attenuation method in typical color imaging applications.

References

See the unified bibliography of the thesis.

3.2 Automated Prescreening of Pigmented Skin Lesions Using Standard Cameras

Cavalcanti, Pablo G. ; Scharcanski, Jacob . Automated prescreening of pigmented skin lesions using standard cameras. Computerized Medical Imaging and Graphics, v. 35, p. 481-491, 2011.

Abstract

This paper describes a new method for classifying pigmented skin lesions as benign or malignant. The skin lesion images are acquired with standard cameras, and our method can be used in telemedicine by non-specialists. Each acquired image undergoes a sequence of processing steps, namely: (1) preprocessing, where shading effects are attenuated; (2) segmentation, where a 3-channel image representation is generated and later used to distinguish between lesion and healthy skin areas; (3) feature extraction, where a quantitative representation for the lesion area is generated; and (4) lesion classification, producing an estimate if the lesion is benign or malignant (melanoma). Our method was tested on two publicly available datasets of pigmented skin lesion images. The preliminary experimental results are promising, and suggest that our method can achieve a classification accuracy of 96.71%, which is significantly better than the accuracy of comparable methods available in the literature.

3.2.1 Introduction

Pigmented skin lesions include both, benign and malignant forms. Melanoma is a kind of malignant pigmented skin lesion, and currently is among the most dangerous existing cancers, resulting in about 10000 deaths from the 40000 to 50000 diagnosed cases per year, just in United States of America (Melanoma Research Project, 2010). According to World Health Organization (World Health Organization, 2011), about 132000 melanoma cases occur globally each year. Benign pigmented skin lesions are called moles, or nevi. However, differentiating benign and malignant lesions can be challenging. For example, there are nevi known as Clark Nevi (also referred as Dysplastic or Atypical Nevi) that present similar characteristics to melanomas (The Skin Site, 2010). It is consensual that the early diagnosis of malignant skin lesions (melanomas) is essential for the patient prognosis.

Recently, telemedicine techniques have been studied as a resource to obtain an early diagnosis of skin lesions. Besides the fact that dermatology probably is the most visual specialty in medicine, the teledermatology consultation brings some benefits, like easier access to health care and faster clinical results (MASSONE et al., 2008). Comparing the physical examination (face-to-face diagnosis) with the remote diagnosis, experiments indicate that teledermatology is effective and reliable (WHITED, 2006). Teledermatology can benefit from image prescreening to help identify potentially malignant cases in their early stages.

To help distinguishing between benign and malignant cases, dermatologists often analyze each lesion with a dermoscope, which is a noninvasive tool that facilitates the evaluation of submacroscopic morphologic and vascular structures. As can be seen in Fig. 3.7, dermoscopy enables the generation of images with constant illumination, different texture patterns, and characteristics that are not measurable in standard camera images, such as lesion area and perimeter. In this way, many research groups developed digital dermoscopy image analysis schemes to help in skin lesion diagnosis (BLUM; ZALAUDEK;

ARGENZIANO, 2008).

In an attempt to prescreen/classify dermoscopy images, Celebi et al. (CELEBI et al., 2007) achieved 92.34% and 93.33% of specificity and sensitivity, respectively, using a JSEG-based segmentation algorithm and Support Vector Machines in the classification. More recently, Iyatomi et al. (IYATOMI et al., 2008) proposed, to the best of our knowledge, the first publicly accessible system ("Dermatologist-like"), based on a region growing segmentation method and an Artificial Neural Network classifier. The user can upload an image at their website '`http://dermoscopy.k.hosei.ac.jp`' and obtain a prescreening result. This system achieved a sensitivity of 85.9% and a specificity of 86.0% for a set of 1258 dermoscopy images (IYATOMI et al., 2008).

Despite the importance of these efforts, a disadvantage of these methods is that they require dermoscopy images, and dermoscopes are not common among non-specialists. Moreover, studies indicate that dermoscopy images do not increase diagnosis accuracy in early stages (SKVARA et al., 2005). So, with the proposal to facilitate the access to health care, also have been developed teledermatology systems making use of images acquired with standard cameras. In this way, patients do not need to go physically to a hospital or a clinic for a preliminary evaluation (even in benign cases), specially in remote areas. However, due to the already mentioned different visible characteristics in standard camera images and in dermoscopy images (see Fig. 3.7), these systems require different segmentation methods and feature extraction techniques.

A recent approach proposed by Alcon et al. (ALCON et al., 2009) is an easy-to-use melanoma prescreening system based on standard camera images. A skin lesion photograph is provided as an input, and its prescreening is automatically produced, using segmentation and classification algorithms. However, often the acquired images contain artifacts, such as uneven illumination, which causes difficulties in the lesion segmentation stage, therefore their system initially corrects the image background. Afterwards, 55 features are extracted and the ABCD rule (Asymmetry, Border irregularity, Color variation and Differential structures) is employed to classify the lesion image as benign or malignant. Their system reaches an accuracy of 86.64% in its best performance (ALCON et al., 2009).

This paper describes a new melanoma prescreening method using standard camera images and new techniques to improve the processing and analysis of such images, which

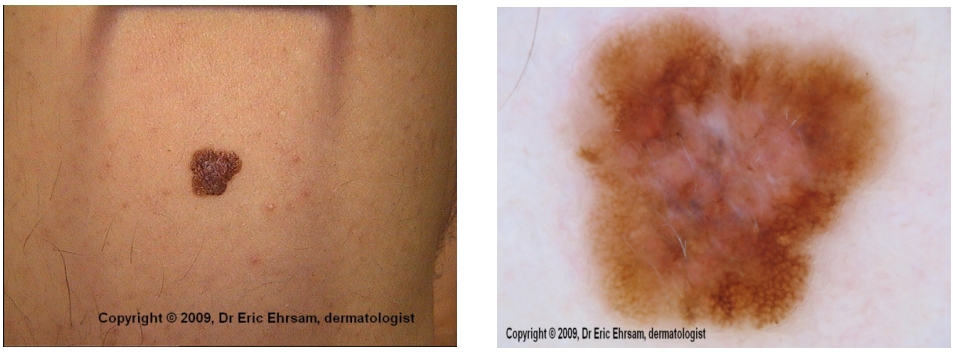

(a) (b)

Figure 3.7: Different visualizations of the same lesion using : (a) macroscopic image acquired with a standard camera; and (b) microscopic image acquired with a dermoscope. (Courtesy of Eric Ehrsam, MD (EHRSAM, 2010)).

were designed to be used remotely by non-specialists. In our experiments (see Section 3.2.6), we used 220 images obtained from two websites, with no special care in image acquisition or postprocessing. Each one of these images was submitted to a sequence of processing steps, namely: (1) preprocessing, where shading effects are attenuated by a new method proposed in this paper, as described in Section 3.2.2; (2) segmentation, where a new 3-channel image representation is proposed and used to discriminate between lesion and healthy skin areas, as described in Section 3.2.3; (3) feature extraction, where a quantitative lesion description containing new features extracted from our 3-channel representation is generated, as described in Section 3.2.4; and finally (4) lesion classification provides a lesion pre-diagnosis (using a hybrid classifier proposed in this paper); this processing step was designed to reduce the number of false negatives in the classification of skin lesion images as malignant or non-malignant, as described in Section 3.2.5. In Section 3.2.6, experimental results are presented, and Section 3.2.7 we present our conclusions.

3.2.2 Preprocessing

As mentioned before, the input image may be affected by illumination artifacts, and if used directly in the segmentation process, shading and lesion regions could be confused. Therefore, shading is attenuated in the input image before the image segmentation.

We start by converting the input image $\bar{I}_i^c$ ($\bar{I}_i^c(x,y) \subset [0,1], i = 1,2,3$) from the original RGB color space to the HSV color space $\bar{I}_i^{HSV}$ (SMITH, 1978). This is justified by the better shading visibility in the Value channel, and also by the simplicity of dealing with monochromatic images.

Let $\bar{I}$ be the Value normalized channel of $\bar{I}_i^{HSV}$ such that $\bar{I} \subset [0,1]$. To correct the illumination in monochromatic images, Soille (SOILLE, 1999) suggests a simple operation:

$$R(x,y) = \bar{I}(x,y) \;/\; M(x,y), \tag{3.9}$$

where, R is the resultant image, $M = \bar{I} \bullet s$ is the morphological closing of $\bar{I}$ by the structuring element s, and (x,y) represents a pixel in each one of these images. The main idea is that the closing operator is used to estimate the local illumination. However, Fig. 3.8(b)-(c) show that the result is unsatisfactory, specially because it is difficult to determine an efficient structuring element for each image location, and the illumination is inhomogeneous.

Skin lesions often occur on curved surfaces (e.g. arms, hands, faces, etc.), and the illumination changes locally due to the surface curvature, generating shading effects. A smoothly darkening surface is presented as one that is turning away from the view direction, and we use use shape from shading concepts (SHAPIRO; STOCKMAN, 2001) to relight the image(instead of determining the illumination variation by a morphological closing operation).

In order to avoid confusing lesion darkness with image shading, we capture shading information based on pixels known to be in healthy skin areas. The imaged lesion is supposed to be in the central part of the image, and healthy skin areas are expected in the four image corners (CELEBI et al., 2008; MELLI; GRANA; CUCCHIARA, 2006). Therefore, we use a set of 400 pixels around each image corner (20×20 square), and find the pixel set S as the union of the four pixel sets, corresponding to 1600 pixels associated to healthy skin. The pixel set S is used to determine the local illumination intensity $z(x,y)$, adjusting the following quadric function:

$$z(x,y) = P_1x^2 + P_2y^2 + P_3xy + P_4x + P_5y + P_6, \tag{3.10}$$

where, the six quadric function parameters P_i $(i = 1, ..., 6)$ are chosen to minimize the error ϵ:

$$\epsilon = \sum_{j=1}^{\#S} \left[\bar{I}(S_{j,x}, S_{j,y}) - z(S_{j,x}, S_{j,y})\right]^2, \tag{3.11}$$

where, $S_{j,x}$ and $S_{j,y}$ are the x and y coordinates of the jth element of the set S, respectively.

Calculating the quadric function $z(x,y)$ for each image spatial location (x,y), we estimate the local illumination intensity in the image $\bar{I}$. Next, we replace $M(x,y)$ by $z(x,y)$ in Eq. 3.9 obtaining $R(x,y)$, significantly attenuating the shading effects (see Eq. 3.9). An illustration of the results obtained with this shading attenuation approach is in Figs. 3.8(d)-(e). Next, the monochromatic image $R(x,y)$ is used to replace the Value channel $\bar{I}_3^{HSV}$ in $\bar{I}_i^{HSV}$, and we obtain the shading corrected image by converting $\bar{I}_i^{HSV}$ to the original RGB color space $\bar{I}_i^c$, $i = 1, 2, 3$. Examples of color images with shading correction are presented in Figs. 3.8(f) and 3.10(b).

3.2.3 Skin Lesion Segmentation

The skin lesion segmentation helps identify the skin lesion area and its rim in monochromatic ((ALCON et al., 2009)(GOMEZ et al., 2008)), or in color ((CELEBI et al., 2007)(IYATOMI et al., 2008)) images. However, skin artifacts such as hair and freckles can be confused with lesions, and affect negatively the prescreening process (e.g., feature extraction and classification). Moreover, segmentation techniques developed for dermoscopy images consider texture and color patterns that are usually not visible in standard camera images. Thus, we propose a new method to segment the lesions in these kind of image.

To minimize segmentation errors, we create a new representation $\bar{I}_i^N (i = 1, 2, 3)$ from the original image to improve the discrimination between healthy and unhealthy skin regions. This new three-channel image contains normalized values ($\bar{I}_i^N(x,y) \subset [0,1]$, $\bar{I}_i^N(x,y) \in \mathbb{R}$), improving its sensitivity to local changes in the lesion characteristics (such as texture, darkness and color variation). The features represented more explicitly in each one of these three channels are discussed in Subsections 3.2.3.1, 3.2.3.2 and 3.2.3.3.

3.2.3.1 Texture Variability Information Channel ($\bar{I}_1^N$)

Skin lesions often have more local textural variability than healthy skin areas, and that can be used to differentiate healthy and unhealthy skin regions. In order to capture the textural variation information we start by computing the normalized Luminance image $\bar{L}$, defined as:

$$\bar{L}(x,y) = \frac{\sum_{i=1}^{3} \bar{I}_i^c(x,y)}{3}, \tag{3.12}$$

where, $\bar{I}_i^c(x,y)$ is a pixel (x,y) of the normalized RGB color image. We quantify the textural variability in $\bar{L}(x,y)$ by computing $\tau(x,y,\sigma)$:

$$\tau(x,y,\sigma) = \frac{\bar{L}(x,y)}{S(x,y,\sigma)} - \bar{L}(x,y), \tag{3.13}$$

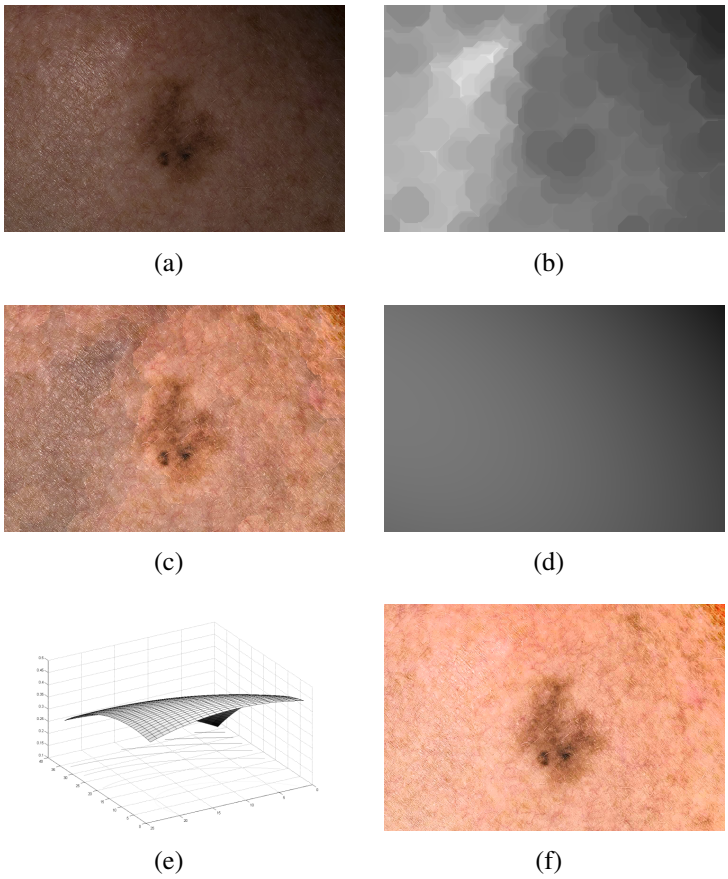

(a) (b) (c) (d) (e) (f)

Figure 3.8: Shading attenuation. (a) Input image; (b) Closing of the input Value channel by a disk of 30 pixels radius; (c) Result obtained by replacing the input Value channel by its normalized version using the image (b) - an unsatisfactory shading attenuation result; (d) Obtained quadric model using the corners of the input Value channel; (e) Obtained quadric model in 3D; (f) Result obtained by substituting the normalized version of the input Value channel using the quadric model in (d) as a normalization factor - our shading attenuation result.

where, $S(x, y, \sigma) = \bar{L}(x, y) * G(\sigma)$ (i.e., the Luminance image $\bar{L}$ is smoothed by a Gaussian filter with standard deviation σ). Re-arranging terms in Eq. 3.13, we obtain:

$$\begin{aligned} \tau(x, y, \sigma) &= \frac{\bar{L}(x, y) - \bar{L}(x, y) \cdot S(x, y, \sigma)}{S(x, y, \sigma)} \\ &= \frac{\bar{L}(x, y)\,(1 - S(x, y, \sigma)))}{S(x, y, \sigma)} \\ &= \bar{L}(x, y)\frac{\widetilde{S}(x, y, \sigma)}{S(x, y, \sigma)}, \end{aligned} \tag{3.14}$$

where, $\widetilde{S}(x, y, \sigma)$ represents the complement of the $\bar{L}$ smoothed image (i.e. $\widetilde{S}(x, y, \sigma) = 1 - S(x, y, \sigma)$). If an image region is dark (i.e. has low pixel intensities, as often occurs in skin lesions), $\widetilde{S}(x, y, \sigma) > S(x, y, \sigma)$ and the ratio $\widetilde{S}(x, y, \sigma)/S(x, y, \sigma)$ tends to increase, and the local region and textural information are emphasized; if the region is bright (e.g. as in healthy skin regions), $\widetilde{S}(x, y, \sigma) < S(x, y, \sigma)$ and the ratio $\widetilde{S}(x, y, \sigma)/S(x, y, \sigma)$

tends to decrease, and the local region and its textural information are de-emphasized. This process is illustrated in Fig. 3.9 for the image shown in Fig. 3.8(f). Usually, most pixels in $\bar{L}$ correspond to healthy skin (i.e., have higher intensity values). However, most skin pixels have lower values (closer to zero) in $\tau(x, y, \sigma)$, while skin lesion pixels tend to be brighter, as the histogram peaks show in Figs. 3.9(b)-(c). [1]

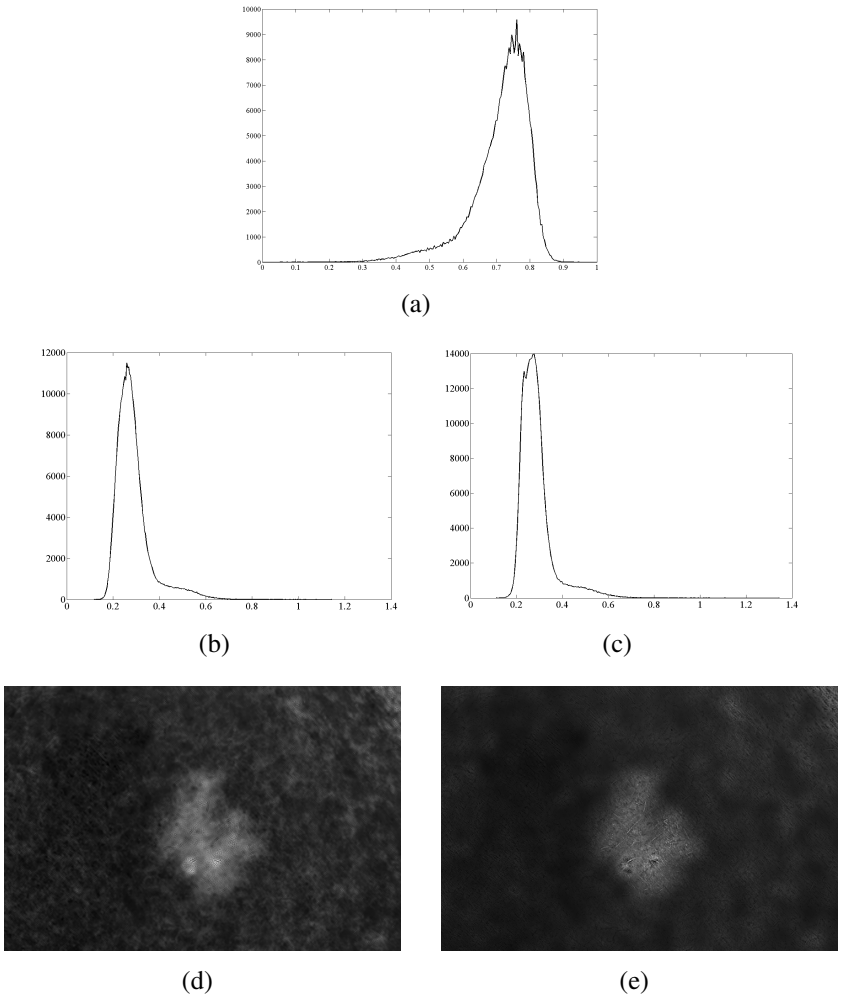

(a) (b) (c) (d) (e)

Figure 3.9: Texture variability quantification $\tau(x, y, \sigma)$. (a) Histogram of the Luminance image $\bar{L}$ of Fig. 3.8(f); (b)-(c) show the histograms of $\bar{L}(x, y)\widetilde{S}(x, y, \sigma)/S(x, y, \sigma)$ (see Eq. 3.14) using $\sigma = 1$ and $\sigma = \frac{43}{7}$, respectively; (d)-(e) show the image instances associated to the histograms shown in Figs. (b)-(c).

A single Gaussian filter may not be sufficient to capture the textural variability of different types of lesions in generic images. So, we calculate $\tau(x, y, \sigma)$ for different σ values[2] $\{\sigma_1, \sigma_2, ..., \sigma_N\}$, and select at each pixel the maximum $\tau(x, y, \sigma)$ among all scales σ:

$$T(x, y) = \max_{\sigma}[\tau(x, y, \sigma)],\ \sigma \in \{\sigma_1, \sigma_2, ..., \sigma_N\}\,. \tag{3.15}$$

Finally, the texture variation channel T is normalized, obtaining $\bar{I}_1^N$:

[1] The higher histogram peaks are associated to healthy skin regions in Figs. 3.9(b)-(c).

[2] In our experiments, we used $\sigma = 1, \frac{11}{7}, \frac{15}{7}, ..., \frac{43}{7}$, and filter window sizes of $7\sigma \times 7\sigma$.

$$\bar{I}_1^N(x,y) = (T(x,y) - \min(T)) \,/\, (\max(T) - \min(T)). \tag{3.16}$$

The procedure described in Eq. 3.15 was designed to capture the local texture variability information at each pixel (x, y), in lesions of different sizes, shapes and texture patterns.

3.2.3.2 Local Skin Darkness Information Channel ($\bar{I}_2^N$)

We use the complement of $\bar{I}_2^N = 1 - \bar{I}_1^c(x,y)$ (i.e. the normalized Red channel) to represent the local skin darkness. Healthy skin tends to be reddish, so brighter $\bar{I}_1^c(x,y)$ pixels occur in healthy skin regions and darker $\bar{I}_1^c(x,y)$ pixels often occur in lesion areas. Therefore, $\bar{I}_2^N$ tends to present lower intensity values in healthy skin areas and higher intensity values in lesion areas. Local skin darkness information $\bar{I}_2^N$ is used to reinforce the local texture variability information in $\bar{I}_1^N$, and to help discriminate between healthy and unhealthy skin areas in the input image.

3.2.3.3 Color Information Channel ($\bar{I}_3^N$)

Usually, healthy and unhealthy skin present different color distributions. To reduce the problem dimensionality and represent the color information in a single channel, we use principal component analysis (PCA) (GOMEZ et al., 2008) (CELEBI et al., 2009).

Let $\vec{c}$ be the color of the normalized RGB image pixel (x, y), where $\vec{c} = [\bar{I}_1^c(x,y)$ $\bar{I}_2^c(x,y)$ $\bar{I}_3^c(x,y)]$. The PCA finds vectors $\vec{u_1}, \vec{u_2}, \vec{u_3}$ (arranged as columns of a matrix U) in such a way that $\vec{c'} = U^{\mathrm{T}}\vec{c}$ maximizes the spread (variance) of the RGB data. We want a single channel representation for the color variability information, and then only use the vector $\vec{u_1}$ pointing along the direction that maximizes the variance of $\vec{c_1'}$ (i.e. first component of $\vec{c'}$). Assuming centered RGB data $\sum \vec{c_i} = 0$ (the color data is centered by subtracting the mean data vector $\vec{\mu}$, i.e. $\vec{c} = \vec{c} - \vec{\mu}$), $\vec{u_1}$ is the vector that maximizes $\sum_{i=1}^{3} (\vec{u_1}^{\mathrm{T}}\vec{c_i})^2$. In other words, the image colors (3D points) projected along the $\vec{u_1}$ direction (1D scalars) are as spread as possible.

To make sure that lesion pixels have higher values than healthy skin pixels, as in the channels $\bar{I}_1^N$ and $\bar{I}_2^N$, we use the PCA property described next. Since the color data is centered around the mean vector $\vec{\mu}$, and the most frequent colors are predominant in the computation of $\vec{\mu}$, the most frequent colors are projected along the $\vec{u_1}$ direction closer to the origin of the space spanned by $\vec{c_i}$. Since the healthy skin pixels usually are more frequent, their corresponding projection magnitudes $\left|\vec{c_1'}\right| \approx 0$, however the lesion pixels projection magnitudes $\left|\vec{c_1'}\right|$ tend to be larger than zero. Therefore, we represent the color variability information C by the magnitudes of $\vec{c_1'}$ (i.e., $C = \left|\vec{c_1'}\right|$).

In the sequence, we normalize the C values at each pixel (x, y) and create the $\bar{I}_3^N$ channel:

$$\bar{I}_3^N(x,y) = (C(x,y) - \min(C)) \,/\, (\max(C) - \min(C)). \tag{3.17}$$

Also, we reduce the noise in $\bar{I}_3^N$ by filtering it with a 5×5 median filter. The color information channel $\bar{I}_3^N$ helps to discriminate between healthy and unhealthy skin pixels, while emphasizing skin lesion regions.

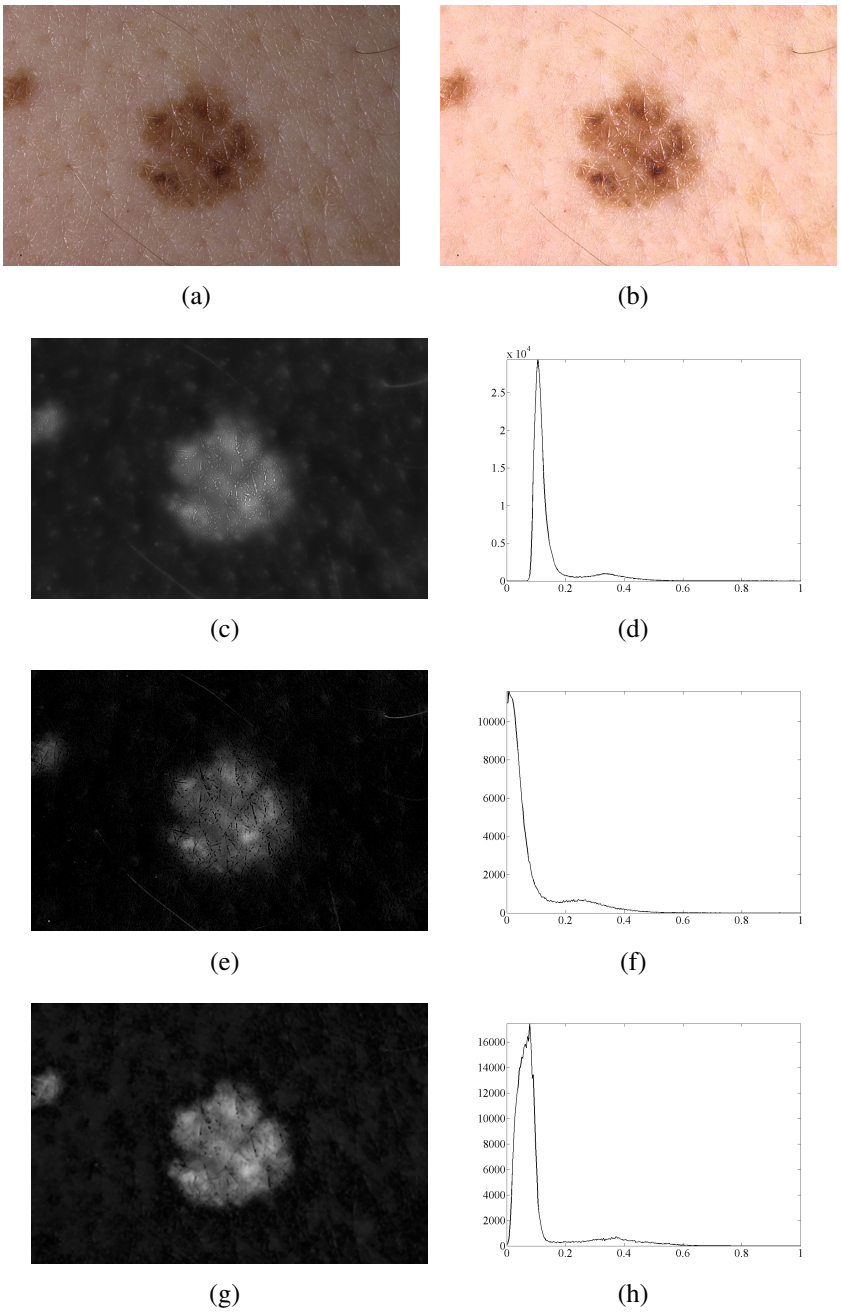

Figure 3.10: Illustration of our 3-channel image representation $\bar{I}_i^N$ and the data bimodal distribution in each channel (in higher peaks of histograms the healthy skin pixels, and in lower peaks the lesion pixels). (a) Input image; (b) Input image after shading attenuation; (c)-(d) show the computed texture variability information channel ($\bar{I}_1^N$) and its histogram; (e)-(f) show the computed local skin darkness information channel ($\bar{I}_2^N$) and its histogram; (g)-(h) show the computed color information channel ($\bar{I}_3^N$) and its histogram.

3.2.3.4 Skin Lesion Area Segmentation and Rim Detection

In all three channels of the input image representation $\bar{I}_i^N(x, y)$, lesion pixels tend to have higher intensities than healthy skin pixels. Fig. 3.10 illustrates these images, and show that their histograms typically are bimodal (lesion/healthy skin). Therefore, we segment the lesion areas using a method inspired on the Otsu's thresholding method (OTSU,

1979). This algorithm assumes two pixel classes, usually healthy and unhealthy skin pixels, and searches exhaustively for the threshold th that minimizes the total intra-class variance $\sigma_w^2(th)$, defined as the weighted sum of variances of the two classes:

$$\sigma_w^2(th) = \omega_1(th)\sigma_1^2(th) + \omega_2(th)\sigma_2^2(th), \tag{3.18}$$

where ω_i are the a priori probabilities of the two classes separated by the threshold th, and σ_i^2 are their intra-class variances. Minimize the intra-class variance is equivalent to maximize the inter-class variance $\sigma_b^2(th)$:

$$\begin{aligned} \sigma_b^2(th) &= \sigma^2 - \sigma_w^2(th) \\ &= \omega_1(th)\omega_2(th)\left[\mu_1(th) - \mu_2(th)\right]^2, \end{aligned} \tag{3.19}$$

where σ^2 is the image pixels variance, and μ_i are the class means. Therefore, we obtain three thresholds th_i by maximizing $\sigma_b^2(th)$ for each $\bar{I}_i^N$ channel, creating a lesion segmentation mask ϕ. A pixel (x, y) is in a lesion region (i.e., $\phi(x, y) = 1$) if its value is higher than the threshold th_i in at least two of the three channels:

$$\phi(x,y) = \begin{cases} 1 & , \text{ if } (\bar{I}_1^N(x,y) > th_1 \ \wedge \ \bar{I}_2^N(x,y) > th_2), \\ 1 & , \text{ if } (\bar{I}_2^N(x,y) > th_2 \ \wedge \ \bar{I}_3^N(x,y) > th_3), \\ 1 & , \text{ if } (\bar{I}_1^N(x,y) > th_1 \ \wedge \ \bar{I}_3^N(x,y) > th_3), \\ 0 & , \text{ otherwise.} \end{cases} \tag{3.20}$$

Although Ganster et al. (GANSTER et al., 2001) stated that a OR-combination of binary masks gives the best segmentation results, we noted in our experiments that this approach tends to merge lesion and healthy skin areas (e.g., areas where the skin have color or texture similar to lesions). In our approach, majority voting is used to eliminate most of these incorrectly segmented areas. After this operation, the remaining skin artifacts (such as freckles and hair) are eliminated more easily. These artifacts usually are form isolated regions that differ in area and perimeter from skin lesions, since lesions often have larger areas and boundaries that are more irregular. Therefore, we compute the perimeter and the area of all thresholded connected pixel sets (i.e. where $\phi(x, y) == 1$), and then partition this set of regions in two clusters. All regions in the cluster with smaller areas (in average) are eliminated, and we set their mask pixels to $\phi(x, y) = 0$.

After eliminating artifacts, the lesion regions still can contain rim imperfections (caused by hair or noise), which we filter out by applying a 5×5 median filter to $\phi(x, y)$. Figs. 3.11(a)-(c) illustrate the lesion area segmentation steps (i.e., thresholding, artifacts elimination and filtering) for the image shown in Fig. 3.10(b). As it can be seen, there can be holes inside the lesion segment, but their pixels are not considered further in the next image processing steps of our method. These holes often occur because of image specularities, and these pixels are not processed further because they could introduce noise in the lesion feature extraction process. The obtained lesion rim (i.e., the outer perimeter of the lesion) is shown in Fig. 3.11(d) superimposed to the original image (Fig. 3.10(b)).

3.2.4 Feature Extraction

We extract a set of image features to distinguish between benign and malignant skin lesions. Given the segmentation results, we compute the local characteristics from the lesion areas according to the ABCD rule (NACHBAR et al., 1994). The ABCD acronym

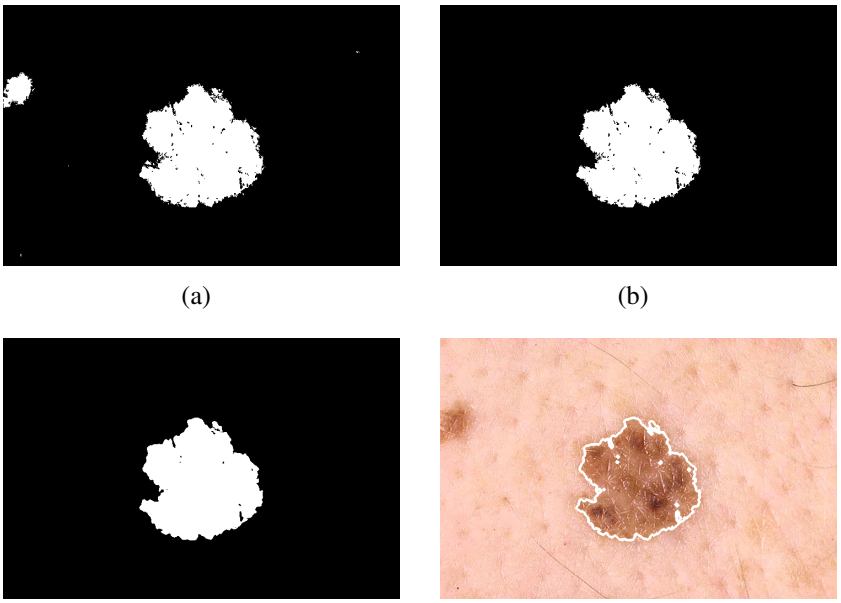

(a) (b) (c) (d)

Figure 3.11: Segmentation process for the image shown in Fig. 3.10(b). (a)-(c) show the lesion segmentation masks $\phi(x, y)$, respectively, after thresholding, artifacts elimination and filtering; (d) shows the computed lesion rim (dilated by two pixels for better visualization) superimposed to the original image.

refers to the four criteria used in this rule, namely: Asymmetry, Border irregularity, Color variation and Differential structures.

The ABCD rule in important dermatology, and most dermatological prescreening systems rely on some scheme for quantifying the four criteria of the ABCD rule. In this work, we propose some image features to quantify this rule, and combine them with other features that have already been proposed in the literature.

3.2.4.1 Features Used for Lesion Asymmetry Characterization

The goal of the proposed features is to quantify the lesion shape, in special the asymmetry of the lesion in relation to the principal axes. The major axis L_1 of the lesion is aligned with its longest diameter, passing through its center; the minor axis L_2 is orthogonal to L_1 and also passes through the shape center. The selected features are:

f_1: Solidity: the ratio between the lesion area (A) and its convex hull area (ALCON et al., 2009);

f_2: Extent: the ratio between the lesion area and its bounding box area (ALCON et al., 2009);

f_3: Equivalent diameter: $4A/(L_1\pi)$ (ALCON et al., 2009; CELEBI et al., 2007);

f_4: Circularity: $4\pi A/(L_1 p)$, where p is the lesion perimeter (ALCON et al., 2009);

f_5: The ratio between the principal axes (L_2/L_1) (ALCON et al., 2009; CELEBI et al., 2007);

f_6: The ratio between sides of the lesion bounding box (ALCON et al., 2009);

f_7**:** The ratio between the lesion perimeter p and its area A (FIKRLE; PIZINGER, 2007);

f_8**:** $(B_1 - B_2)/A$, where, B_1 and B_2 are the areas in each side of axis L_1;

f_9**:** Similar to f_8, but makes use of the shorter axis L_2;

f_{10}**:** $B1/B2$ with respect to the axis L_1;

f_{11}**:** Similar to f_{10}, but makes use of the shorter axis L_2.

3.2.4.2 Features Used for Lesion Boundary Irregularity Characterization

The boundary sharpness is quantified by the magnitude of the gradient $\left|\vec{\nabla}\bar{I}_i^N\right|$ at each pixel using the Sobel operator. However, instead of using pixels only at the lesion rim, we analyze pixels in an extended (dilated) rim[3] (ALCON et al., 2009). Consequently, lesions that have a smooth boundary (usually nevi) are better characterized. Also, the lesion boundary dilation makes the boundary representation more robust to the inaccuracies of the segmentation process. To characterize the lesion boundary irregularity, we use the following features:

f_{12}-f_{14}**:** Average gradient magnitude of the pixels in the lesion extended rim (ALCON et al., 2009), in each one of the three I_i^N channels;

f_{15}-f_{17}**:** Variance of the gradient magnitude of the pixels in the lesion extended rim (ALCON et al., 2009), in each one of the three I_i^N channels;

The lesion rim irregularity is characterized in the ABCD rule by dividing the rim in 8 symmetric regions (NACHBAR et al., 1994). In addition to the two principal axes L_1 and L_2, we rotate these orthogonal axes by 45 degrees and obtain two additional axes. Therefore, 8 symmetric regions $R = 1, ..., 8$ are obtained. For each channel $\bar{I}_i^N$, the average gradient magnitudes of the extended rim pixels $\mu_{R,i}(R = 1, ..., 8)$ are computed. Therefore, we calculate 6 more features:

f_{18}-f_{20}**:** Average of the 8 $\mu_{R,i}$ values in each one of the three $\bar{I}_i^N$ channels;

f_{21}-f_{23}**:** Variance of the 8 $\mu_{R,i}$ values in each one of the three $\bar{I}_i^N$ channels;

3.2.4.3 Features Used for Lesion Color Variation Characterization

Their goal is to quantify the color variation in the lesion. The selected features are:

f_{24}-f_{27}**:** Maximum, minimum, mean and variance of the pixels intensities inside the lesion segment in the color variation channel $\bar{I}_3^N$;

f_{28}-f_{39}**:** Maximum, minimum, mean and variance of the pixels intensities inside the lesion segment in each one of three original $\bar{I}_i^c$ channels;

f_{40}-f_{42}**:** Ratios between mean values of the tree original $\bar{I}_i^c$ channels: $mean(\bar{I}_1^c)/mean(\bar{I}_2^c)$, $mean(\bar{I}_1^c)/mean(\bar{I}_3^c)$ and $mean(\bar{I}_2^c)/mean(\bar{I}_3^c)$, considering only pixels inside the lesion segment.

[3]The rim is dilated by 2 pixels, producing a 5 pixels wide region centered at the lesion rim, as suggested in (ALCON et al., 2009).

Table 3.1: Six Possible Colors of a Lesion on the RGB Color Space.

Color	Red	Green	Blue
White	1	1	1
Red	0.8	0.2	0.2
Light Brown	0.6	0.4	0
Dark Brown	0.2	0	0
Blue-Gray	0.2	0.6	0.6
Black	0	0	0

Physicians usually identify six distinct hues in skin lesions: white, red, light and dark brown, blue-gray, and black (NACHBAR et al., 1994). Lesions containing more of these hues are more likely to be malignant. We quantify the lesion color variability by computing the occurrence of these typical hues within a lesion segment. Given a pixel in the lesion segment, we find the the nearest reference color (associated with a typical hue, see Table 3.1 (ALCON et al., 2009)) by the Euclidean distance to the pixel color in $\bar{I}_i^c$. A hue occurrence counter is created, one cell per typical hue. For each lesion pixel, the nearest typical hue counter is increased by 1. Finally, typical hues counters are normalized/divided by the lesion area A, and generate the 6 additional features f_{43}-f_{48}.

3.2.4.4 Features Used for Lesion Differential Structures Characterization

The lesion differential structures refer to submacroscopic morphologic and vascular structures only visible in dermoscope images, as already mentioned in Section 3.2.1. However, we measure differences between benign and malignant lesions using texture features found in macroscopic images. We extract the 4 features f_{49}-f_{52}, namely the maximum, minimum, mean and variance of the pixels intensities inside the lesion segment to represent the textural variation in the channel $\bar{I}_1^N$.

3.2.5 Lesion Classification

After segmenting a lesion segment, and extracting 52 features $f_1 - f_{52}$ (see Section 3.2.4), we can discriminate a benign from a malignant pigmented skin lesion by classification. We present in the following subsections our classification scheme.

3.2.5.1 Feature Normalization

The extracted features may generate values in different ranges. However, often classifiers are more efficient if these feature values are normalized (produce values within a specified range). We normalize feature values with the z-score transformation (AKSOY; HARALICK, 2000):

$$Z_{i,j} = \frac{((v_{i,j} - \mu_j)/(3\sigma_j) + 1)}{2}, \tag{3.21}$$

where, $v_{i,j}$ is the value of the jth feature of the ith sample (image), μ_j and σ_j are the mean and standard deviation of the jth feature, respectively. After the z-score transformation, most of the $Z_{i,j}$ values are in the [0,1] range. The out-of-range values are saturated to either 0 or 1.

3.2.5.2 Training Samples Selection

We applied holdout validation for the training samples selection. That is, half of the samples in each class (benign or malignant) are randomly selected for the training set. Additionally, since we have a relatively small image data set (with N samples described in Section 3.2.6), new samples are added to the training set using the Smoothed Bootstrap Resampling method[4] to obtain η samples, where $\eta > N$. In our experiments, we used zero mean Gaussian noise with $\sigma = 0.1$, obtained $\eta/2$ samples for each class, and η was set to the interval [1000, 5000] (see Section 3.2.6).

3.2.5.3 Classification

Two classifiers are used in our experiments (see Section 3.2.6), namely the K-Nearest Neighbors (KNN) classifier, and the KNN followed by a Decision Tree classifier (KNN-DT). Detailed information about these classifiers are presented next.

- KNN: classifies samples based on K closest matches in feature space (ALPAYDIN, 2004). We use K=1 and each sample is assigned to its nearest neighbor class by using the Euclidean distance. Although this classifier have already been used in previous approaches (DREISEITL et al., 2001; BURRONI et al., 2004), it never resulted in a high accuracy as with our proposed features (see our experiments in Section 3.2.6).

- KNN-DT: uses the KNN classifier (above) and a set of Bayes' Classifiers, each specialized on a feature sub-set identified by the Classification and Regression Tree (CART) method (BREIMAN et al., 1984). It is used to reduce false negatives by confirming each case pre-diagnosed as benign by KNN.

The Bayes' Classifier (ALPAYDIN, 2004) assigns a sample to the most likely class, given the extracted features. Let L be a set of features [f_1,...,f_{52}] extracted from a lesion, $D = 1$ the malignant class and $D = 0$ the benign class. Using Bayes' rule, the posterior probability $P(D|L)$ of assigning a lesion to class D can be written as:

$$P(D|L) = \frac{P(L|D)P(D)}{P(L)}, \tag{3.22}$$

where, $P(L|D)$ is the conditional probability of finding a lesion with these L characteristics in class D, $P(D)$ is the prior probability of class D, and $P(L)$ is the evidence and it represents the probability that these features L are seen.

Since the evidence $P(L)$ will be the same for both classes ($D = 0$ and $D = 1$), we discard this term from Eq. 3.22, and the lesion can be considered malignant if $P(D = 1|L) > P(D = 0|L)$, or if:

$$P(L|D = 1)P(D = 1) > P(L|D = 0)P(D = 0), \tag{3.23}$$

and it is benign otherwise. Given adequate estimates of $P(L|D)$ and $P(D)$, the Bayes' Classifier is optimal since it minimizes the probability classification error. Considering that we have a relatively large feature set to discriminate reliably between nevi and melanomas (see the proposed 52 features in Section 3.2.4), the estimate of the joint feature

[4]The Smoothed Bootstrap Resampling method is used when data is not sufficient to guarantee statistical significance. Specifically, original samples are randomly selected, and new ones are created adding a small amount of zero-centered noise to their feature values, enlarging the data set (YOUNG, 1990).

probability in Eq. 3.23 can be quite challenging without simplifications. Also, $P(D)$ is not known (to the best of our knowledge). Therefore, we use CART to simplify these estimates. Recall that we are interested in minimizing false malignant (melanoma) negatives by confirming the pre-diagnosis of nevi cases. Therefore, we simplify our problem by choosing feature sub-sets (with less than 52 features) that independently can provide evidences for a melanoma pre-diagnosis, and the CART method is used to find these feature sub-sets.

The CART algorithm (BREIMAN et al., 1984) iteratively builds a decision tree by creating feature space binary partitions (i.e. selecting linear feature discriminants), in a way that sample class homogeneity is increased within each sub-space partition at each decision tree construction step (ALPAYDIN, 2004). Each step introduces a new tree node (i.e. a new linear discriminant), and the CART algorithm converges when the sample class partitions are maximally homogeneous, given an error criterion. Traversing the decision tree from its root to a given leaf node, implies in testing a set of feature thresholds to reach a decision at the leaf node (i.e. reaching a sub-space partition that allows to decide whether to assign a sample to the nevus or the melanoma class). However, in this work we are only interested in the CART feature sub-sets involved in the leaf nodes decisions (not in the linear discriminants produced by CART). Particularly, we are interested in the feature sub-sets found by the CART algorithm that are used to decide the assignment of a sample to the melanoma class (i.e. in the decision tree leaf nodes involved in class $D = 1$ assignment decisions). In this way, we reduce the computation of joint feature probability estimates in Eq. 3.23 to a set of feature sub-sets joint probability estimates, specially those used in $D = 1$ class assignments. In this way, we reduce the chance of occurring false negatives in the final classification results (i.e. malignant samples incorrectly assigned to the non-malignant class can be re-assigned correctly). In our experiments, we obtained feature sub-sets of 4 to 19 features.

The following steps are used to estimate the joint probabilities in these feature sub-sets:

- Considering all training samples assigned to classes $D = 1$ and $D = 0$ at all the leaf nodes, we estimate the prior class probabilities $P(D = 1)$ and $P(D = 0)$ as the percentages of the training set samples that are assigned to each one of these two classes. In fact, we calculate the probability that a sample is assigned to the class $D = 1$ at any decision tree leaf node $P(D = 1)$, and then estimate the prior probability that a sample is assigned to class $D = 0$ as $P(D = 0) = 1 - P(D = 1)$;

- Considering all training samples, we compute the feature joint probabilities $P(L_l|D = 0)$ and $P(L_l|D = 1)$ for each decision tree branch leading to a $D = 1$ class assignment, where L_l denotes all the features involved in the decision sequence at a given tree branch, from the tree root to its leaf node l. The feature joint probabilities $P(L_l|D = 0)$ and $P(L_l|D = 1)$ are multivariate distributions, that we model using an unsupervised modified EM (Expectation-Maximization) algorithm to identify finite Gaussian mixture models based on the method proposed by Figueiredo and Jain (FIGUEIREDO; JAIN, 2002);

- We calculate $P(L_l|D = 0)$ and $P(L_l|D = 1)$ for each leaf node l. Given the estimated $P(D = 1)$ and $P(D = 0)$, we use Eq. 3.23 at each leaf node with a $D = 1$ class assignment decision, and decide if the sample in question should be re-assigned to the class $D = 1$. For all leaf nodes l involving a $D = 1$ assignment

decision, we increment the counter $C_{D=1}$ if $P(L_l|D=1)P(D=1) > P(L_l|D=0)P(D=0)$;

- Finally, a sample previously assigned to the nevus class ($D=0$) is re-assigned to the melanoma class ($D=1$) is $C_{D=1} \geq \frac{N_l}{2}$, i.e. when at least half of the leaf node tests ($D=1$) indicate that the sample should be assigned to class $D=1$.

3.2.6 Experimental Results and Discussion

We used in our experiments two datasets of images publicly available: (a) the dataset used by Alcon et al.(ALCON et al., 2009), with 152 images from the Dermnet dataset (DERMNET, 2010) (i.e., 45 benign Clark Nevi and 107 Melanomas); and (b) an extended dataset, which was built by adding 68 extra images from the DermQuest dataset (DERMQUEST.COM, 2010) (37 Clark Nevi and 31 Melanomas), constituting a total of 82 Clark Nevi and 138 Melanomas in this extended dataset. The idea is to test our method in a dataset used in the literature (i.e. the Alcon et al. dataset (ALCON et al., 2009)), and then confirm our method performance in a larger dataset. Most of the 220 images are 720 pixels in width (in landscape format), or have a height of 720 pixels in portrait format. All images with different dimensions were resized to have 720 pixels in width or in height, preserving the aspect ratio. Recall that the imaged lesion is supposed to be in the central part of the image, and healthy skin areas are expected in the four image corners (CELEBI et al., 2008; MELLI; GRANA; CUCCHIARA, 2006).

We compare our results with two representative methods reported in the literature: (a) Iyatomi et al. (known as the "Dermatologist-like") (IYATOMI et al., 2008), which is a publicly accessible teldermatology prescreening system; and (b) Alcon et al. (ALCON et al., 2009), which is the most recent approach dealing with standard camera images. In order to compare our results with these methods results, we use sensitivity ($Sens = (TP)/(TP+FN)$), specificity ($Spec = (TN)/(TN+FP)$) and accuracy ($Acc = (TP+TN)/(TP+FP+TN+FN)$), where TP, TN, FP and FN are, respectively, the number of True Positive, True Negative, False Positive and False Negative cases. Specifically, $Sens$ indicates the percentage of malignant (melanomas) cases correctly classified; $Spec$ indicates the percentage of benign (nevi) cases correctly classified; and, Acc indicates the overall percentage of cases correctly classified.

In order to illustrate that dermoscopy image prescreening methods actually are not adequate to analyze skin lesion images acquired with standard cameras, we tested the Alcon et al. dataset (ALCON et al., 2009) in the web-based "Dermatologist-like" system (IYATOMI et al., 2008). From 107 melanomas images, 68 images were classified correctly, 26 images were mistakenly classified as nevi, 11 images were classified as 'suspicious' and 2 images resulted in no classification at all (probably because of segmentation errors). From 45 nevi images, 31 images were classified correctly, 7 images were incorrectly classified as melanomas, 5 images were classified as 'suspicious', and 2 resulted in no classification at all. In this experiment, we obtained a sensitivity of 75.24%, a specificity of 80%, and an accuracy of 75.66%, considering 'suspicious' as a correct classification, and the situation of 'no classification' as an incorrect classification.

Alcon et al.(ALCON et al., 2009) tested two classification procedures in their dataset. They tested the Correlation-Based Feature Selector (CFS) with Logistic Model Tree (LMT), and classified correctly 101 images in 107 melanoma images, and 31 images in 45 nevi images, achieving a sensitivity of 94.39%, a specificity of 68.89% and an accuracy of 86.84%. The second classifier, CFS with Adaboost and LMT, classified correctly 96 im-

ages in 107 melanoma images, and 34 images in 45 nevi images, achieving a sensitivity of 89.72%, a specificity of 75.56% and an accuracy of 85.53%.

Table 3.2 presents the results obtained using the Alcon et al. dataset, our lesion image representation with the KNN classifier and the KNN-DT classifier, and different values of η (the number of Bootstrap samples). Although results are very similar, $\eta = 5000$ was selected since it maximize the accuracy (i.e., generated the highest number of correctly classified images). It is important to observe that the KNN-DT classifier reduced the number of false negatives compared with the KNN classifier (i.e., correctly classified a higher number of melanoma/malignant images), obtaining higher sensitivity values.

Table 3.2: Our experiments results using the Alcon et al. dataset.

Method	η	Melanomas class (total:107)	Nevi class (total:45)	$Sens$	$Spec$	Acc
KNN	1000	103	43	96.26%	95.55%	96.05%
KNN-DT	1000	103	43	96.26%	95.55%	96.05%
KNN	2000	102	43	95.32%	95.55%	95.39%
KNN-DT	2000	103	43	96.26%	95.55%	96.05%
KNN	3000	100	43	93.45%	95.55%	94.07%
KNN-DT	3000	102	43	95.32%	95.55%	95.39%
KNN	4000	101	43	94.39%	95.55%	94.73%
KNN-DT	4000	103	43	96.26%	95.55%	96.05%
KNN	5000	103	44	96.26%	97.77%	96.71%
KNN-DT	5000	103	44	96.26%	97.77%	96.71%

Table 3.3 compares the results obtained by our approach ($\eta = 5000$), and by the methods proposed by Iyatomi et al.(IYATOMI et al., 2008) and Alcon et al.(ALCON et al., 2009), using the Alcon et al. dataset. As it can be seen that for the same dataset, the results obtained by our approach improved significantly on the results obtained by the other methods (IYATOMI et al., 2008) (ALCON et al., 2009). The highest accuracy result achieved by these methods was 86.84% (ALCON et al., 2009), and our approach achieved 96.71%, meaning that nearly 10% more cases of the dataset were classified correctly by our method. Also, our approach classified correctly more melanoma images, increasing the sensitivity values, and more nevi images, increasing significantly the specificity values.

As already mentioned, we also realized experiments with an extended image dataset. The KNN classifier classified correctly 131 images in 138 melanoma images, and classified correctly 75 images in 82 nevi images, achieving a sensitivity of 94.92%, a specificity of 91.46% and an accuracy of 93.63%. The KNN-DT classifier classified correctly 133 images in 138 melanoma images, and classified correctly the same 75 images in 82 nevi

Table 3.3: Comparison of results.

Method	$Sens$	$Spec$	Acc
"Dermatologist-like" (IYATOMI et al., 2008)	75.24%	80%	75.66%
CFS with LMT (ALCON et al., 2009)	94.39%	68.89%	86.84%
CFS with Adaboost and LMT (ALCON et al., 2009)	89.72%	75.56%	85.53%
KNN	96.26%	97.78%	96.71%
KNN-DT	96.26%	97.78%	96.71%

images (as the KNN classifier), achieving a sensitivity of 96.37.%, a specificity of 91.46% and an accuracy of 94.54%. These results are shown in Table 3.4. Comparing with the results for the Alcon et al. dataset, we had a decrease in specificity. However, it is important to observe that our approach can correctly classify higher percentages of melanoma and nevi images, achieving better sensitivity, specificity and accuracy measures as compared to the other methods (IYATOMI et al., 2008)(ALCON et al., 2009) (see Table 3.3).

Table 3.4: Results to the extended dataset.

Method	Melanomas class (total:138)	Nevi class (total:82)	$Sens$	$Spec$	Acc
KNN	131	75	94.92%	91.46%	93.63%
KNN-DT	133	75	96.37%	91.46%	94.54%

As can be seen in Tables 3.2 and 3.4, the KNN-DT classifier is able to classify correctly a higher number of melanoma images, and consequently it can provide higher sensitivity values than the KNN classifier. However, the KNN classifier can be useful since it has lower computational cost.

3.2.7 Conclusion

This paper presented a method for classifying pigmented skin lesions as benign or malignant. It is assumed that the imaged lesion is located in the central part of the image, and healthy skin areas are expected in the four image corners (CELEBI et al., 2008; MELLI; GRANA; CUCCHIARA, 2006). Besides the lesion position in the image, no special care is required in the skin lesion image acquisition (e.g. dermoscopy is not used), and the images are acquired with standard cameras and standard illumination, making our method suitable for telemedicine applications. Since our skin lesion pre-screening method is automatic, it can be used by non-specialists.

In this paper, new techniques to improve the processing and analysis of skin images acquired with standard cameras were proposed, such as: (1) a new image data-driven shading attenuation methodology was introduced to improve the image preprocessing stage; (2) to facilitate the lesion rim detection, a new 3-channel image representation that maximizes the discrimination between the lesion and healthy skin regions was presented; (3) to better discriminate between malignant and non-malignant lesions, new features based in our new 3-channel representation were introduced; and, (4) a hybrid classifier was proposed to reduce the number of false negatives in the classification of skin lesion images as malignant or non-malignant.

The preliminary experimental results suggest that our approach potentially can achieve better classification results than comparable methods available in the literature (IYATOMI et al., 2008; ALCON et al., 2009). Using the dataset proposed in (ALCON et al., 2009), the highest accuracy result achieved by these other methods was 86.84% (ALCON et al., 2009), while our approach achieved 96.71% of accuracy. Also, our approach classified correctly more melanoma images, increasing the sensitivity values, and more nevi images, increasing the specificity values. Also, using an extended dataset, our method classified correctly 94.54% of the 220 images, with a sensitivity of 96.37% and a specificity of 91.46%. These experimental results are encouraging, and we plan to further develop our approach by improving the lesion segmentation and classification stages, and test more extensively our method in clinical trials.

References

See the unified bibliography of the thesis.

3.3 An ICA-Based Method for the Segmentation of Pigmented Skin Lesions in Macroscopic Images

Cavalcanti, P. G. ; Scharcanski, J. ; Persia, L. E. ; Milone, D. H. . An ICA-Based Method for the Segmentation of Pigmented Skin Lesions in Macroscopic Images. In: 33rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 11), 2011, Boston, EUA. Proceedings of the Int. Conf. of the IEEE EMBS, 2011.

Abstract

Segmentation is an important step in computer-aided diagnostic systems for pigmented skin lesions, since that a good definition of the lesion area and its boundary at the image is very important to distinguish benign from malignant cases. In this paper a new skin lesion segmentation method is proposed. This method uses Independent Component Analysis to locate skin lesions in the image, and this location information is further refined by a Level-set segmentation method. Our method was evaluated in 141 images and achieved an average segmentation error of 16.55%, lower than the results for comparable state-of-the-art methods proposed in literature.

3.3.1 Introduction

Pigmented skin lesions include both, benign and malignant forms. Just in United States of America occur about 10000 deaths per year from the 40000 to 50000 diagnosed cases of melanoma, a dangerous kind of malignant pigmented skin lesions. Early diagnosis is of fundamental importance to improve the patient prognosis, nevertheless discriminating benign from malignant skin lesions has been proven to be a challenging task.

To facilitate the diagnosis, physicians often use dermoscopy, which is a non-invasive technique that magnifies submacroscopic structures with the help of an optical lens (a dermoscope) and liquid immersion. According to Mayer (MAYER, 1997), the use of dermoscopy can increase the diagnosis sensitivity in 10-27% with respect to the clinical diagnosis. The early diagnosis of melanomas is very important for the patient prognosis, since most malignant skin lesion cases can be treated successfully in their early stages. However, even with the help of dermoscopy, differentiating malignant and benign lesions is a hard task. In fact, specialists affirm that in the early stages of the evolution of malignant lesions, dermoscopy may be useless as a tool to help the diagnosis (SKVARA et al., 2005).

Still considering early stage cases, there are practical situations where a non-specialist (e.g. a physician not trained on Dermatology) wishes to have a qualified opinion about a suspect skin lesion, but only standard camera imaging is available on site. In such situations, telemedicine is justifiable, and the non-specialist can capture an image of the suspect skin lesion and send it to an specialist, who can analyze it in higher detail. In this particular situation, a teledermatology consultation brings benefits, like the easier access to health care and faster clinical results (MASSONE et al., 2008). Besides, comparing the physical (face-to-face) patient diagnosis with the remote diagnosis by teledermatology, recent results suggest that teledermatology also tends to be effective and reliable (WHITED, 2006). Therefore, there is a growing interest in methods for diagnosing pigmented skin lesions remotely, and segmentation is important in this context since the lesion must be correctly located before it is analyzed and diagnosed.

Many segmentation methods have been proposed in the literature for dermoscopy im-

ages. Such techniques often are based on different strategies, namely: region-based, using mainly region-growing approaches (IYATOMI et al., 2008) (CELEBI et al., 2008); clustering algorithms, separating healthy and unhealthy pixels in homogeneous regions (GOMEZ et al., 2008); thresholding methods, computing values that can identify the lesion in histograms (GANSTER et al., 2001); and, active contours procedures, where the lesion borders are detected iteratively (SILVEIRA et al., 2009). In a literature review, Celebi et al. (CELEBI et al., 2009) reported that most of these methods are automatic and operate on color images, but usually can result on segmentation errors, specially closer to the lesion borders.

Unfortunately, macroscopic pigmented skin lesion (MPSL) image segmentation did not receive much attention in the literature. Recently, some experiments (TANG, 2009; PAROLIN; HERZER; JUNG, 2010) demonstrated that active contours-based techniques can achieve good results. However, these methods usually depend on good initialization, and in our case this is not easily attainable. MPSL images usually contain several artifacts (as uneven illumination condition, hair, freckles, etc.) that make the automatic segmentation of the lesion region more difficult, and manual initialization often becomes necessary. In this paper, we propose a method based on Independent Component Analysis (ICA), that detects the skin lesion in macroscopic images, and uses this approximation as an initialization of a Level-set method that detects the lesion boundaries more accurately.

In Section 3.3.2 we describe all the steps of this method. Then, in Section 3.3.3 we present our segmentation results for publicly available datasets. Finally, Section 3.3.4 presents conclusions and directions for future work.

3.3.2 Proposed Method for Skin Lesion Segmentation

As mentioned before, MPSL images usually contain artifacts that make the automatic lesion segmentation more difficult. Shading is a particularly challenging artifact, since shading areas often are darker than healthy skin and may be confused with lesion areas. To avoid confusing lesion and shading areas, the first step of our segmentation method is to attenuate the local shading effects. After that, we use ICA to separate the information contained in the image, usually healthy skin, lesion and possibly other artifacts. Given the ICA results we obtain an initial lesion localization, and then the lesion boundary is determined more accurately by using a Level-set method (CHAN; SANDBERG; VESE, 2000) and a few post-processing steps. All these steps are detailed in the following subsections.

3.3.2.1 Shading Attenuation

Cavalcanti et al. (CAVALCANTI; SCHARCANSKI; LOPES, 2010) (CAVALCANTI; SCHARCANSKI, 2011) recently proposed a method to significantly attenuate shading effects in MPSL images. The method assumes that images are acquired in a way that the lesion appears in the image center, and it does not touch the image outer borders. The first step of the method is to convert the image from the original RGB color space to the HSV color space, and retain the Value channel V. This is justified by the fact that this channel presents the higher visibility of the shading effects. We extract 20×20 pixels in each V corner and define S as the union of these four sets. This pixel set is used to adjust the following quadric function $z(x, y)$:

$$z(x, y) = P_1x^2 + P_2y^2 + P_3xy + P_4x + P_5y + P_6, \tag{3.24}$$

where the six quadric function parameters P_i $(i = 1, ..., 6)$ are chosen to minimize the

error ϵ:

$$\epsilon = \sum_{j=1}^{N_s} \left[V(S_{j,x}, S_{j,y}) - z(S_{j,x}, S_{j,y}) \right]^2 , \tag{3.25}$$

where, $S_{j,x}$ and $S_{j,y}$ are the x and y coordinates of the jth element of the set S, respectively, and N_s is the total number of pixels of the four corners ($N_s = 1600$ in our experiments).

Calculating the quadric function $z(x, y)$ for each image spatial location (x, y), we have an estimate $z(x, y)$ of the local illumination intensity in the image $V(x, y)$. Dividing the original $V(x, y)$ channel by $z(x, y)$, we obtain a new Value channel where the shading effects have been attenuated. The final step is to replace the original Value channel by this new Value channel, and convert the image from the HSV color space to the original RGB color space.

3.3.2.2 ICA-Based Lesion Localization

Independent Component Analysis (ICA) is a method to process multivariate data, producing projections that are as statistical independent as possible (HYVäRINEN; KARHUNEN; OJA, 2001). In this way, ICA searches for a data projections that maximize the degree of independence of such data projection. In this work, we use the FastICA algorithm (HYVÄRINEN, 1999) with a cost function that maximizes non-Gaussianity. To apply FastICA, we resize each $n \times m$ image channel to a $1 \times nm$ vector, by scanning the image top to bottom, left to right, and normalize color image triplet components to unity length [5], obtaining the $3 \times nm$ measure matrix X. Then, the FastICA algorithm searches for a separation matrix W such that the output matrix $Y = WX$ has row components that are as statistically independent as possible.

Applying the ICA to MPSL images three independent components are obtained, one corresponds mainly to the lesion area, another to the healthy skin, and the third component corresponds to noise and other artifacts. Nevertheless, there is an ordering indeterminacy inherent to the ICA method, and it is not possible to know in advance which component will show the lesion more clearly. However, due to the lesion variability, the histogram of the component that shows more clearly the lesion often has a non-Gaussian histogram (frequently multimodal). The histogram of the component showing predominantly noise and artifacts tends to be non-Gaussian, and the component that shows healthy skin more clearly tends to have a Gaussian histogram. Thus, we measure the non-Gaussianity of the ICA histogram components with differential entropy, i.e. $\mathcal{J}(\mathcal{X}) = |H(\mathcal{X}) - H(\mathcal{X}_g)|$, where $\mathcal{X}_g$ is a Gaussian distributed random variable with the same variance as $\mathcal{X}$. The component that produces the largest differential entropy is identified as the one containing the lesion information more clearly, and the smallest differential entropy component carries basically healthy skin.

After reordering the channels, we have the lesion region best represented in the first channel. Next, we normalize its values in the range $[0, 1]$ and use the Otsu's thresholding method (OTSU, 1979) to segment the skin lesion in this channel. This algorithm assumes two pixel classes, usually background and foreground pixels (specifically in our case, healthy and unhealthy skin pixels), and searches exhaustively for the threshold th that maximizes the inter-class variance $\sigma_b^2(th)$:

[5] The ICA assumes that samples are iid, and scanning order is irrelevant.

$$\sigma_b^2(th) = \omega_1(th)\omega_2(th)\left[\mu_1(th) - \mu_2(th)\right]^2, \tag{3.26}$$

where, ω_i are the a priori probabilities of the two classes separated by the threshold th, and μ_i are the class means. Given the ICA results, the lesion information can be emphasized (closer to value 1) or de-emphasized (closer to value 0) in this channel, and consequently the thresholded area may correspond to either, the lesion or the background (the only constrain is to maximize σ_b^2). To guarantee that we capture the lesion in the thresholded area, the corners pixels (used in the shading attenuation step, that are known to correspond to healthy skin) are tested to check if they are thresholded as '1's or '0's. If most corner pixels are thresholded as '1's, the thresholded area corresponds to healthy skin, and the logical complement is used to obtain the lesion localization mask. Since the thresholded first re-ordered ICA component now contains the lesion approximately, a morphological opening is performed on this image to better approximate the lesion boundary and eliminate residual artifacts (using a disk of 3 pixels of radius as the structuring element).

3.3.2.3 Lesion Boundary Detection

After obtaining the lesion localization mask (see Section 3.3.2.2), we determine the lesion boundary more precisely using the Chan-Vese Level-set method for vector-valued images (CHAN; SANDBERG; VESE, 2000). It assumes that the color image I_i is formed by two regions of approximately constant intensities c_1 and c_2, separated by a curve C. This lesion localization mask is used as an initialization indicating approximately the region to be segmented. Afterwards, the Level-set method iteratively tries to minimize the energy function $F(c_1, c_2, C)$ in the color image I_i (after shading effects have been attenuated):

$$\begin{aligned} F(c_1, c_2, C) = {} & \mu\, length(C) + \\ & \lambda_1 \int_{inside(C)} \frac{1}{3}\sum_{i=1}^{3} \left|I_i(x,y) - c_{1,i}\right|^2 dxdy + \\ & \lambda_2 \int_{outside(C)} \frac{1}{3}\sum_{i=1}^{3} \left|I_i(x,y) - c_{2,i}\right|^2 dxdy, \end{aligned}$$

where μ, λ_1 and λ_2 are weighting parameters (we used λ_1=λ_2=1, as suggested by the authors (CHAN; SANDBERG; VESE, 2000), and μ=0.2). Using the Level-set formulation, it is possible to minimize the energy function embedding the curve C, obtaining the zero level set $C(t) = \{(x,y)|\phi(t,x,y) = 0\}$ of a higher dimensional Level-set function $\phi(t,x,y)$. The evolution of $\phi(t,x,y)$ is given by the following motion Partial Differential Equation:

$$\frac{\partial \phi}{\partial t} = \delta_\epsilon(\phi) \left[\begin{array}{l} \mu\, div\left(\frac{\nabla\phi}{|\nabla\phi|}\right) - \\ \frac{1}{3}\sum_{i=1}^{3} \lambda_1 \left|I_i(x,y) - c_{1,i}\right|^2 + \\ \frac{1}{3}\sum_{i=1}^{3} \lambda_2 \left|I_i(x,y) - c_{2,i}\right|^2 \end{array} \right], \tag{3.27}$$

where $\delta_\epsilon(\phi)$ is the Dirac delta function, $c_{1,i}$ and $c_{2,i}$ are the averages inside and outside of the curve C in the i-th channel I_i, respectively.

It is possible that the final curve C contains regions beyond the lesion area. So, if the number of regions segmented by the Chan-Vese method is higher than one, local artifacts are eliminated. We compute the area and the perimeter of each segmented region, and cluster these values with K-Means, where K=2. The regions in the cluster with the smaller areas (in average) are eliminated as artifacts and the other regions are kept. The regions kept are hole filled to improve their connectivity, forming the final segmentation mask.

The final post-processing step is a morphological dilation (with a disk of 5 pixels of radius as the structuring element). As already observed (CELEBI et al., 2009) (GARNAVI et al., 2011), the hand labeled lesion ground truth created by specialists tends to be slightly larger than the result of our automatic segmentation, and this dilation operation helps to suppress this difference. Fig. 3.12 presents images that demonstrate all these steps, from the image shading attenuation until the obtained of lesion segmentation.

3.3.3 Experimental Results

In order to evaluate the performance of our proposed segmentation method, we used 141 macroscopic images of pigmented skin lesions from the Dermnet database (DERMNET, 2010). In these 141 images, 97 are malignant melanomas images, and 44 are Clark Nevi images (which is a benign lesion difficult to diagnose since it contains similar characteristics to melanomas). The segmentation error measure ϵ was chosen because it was also used by the authors of the four state-of-the-art segmentation approaches used in our comparisons (IYATOMI et al., 2008) (CELEBI et al., 2008) (GOMEZ et al., 2008) (CAVALCANTI; YARI; SCHARCANSKI, 2010):

$$\epsilon = \frac{S \oplus G}{G} \times 100\%, \tag{3.28}$$

where S is the result of segmentation by the method under test, G is the manual segmentation of the same lesion and $\oplus$ indicates the exclusive-OR, which gives the pixels for which S and G disagree.

We also selected four segmentation methods to compare with the results of our proposed method. In a recent paper, Celebi et al. (CELEBI et al., 2009) reviewed several different segmentation approaches for dermoscopy, and indicated two methods among the best performing published approaches: (1) Statistical Region Merging, proposed by Celebi et al. (CELEBI et al., 2008); and (2) Independent Histogram Pursuit, proposed by Gómez et al. (GOMEZ et al., 2008). Also recently, Iyatomi et al. (IYATOMI et al., 2008) proposed a method named "Dermatologist-like", which they claim to obtain segmentation results similar to those obtained by the Celebi et al.(CELEBI et al., 2008). We also included in our comparisons a method recently proposed by Cavalcanti et al. (CAVALCANTI; YARI; SCHARCANSKI, 2010), designed specifically for macroscopic images.

Analyzing the segmentation results, we observe that the methods proposed by Celebi et al. (CELEBI et al., 2008) and Gómez et al. (GOMEZ et al., 2008) confuse shading and lesion areas in several images. Probably because these methods have been proposed for dermoscopy images, which are not affected by shading. Also, we generated segmentation results for the 141 images with the three methods mentioned above, using our shading attenuation method as a pre-processing step. The segmentation error averages are presented in Table 3.5, as well as the percentages of the images in the database that generated segmentation errors lower than 5%, 10%, 20%, 30%, 40% and higher than 100%[6], respectively. As it can be seen, our proposed method achieved the best results

[6]The segmentation error is higher than 100% if the segmented area that is not part of the lesion ground

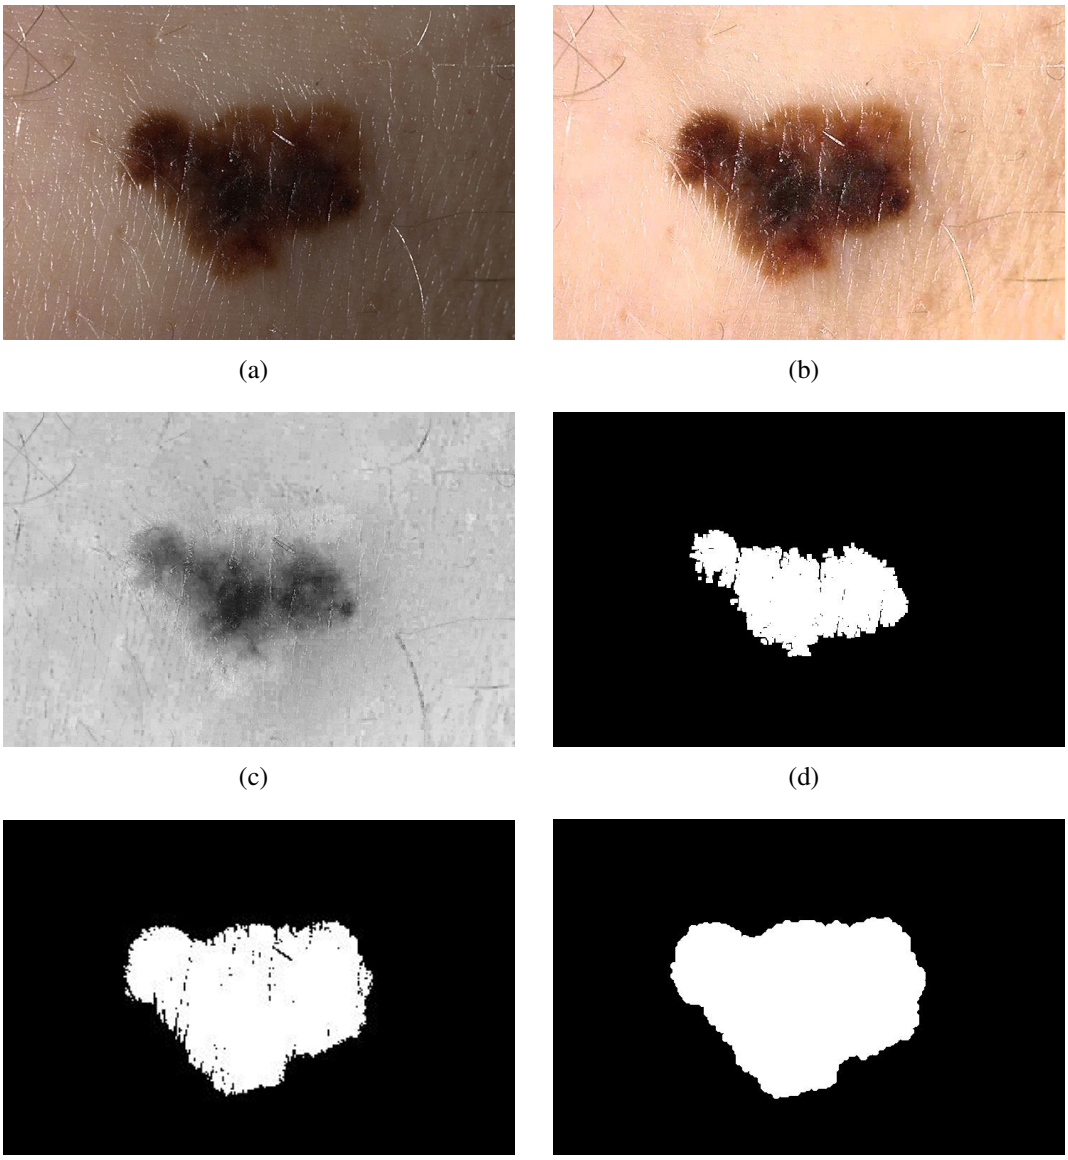

Figure 3.12: Illustration of the segmentation process for a pigmented skin lesion image. (a) Original image. (b) Image in (a) after shading attenuation. (c) First re-ordered independent component/channel of image (b). (d) Lesion localization mask. (e) The Level-set segmentation result. (f) Final lesion segmentation, after post-processing image (e).

in all tests in comparison with skin lesion segmentation approaches representative of the state-of-the-art.

3.3.4 Conclusions and Future Work

Several MPSL segmentation methods have been proposed for dermoscopy images, however skin lesion segmentation on macroscopic images have not received much attention. This paper proposes a new method for segmenting pigmented skin lesions on macroscopic images acquired with standard cameras. Our proposed method uses the ICA approach for lesion localization, and a Level-set based segmentation algorithm to obtain

truth is larger than the lesion itself.

Table 3.5: Comparison of the obtained segmentation errors (in average).

Approach	ϵ in average	ϵ in average w/ Sh. Att.
(CELEBI et al., 2008)	64.80%	42.93%
(GOMEZ et al., 2008)	245.70%	123.65%
(IYATOMI et al., 2008)	29.34%	21.64%
(CAVALCANTI; YARI; SCHARCANSKI, 2010)	25.91%	-
Proposed method	**16.55%**	-

Table 3.6: Segmentation errors in terms of error percentages.

Approach	$\epsilon < 10\%$	$\epsilon < 20\%$	$\epsilon < 30\%$	$\epsilon < 40\%$	$\epsilon > 100\%$
(CELEBI et al., 2008)	4.96%	12.77%	20.57%	29.79%	0.71%
Celebi w/ Sh. Att.	2.84%	25.53%	46.81%	58.16%	2.84%
(GOMEZ et al., 2008)	14.18%	39.01%	44.68%	46.81%	43.97%
Gómez w/ Sh. Att.	21.28%	58.87%	67.38%	70.92%	20.57%
(IYATOMI et al., 2008)	14.18%	52.48%	74.47%	83.69%	1.42%
Iyatomi w/ Sh. Att.	19.86%	60.99%	80.85%	90.07%	0.00%
(CAVALCANTI; YARI; SCHARCANSKI, 2010)	22.70%	57.45%	71.63%	79.43%	0.00%
Proposed method	**30.50%**	**75.18%**	**87.94%**	**95.74%**	**0.00%**

more precisely the lesion boundary, which is an important feature to discriminate benign and malignant lesions. In the next stage of this work, we intend to integrate and test our proposed lesion segmentation and boundary detection method in a MPSL pre-screening system, that can be used by non-specialists.

References

See the unified bibliography of the thesis.

3.4 A Two-Stage Approach for Discriminating Melanocytic Skin Lesions Using Standard Cameras

Cavalcanti, Pablo G. ; Scharcanski, Jacob ; Baranoski, G. V. G. . A Two-Stage Approach for Discriminating Melanocytic Skin Lesions Using Standard Cameras. Expert Systems with Applications, In Press, 2013.

Abstract

In this paper, we propose a novel approach to discriminate malignant melanomas and benign atypical nevi, since both types of melanocytic skin lesions have very similar characteristics. Recent studies involving the non-invasive diagnosis of melanoma indicate that the concentrations of the two main classes of melanin present in the human skin, eumelanin and pheomelanin, can potentially be used in the computation of relevant features to differentiate these lesions. So, we describe how these features can be estimated using only standard camera images. Moreover, we demonstrate that using these features in conjunction with features based on the well known ABCD rule, it is possible to achieve 100% of sensitivity and more than 99% accuracy in melanocytic skin lesion discrimination, which is a highly desirable characteristic in a prescreening system.

3.4.1 Introduction

Malignant melanoma is a type of melanocytic skin lesion, and it is among the most dangerous forms of cancer. According to World Health Organization (2011), approximately 132,000 melanoma cases occur globally each year. The early diagnosis of melanomas is essential for the patient prognosis since most malignant skin lesion cases can be treated successfully in their initial stages. However, a benign atypical melanocytic nevus shares at least some, and sometimes all, of the clinical characteristics of a malignant melanoma, and discriminating benign from malignant cases is often challenging, sometimes in the extreme (RAO et al., 1997; FIKRLE; PIZINGER, 2007).

It would be desirable that methods for detecting this malignancy approach to 100% of sensitivity, meaning that almost no malignant melanomas are missed. Ideally, this high degree of sensitivity should be coupled with a high degree of specificity, and consequently a high degree of accuracy (RAO et al., 1997). Currently, the most reliable method is the histopathology analysis. However, this requires that the patient undergoes surgical excision. Although this may not be a significant problem in some cases, this is a significant problem for patients with the atypical-mole syndrome, presenting several melanocytic nevi (i.e. 100 nevi or more (RAO et al., 1997), not viable of doing biopsy in all nevi). Moreover, this option is very time consuming, considering the patient and medical staff dislocation, the surgery preparation and the posterior patient monitoring.

To help diagnosing melanocytic skin lesion, physicians often use dermoscopy. This non-invasive technique allows the magnification of submacroscopic structures through the use of an optical device (a dermoscope) and liquid immersion. According to MAYER (1997), the use of dermoscopy can increase the diagnosis sensitivity in 10-27% with respect to the clinical diagnosis. Moreover, some computer-aided diagnosis systems have been proposed in the literature to help the analysis of dermoscopy images. An example is the approach proposed by CELEBI et al. (2007), which achieved 92.34% and 93.33% of sensitivity and specificity, respectively. IYATOMI et al. (2008) proposed a web-based system, obtaining a sensitivity of 85.9% and a specificity of 86%. More recently, RUIZ et al. (2011) combined three classification algorithms and obtained classification rates of

87.76%. However, even with the assistance of dermoscopy, differentiating malignant and benign lesions is a difficult task. In fact, specialists state that in the early evolution stages of malignant lesions, dermoscopy may not be helpful since it often does not improve the diagnosis accuracy (SKVARA et al., 2005).

Still considering early stage cases, there are practical situations where a non-specialist (e.g., a physician not trained on dermatology) wishes to have a qualified opinion about a suspect skin lesion, but only standard camera imaging is available on site. In such situations, telemedicine is justifiable, and the non-specialist can capture a macroscopic image of the suspect melanocytic skin lesion and send it to a specialist, who can analyze it in higher detail. In this particular situation, a teledermatology consultation brings benefits such as the easier access to health care and faster clinical results (MASSONE et al., 2008). Besides these benefits, recent results also suggest that teledermatology tends to be effective and reliable (WHITED, 2006).

We also can find in the literature systems designed to assist the diagnosis of melanocytic skin lesions using standard camera images. MANOUSAKI et al. (2006) proposed a system that achieved a sensitivity of 60.9% and a specificity of 95.4% in their experiments. The approach proposed by TABATABAIE; ESTEKI; TOOSSI (2009) resulted in 85% and 92.5% of sensitivity and specificity, respectively. ALCON et al. (2009) described a methodology to obtain a lesion classification with sensitivity of 89.72% and specificity of 75.56%. CAVALCANTI; SCHARCANSKI (2011) recently proposed a system even more accurate, achieving a sensitivity of 96.2% and a specificity of 97.7%.

Regardless of the type of image and the devices employed in the data (image) acquisition, the systems proposed to assist the diagnosis of skin lesions usually perform the following operations: (1) preprocessing, in which image artifacts, such as hair or uneven illumination, are eliminated; (2) segmentation, in which the lesion boundaries are determined; (3) feature extraction, in which a quantitative representation for the lesion area is generated; and (4) classification, in which an estimate of whether the lesion is benign or malignant is produced. In general, we can find very different methodologies to perform the first two operations. On other hand, all systems employ the same methodology to perform the last two operations. It consists in reproducing ABCD rule of dermoscopy, a medical criterion that tries to differentiate malignant and benign melanocytic lesions using a set of lesion characteristics, namely Asymmetry, Border irregularity, Color variation and Differential structures.

Beyond the ABCD rule of dermoscopy, there are other dermatological methods, such as the Menzies Scoring Method or the 7-point Checklist (JOHR, 2002), that specialists can employ to identify malignant cases. These criteria share some similarities, and seek to determine which features should be visually identified by the physician to perform the diagnosis. However, even experienced dermatologists do not obtain 100% of accuracy, independent of the method in use (ARGENZIANO; ZALAUDEK; SOYER, 2004).

In opposite side to these methods of visual analysis, recent studies have been made using features that are not visible to the naked eye. Melanomas are malignant tumors of melanocytes, the cells responsible for the biochemical process that leads to the production of melanin pigments. And, this process (known as melanogenesis) takes place in highly specialized organelles known as melanosomes, and it results in the two main classes of melanin present in the human skin, namely the eumelanin (brown and black) and the pheomelanin (red and yellow). The color of human skin is largely determined by the presence of these pigments, whose relative concentration may vary significantly in skin lesions, notably in melanocytic tumors (JIMBOW et al., 1995). Accordingly, rele-

vant efforts have been directed toward the identification of variations of eumelanin and/or pheomelanin in melanocytic skin lesions with the purpose of improving the non-invasive diagnosis of these lesions. For example, MARCHESINI; BONO; CARRARA (2009) performed a retrospective analysis on 1671 lesions using diffuse reflectance spectroscopy. They observed a decrease in pheomelanin and an increase in eumelanin concentration in passing from benign to malignant cases. Their results are consistent with results later obtained by ZONIOS et al. (2010) from reflectance measurements performed on 1379 lesions. However, Zonios et al. observed that the spectral responses obtained from different melanin concentrations do not appear strong enough to provide a definite criterion for the diagnosis of melanocytic skin lesions. However, Zonios et al. also suggested that these responses could be potentially useful for the characterization of melanocytic lesions and the early diagnosis of melanoma if used in combination with other parameters and diagnostic criteria. More recently, MATTHEWS et al. (2011) proposed a pump-probe imaging system. Examining slices from 42 lesions, they also observed that melanomas have higher amounts of eumelanin in comparison to benign nevi, concluding that melanin features can improve diagnostic accuracy if used in conjunction with current diagnostic techniques.

The main goal of this work is to propose a two-stage classification algorithm that potentially can obtain highly accurate results, and achieve 100% of sensitivity, which is, as already mentioned, a very desirable characteristic for prescreening systems. Firstly, in Section 3.4.2 we present the algorithms for the first two initial processing stages: preprocessing and segmentation. In Section 3.4.3 is presented the feature extraction process, first computing 52 features to reproduce the ABCD rule of dermoscopy, and then computing 12 additional features that represent the eumelanin and pheomelanin lesion contents. As far as we know, this is the first time that melanin features are used in melanocytic skin lesion image analysis and classification, specially using standard camera images only. In Section 3.4.4 we present our proposed two-stage melanocytic skin lesion discrimination scheme, designed to improve on the ABCD rule for malignant cases identification. Section 3.4.5 presents a general overview of the processing steps of our melanocytic skin lesion pre-screening scheme, while Section 3.4.6 presents experimental results demonstrating how accurate this scheme can be, specially because melanin-based features are used to complement the ABCD rule-based features. Finally, Section 3.4.7 presents our conclusions and outlines directions for future work.

3.4.2 Skin Lesion Detection

Before we start quantifying medical or physiological characteristics to differentiate benign from malignant cases, it is crucial that we identify the lesion area in the input image. This task is performed by the segmentation process, which is usually preceded by a image preprocessing step. In the following subsections, these two procedures are detailed.

3.4.2.1 Preprocessing

The input image acquired with a standard camera may be affected by illumination artifacts, and if used directly in the segmentation process, shading and lesion regions could be confused. Therefore, shading is attenuated in the input image before the image segmentation.

Our preprocessing step is based on the approach proposed by CAVALCANTI; SCHARCANSKI; LOPES (2010), which is based on the idea that a physician (or a person respon-

sible for taking the lesion picture) focuses a camera on the skin lesion, placing the lesion area in the center of the image, and healthy skin pixels appear at the image corners.

The first step of our shading attenuation method is to convert the image from the original RGB color space to the HSV color space, and retain the Value channel V_{orig}, which presents the higher visibility of the shading effects. We extract 20×20 pixels in each V corner and define S as the union of these four sets. This pixel set with 1600 pixels is used to adjust the quadric function $z(x, y) = P_1x^2 + P_2y^2 + P_3xy + P_4x + P_5y + P_6$, where the six quadric function parameters P_i $(i = 1, ..., 6)$ are chosen to minimize the quadratic error.

Calculating the quadric function $z(x, y)$ for each image spatial location (x, y), we get an estimate $z(x, y)$ of the local illumination intensity in the image V_{orig}. Dividing the original V_{orig} channel by $z(x, y)$, we obtain a new Value channel V_{proc} where the shading effects have been attenuated. The final step consists in replacing V_{orig} by V_{proc}, and converting the image from the HSV color space to the original RGB color space. This shading correction method has limited effect on local cast shadows, and may fail to reduce shading artifacts in areas that are not locally smooth since it is based on a second-order quadric function. However, as can be seen in Figures 3.13(a)-(b), this algorithm tends to relight the healthy skin area without loosing lesion information. In this way, we obtain an image where the skin lesion can be segmented more easily.

We consider the healthy skin region surrounding the lesion to compute lesion characteristics, specially the melanin features. The shading attenuation method may generate a highly contrasted image, with a bright skin color, and that may influence negatively the melanin features estimation. Thus, to facilitate the feature extraction stage, we perform one more image preprocessing step, correcting the illumination to obtain more realistic skin colors (more similar to those in the original image), while trying to obtain an even illumination condition. We extract again the Value channels (V from the HSV color space) from the original image (V_{orig}) and from the processed image (V_{proc}), and generate a new Value channel (V_{new}):

$$V_{new}(x, y) = \frac{V_{proc}(x, y) * \mu_{V_{orig}}}{\mu_{V_{proc}}} \tag{3.29}$$

where (x, y) indicates the pixel coordinates, and μ_V represents the mean value of the referred V channel. After that, V_{new} replaces V_{proc}, and we obtain an image with even illumination and correct colors representation by converting from the HSV to the RGB color space. Fig 3.13 shows an example of this process, presenting an original melanocytic skin lesion image, the result of applying the shading attenuation method in this image, and also shows the image obtained by applying Eq. 3.29 (to improve the image quality for the feature extraction process).

3.4.2.2 Segmentation

Several techniques have been proposed for segmenting melanocytic skin lesions in dermoscopy images (CELEBI et al., 2008, 2009; IYATOMI et al., 2008; GOMEZ et al., 2008; WIGHTON et al., 2009, 2011). However, dermoscopy obtains skin lesion images with constant illumination. Consequently, these methods do not perform well on standard camera images, and other approaches have been proposed specifically for these images. Usually, the input standard camera image is converted to grayscale and a thresholding-based algorithm is used to identify the lesion area (MANOUSAKI et al., 2006; RUIZ et al., 2011; ALCON et al., 2009). Also, segmentation approaches have been proposed

(a)

(b)

(c)

Figure 3.13: Illumination and color corrections of a melanocytic skin lesion image. (a) The original image. (b) Fig. (a) after the shading attenuation step. (c) Fig. (b) after the color correction step.

using the original color image obtained with a standard camera (TANG, 2009; WONG; SCHARCANSKI; FIEGUTH, 2011). However, these image representations may confuse healthy and lesion pixels due to its color or grayscale similarity. So, in our experiments we used a multichannel image representation that maximizes the discrimination between healthy and unhealthy skin regions as proposed in (CAVALCANTI; SCHARCANSKI, 2011).

As proposed in (CAVALCANTI; SCHARCANSKI, 2011), we create a new 3-channel normalized image $\bar{I}_i^N$ (i.e., $\bar{I}_i^N(x,y) \subset [0,1], \forall i$) based on the normalization of the RGB channels $\bar{I}_i^C$ of the input image. The first channel is a representation of the image darkness, relying on the fact that lesion areas are hyper-pigmented skin regions. Each pixel is defined as $\bar{I}_1^N(x,y) = 1 - \bar{I}_1^C(x,y)$, i.e., the complement of the normalized Red channel.

The second channel is a intensity variation representation, since the intensity variability usually is higher in lesions than in healthy skin areas. Being $\bar{L}$ a normalized Luminance image defined by the average of the three $\bar{I}_i^C$ channels, we quantify the textural variability in $\bar{L}(x,y)$ by computing $\tau(x,y,\sigma)$:

$$\tau(x,y,\sigma) = \bar{L}(x,y)\frac{\widetilde{S}(x,y,\sigma)}{S(x,y,\sigma)}, \tag{3.30}$$

where $S(x,y,\sigma) = \bar{L}(x,y) * G(\sigma)$ (i.e., the Luminance image $\bar{L}$ is smoothed by a Gaussian filter with standard deviation σ), and $\widetilde{S}(x,y,\sigma)$ represents its complement. In this way, if an image region is dark, its intensity variation information is emphasized; if the region is bright, it is de-emphasized. However, since a single Gaussian filter may not be

sufficient to capture the textural variability, $\tau(x, y, \sigma)$ is calculated for different σ values[7], and we select its maximum value at each pixel. Finally, the $\bar{I}_2^N(x, y)$ channel is obtained normalizing these values to the range $[0, 1]$.

The third channel $\bar{I}_3^N(x, y)$ of the representation describes the color variation, assuming that healthy and unhealthy skin regions present different color distributions. The Principal Component Analysis (PCA) method is applied on the normalized colors of the image $\bar{I}_i^C(x, y)$. We observed that the first principal component (i.e., the component that maximizes the local data variance) tends to project the healthy skin pixels nearer to zero and lesion pixels tend to have larger magnitudes. So, the $\bar{I}_3^N$ channel is computed normalizing the magnitudes of the first component, and then filtering with a 5×5 median filter to reduce noise.

Once obtained this multichannel representation, the Otsu's thresholding method (OTSU, 1979) is used to segment the image. A pixel is defined as part of a lesion region if its value is higher than the threshold in at least two of the three channels $\bar{I}_i^N(x, y)$. After thresholding, the remaining skin artifacts (such as freckles and hair) are eliminated more easily. The perimeter and the area of all thresholded connected pixel sets are computed, and then this set of regions is partitioned in two clusters. All regions in the cluster with smaller areas (in average) are eliminated, and their correspondent mask pixels are set to zero. At the end, the resultant mask is filtered by a 5×5 median filter, eliminating any possible remaining artifacts that may originate rim imperfections. In Figure 3.14, we present the results for all steps of the skin lesion segmentation method, including the multichannel representation generation, the thresholding and post-processing steps.

3.4.3 Melanocytic Skin Lesion Characterization

The ABCD rule of dermoscopy is the methodology normally used to differentiate benign and malignant melanocytic skin lesions. Therefore, we propose in Section 3.4.3.1 a set of 52 features aiming at capturing the ABCD rule of dermoscopy criteria, namely, Asymmetry, Border irregularity, Color variation and Differential structures. As mentioned before, although prescreening systems usually try to reproduce the ABCD rule of dermoscopy criteria, the proposed approaches usually are not able to achieve 100% of sensitivity, as it would be desirable. So, in addition to the ABCD features, we propose 12 features based on the concentrations of eumelanin and pheomelanin, the two kinds of melanin that can potentially be used for the discrimination of melanocytic lesions.

3.4.3.1 ACBD Rule Characterization

The acronym ABCD refers to the criteria commonly used by physicians to differentiate malignant melanomas from benign melanocytic nevi; that is, malignant lesions tend to exhibit **A**symmetry, **B**order irregularity, **C**olor variation, and a **D**iameter greater than 6 millimeters (RAO et al., 1997). However, we do not adopt any kind of constraint during the image acquisition, and the diameter measurement become impracticable. So, our ABCD rule characterization is based on the ABCD rule of dermoscopy (NACHBAR et al., 1994), which changes the D letter to **D**ifferential structures, submacroscopic morphologic and vascular structures (e.g., pigment network, dots, globules) that occur more frequently in malignant cases. And, to perform this characterization, we extract 52 features as described next. It is important to observe that the $\bar{I}_i^N$ channels refer to the same channels

[7] We tested different σ values, and based on our experiments we suggest using $\sigma = 1, \frac{11}{7}, \frac{15}{7}, ..., \frac{43}{7}$, and filter window sizes of $7\sigma \times 7\sigma$.

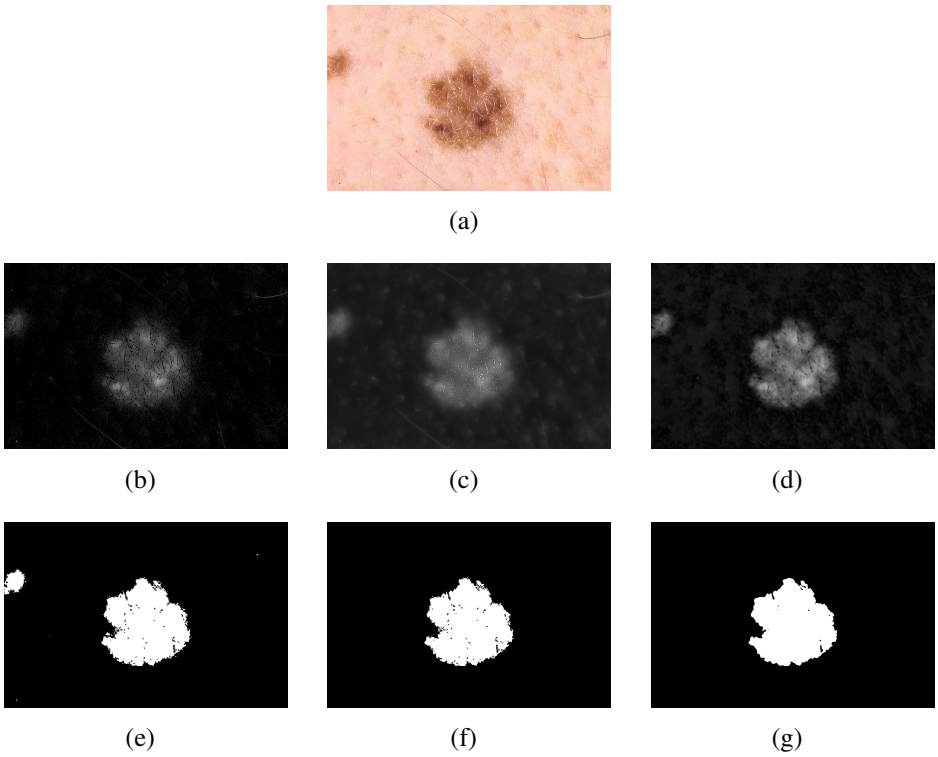

(a)

(b) (c) (d)

(e) (f) (g)

Figure 3.14: Segmentation process for the image shown in (a). In (b)-(d), respectively, the $\bar{I}_i^N$ channels representing darkness, texture variation and color variation. In (e)-(g), respectively, binary masks after thresholding, artifacts elimination and filtering.

that have been created during the segmentation process (see Section 3.4.2.2).

We start by quantifying the lesion shape, characterizing the asymmetry and irregularity characteristics of the lesion. Let A, Ac, Ab be, respectively, the areas of the segmented lesion, the convex hull and the bounding box. Moreover, p is the lesion perimeter, $L1$ the major axis of the lesion aligned with its longest diameter, and L_2 is the minor axis orthogonal to L_1. Considering these measurements, we consider the following 11 parameters (ALCON et al., 2009; CELEBI et al., 2007; FIKRLE; PIZINGER, 2007; CAVALCANTI; SCHARCANSKI, 2011): A/Ac (solidity), A/Ab (extent), $4A/(L_1\pi)$ (equivalent diameter), $4\pi A/(L_1 p)$ (circularity), p/A, ratio between the principal axes (L_2/L_1), ratio between sides of the axis L_1 and of the axis L_2, the difference between the areas of each side of the axes L_1 and L_2 normalized by the whole area A, and ratio between sides of the lesion bounding box.

The boundary sharpness is quantified by the magnitude of the gradient using the Sobel operator. We dilate the lesion rim by 2 pixels, obtaining a 5 pixels wide region, and compute the following 6 features: the average and the variance of the gradient magnitudes of all pixels in the extended rim in each one of the three $\bar{I}_i^N$ channels.

To reproduce closely the application of the ABCD rule in terms of asymmetry and border irregularity, the next features are based on a sub-division of the lesion rim in 8 symmetric regions (NACHBAR et al., 1994). In addition to the two principal axes L_1 and L_2, we rotate these orthogonal axes by 45 degrees and obtain two additional axes, obtaining 8 symmetric regions $R = 1, ..., 8$. For each channel $\bar{I}_i^N$, the average gradient magnitudes of the extended rim pixels $\mu_{R,i}(R = 1, ..., 8)$ are computed. Therefore, in this way we calculate 6 more features: the average and the variance of the 8 $\mu_{R,i}$ gradient magnitude values in each one of the three $\bar{I}_i^N$ channels.

Table 3.7: Six Possible Colors of a Lesion on the RGB Color Space. (ALCON et al., 2009)

Color	Red	Green	Blue
White	1	1	1
Red	0.8	0.2	0.2
Light Brown	0.6	0.4	0
Dark Brown	0.2	0	0
Blue-Gray	0.2	0.6	0.6
Black	0	0	0

To represent the C letter of the ABCD rule, we start by computing maximum, minimum, mean and variance of the intensities in the color variation channel $\bar{I}_3^N$. The same four features are also computed for each one of the three original $\bar{I}_i^C$ channels. We also consider the ratios between the mean values: $mean(\bar{I}_1^C)/mean(\bar{I}_2^C)$, $mean(\bar{I}_1^C)/mean(\bar{I}_3^C)$ and $mean(\bar{I}_2^C)/mean(\bar{I}_3^C)$, totalizing 19 features.

Furthermore, physicians usually identify six distinct hues in skin lesions: white, red, light and dark brown, blue-gray, and black (NACHBAR et al., 1994). So, we compute the occurrence of these typical hues within the lesion. Given a pixel in the lesion segment, we find the nearest reference color (associated with a typical hue, see Table 3.7 (ALCON et al., 2009)) using the Euclidean distance to the pixel color in $\bar{I}_i^C$. A hue occurrence counter is created, one cell per typical hue. For each lesion pixel, the nearest typical hue counter is increased by 1. Finally, typical hue counters are normalized/divided by the lesion area A, and generate 6 additional features.

Finally, we quantify the differential structures. The differential structure characteristics are not easily discriminated based on macroscopic images, however these structures tend to generate different intensity variations in benign and malignant lesions. So, we extract 4 features, namely the maximum, minimum, mean and variance of the intensity variation channel $\bar{I}_2^N$, considering only the pixels inside the lesion segment.

3.4.3.2 Melanin Content Variation Characterization

As already mentioned in Section 3.4.1, recent studies have analyzed the melanin variation in melanocytic skin lesion to differentiate benign and malignant cases through spectroscopy. In this work, we follow a different approach. Instead of using spectroscopy, we propose to estimate the melanin variation in skin lesions using standard camera images. In order to perform this task, we employ a biophysically-based spectral model of light interaction with human skin developed by Krishnaswamy and Baranoski (KRISHNASWAMY; BARANOSKI, 2004; BARANOSKI; KRISHNASWAMY, 2010) known as BioSpec (see more details about BioSpec in Appendix A, Section 3.4.8). This model simulates the reflectance spectra of skin specimens within the visible spectral domain (i.e., from 400 to 700 nm). It takes into account the specimens' biophysical characteristics described through a set of biophysical parameters. In this work, unless otherwise stated, we employed the default values assigned for these parameters in the online version of the model (Natural Phenomena Simulation Group, 2011). This set of parameters includes the concentration of eumelanin and pheomelamin. Although BioSpec was developed for healthy skin tissue, we believe that by varying the concentration of eumelanin and pheomelanin we can obtain sound estimates for the reflectance spectra of melanomas since the main factor responsible for their distinct spectral appearance, notably in their

initial stages, is their high concentration of melanins (LAZOVA; PAWELEK, 2009).

In Fig. 3.15, we present a plate showing the skin colors that were obtained by converting the reflectance spectra provided by BioSpec given a wide range of parameter values (see Appendix B, Section 3.4.9 for more details about how this conversion is performed). We varied the concentration of eumelanin from 20 to 300 g/L, and the concentration of pheomelanin from 4 to 60 g/L, both with steps of 4 g/L, and used the default values for the other model parameters. The choice of using eumelanin predominantly in these simulations is supported by optical scanning studies in which the spectrum of malignant melanomas tend to show an eumelanin predominance in their spectral characteristics (ZONIOS et al., 2008; DIMITROW et al., 2009). Thus, we used 71 and 15 values for the eumelanin and pheomelanin parameters, respectively, obtaining 1065 skin colors with BioSpec. The colors appearing in the lesion images (photographs) are matched with the skin colors derived from the reflectance spectra provided BioSpec, and the estimates of the eumelanin and pheomelanin parameters are obtained.

Figure 3.15: The possible colors of skin, varying the concentration of eumelanin from 20 to 300 g/L (in the horizontal axis) and the concentration of pheomelanin from 4 to 60 g/L (in the vertical axis).

3.4.3.3 Extracting Melanin Variation Features from Skin Lesion Images Captured with Standard Cameras

Given a pixel, its color is compared to each one of the 1065 possible modeled colors obtained using BioSpec. This comparison, consists of using the minimum Euclidean distance to find the modeled color nearest to the pixel color. This information allows us to estimate the concentration of eumelanin and pheomelanin in the skin location that corresponds to this pixel. To compare the variation of the pigment concentrations within the lesion with the concentrations in healthy skin, we also extract these concentration values in the healthy skin in the surround area outside the lesion. In order to locate the healthy skin in the area surrounding the lesion, we dilate the binary mask obtained in the segmentation step with a disk structuring element of 30 pixels of radius. Then, we subtract the original segmented lesion binary mask from the dilated segmentation binary mask, obtaining just the region surrounding the lesion with a width of approximately 30 pixels. An example of this operation can be seen in Fig. 3.16. Located the healthy skin region surrounding the lesion, we estimate the concentrations of eumelanin and pheomelanin in each inner pixel in the way described above. Given a healthy skin pixel, we also estimate these values by minimum Euclidean distance to the 1065 modeled colors as mentioned before.

After computing the estimates of the eumelanin and pheomelanin concentrations associated with a set of pixels of the given image, we extract the following melanin features

(a) (b) (c)

Figure 3.16: Examples of the binary masks used in the melanin features computation. (a) The original image. (b) The lesion segmentation binary mask. (c) The binary mask corresponding to the healthy skin area surrounding the lesion.

mf_i:

mf_1**:** Average of the eumelanin concentration of all pixels inside the lesion segment;

mf_2**:** Average of the eumelanin concentration of all pixels inside the lesion segment as well as all pixels inside the healthy skin surrounding area;

mf_3**:** Standard deviation of the eumelanin concentration of all pixels inside the lesion segment;

mf_4**:** Standard deviation of the eumelanin concentration of all pixels inside the lesion segment as well as all pixels inside the healthy skin surrounding area;

mf_5**:** The difference between the average of the eumelanin concentration inside the lesion segment and the average of the eumelanin concentration inside the healthy skin surrounding area;

mf_6**:** The ratio between the average of the eumelanin concentration inside the lesion segment and the average of the eumelanin concentration inside the healthy skin surrounding area;

We also consider six more pheomelanin features mf_7-mf_{12}, similar to mf_1-mf_6, but instead of using eumelanin concentration we use the pheomelanin concentration.

3.4.4 Proposed Two-Stage Melanocytic Skin Lesion Discrimination Scheme

Once the features that characterize the melanocytic lesions are extracted, we use these information to discriminate the malignant and the benign cases. The following Sections detail our two-stage discrimination approach, first using the ABCD rule-based features, and next refining by the discrimination results using the melanin variation features. The melanocytic skin lesion discrimination approaches proposed in the literature usually contain a single discrimination step based on the ABCD rule, and as already mentioned in Section 3.4.1, are not able to obtain the desirable 100% of sensitivity. So, we add a second discrimination stage to our scheme, trying to correct cases that incorrectly received a benign label in the initial stage. The melanin variation features improve the discrimination results obtained with the ABCD rule (first stage), helping achieve a higher degree of sensitivity (i.e. so virtually no malignant case would be missed).

3.4.4.1 First Stage: Melanocytic Skin Lesion Discrimination Using the ABCD Rule

Based on the 52 ABCD features that have been extracted (see Section 3.4.3.1), we generate a preliminary discrimination of benign and malignant melanocytic skin lesions. However, classifiers tend to be more efficient if these feature values are normalized (i.e. produce values within a specified range). We normalize feature values with the z-score transformation (CELEBI et al., 2007; AKSOY; HARALICK, 2000):

$$Z_{i,j} = \frac{((v_{i,j} - \mu_j)/(3\sigma_j) + 1)}{2}, \tag{3.31}$$

where $v_{i,j}$ is the value of the jth feature of the ith sample (image), μ_j and σ_j are the mean and standard deviation of the jth feature, respectively, of all training set samples. After the z-score transformation, most of the $Z_{i,j}$ values are in the [0,1] range. The out-of-range values are saturated to either 0 or 1.

To create the training set, half of the samples in each class (benign or malignant) are randomly selected. Additionally, since we have a relatively small image data set, new samples are added to the training set using the Smoothed Bootstrap Resampling method (i.e., original samples are randomly selected, and new ones are created adding a small amount of zero-centered noise to their feature values, enlarging the data set (YOUNG, 1990)). In our experiments, we used zero mean Gaussian noise with $\sigma = 0.1$, obtaining from 500 to 2500 samples for each class so the classes are balanced (i.e., have the same number of samples). Using Normal probability plots, we observed that Normality is a reasonable assumption for our features , and by introducing small Gaussian noise we avoid "ties" that may lead to overfitting in random oversampling procedures.

Finally, we apply a classifier to determine the class (benign or malignant) of each sample/image. We used a K-Nearest Neighbor Classifier (KNN) with K=1, or simply a Nearest Neighbor Classifier, because this is a very simple method where each sample/image is assigned to its closest neighboring class in Euclidean (feature) space. Besides its simplicity, this method has been used in melanocytic skin lesion image classification research (DREISEITL et al., 2001; BURRONI et al., 2004), producing very accurate results (CAVALCANTI; SCHARCANSKI, 2011). Also, this classifier also has been successfully used for the classification of non-melanoma skin lesions (BALLERINI et al., 2012a,b).

3.4.4.2 Second Stage (Beyond the ABCD Rule): Melanocytic Skin Lesion Discrimination with Melanin Variation Features

A desirable characteristic for prescreening systems is to achieve 100% of sensitivity, i.e. to not identify mistakenly any malignant case as benign. However, we observed in our experiments that considering the proposed melanin variation features mf_1-mf_{12} in the first classification stage, i.e. when using the KNN classifier, is not efficient to achieve this goal (i.e. of obtaining high sensitivity). So, in order to reach high sensitivity rates, we propose to include a second classification stage. After obtaining the first melanocytic skin lesion classification result, we re-classify all samples/images that have been initially identified as benign using the melanin variation features mf_1-mf_{12}. This second stage is based on the Bayes's Classifier (ALPAYDIN, 2004) (which provided the best classification results in our experiments), and assigns a sample to the most likely class based on the extracted features. Let L be a set of melanin features [mf_1,...,mf_{12}] extracted from a lesion, $D = 1$ the malignant class and $D = 0$ the benign class. Using Bayes' rule, the posterior probability $P(D|L)$ of assigning a lesion to class D can be written as:

$$P(D|L) = \frac{P(L|D)P(D)}{P(L)}, \tag{3.32}$$

where, $P(L|D)$ is the conditional probability of finding a lesion with the L characteristics in class D, $P(D)$ is the prior probability of class D, and $P(L)$ is the evidence and it represents the probability that the features L are seen. To the best of our knowledge, $P(D)$ is not known, and considering that our training sets have an equal number of samples per class, we consider $P(D = 0) = P(D = 1)$ and this term can be discarded from Eq. 3.32. Since $P(L)$ is constant with respect to $argmax_D(P(L|D)/P(L))$ and therefore has no effect on the solution, it can be also discarded, and a lesion can be considered malignant if:

$$P(L|D = 1) \ > \ P(L|D = 0), \tag{3.33}$$

.
and it is benign otherwise. This classification method is usually referred to as the Maximum Likelihood Classifier. So, the challenge is to properly estimate $P(L|D)$ in a way that minimizes the probability classification error. To perform this task and generate accurate classification results, we propose the following discrimination scheme:

- To generate the training set, we use half of the samples of each class (the same samples that have been randomly selected to train the first discrimination stage). Since the number of sample is small, specially in the benign class, we use again Smoothed Bootstrap Resampling (see Section 3.4.4.1) to generate 2500 samples for each class (i.e., benign and malignant), adding Gaussian noise ($\mu = 0$ and $\sigma = 0.1$) to the original mf_i features. In this way, we obtain more statistical significant data;

- Given the training set and a subset ψ containing the mf_i features, we generate two multivariate probabilities density functions (pdfs): one for the benign class $P(L|D = 0)$ and one for the malignant class $P(L|D = 1)$. We model these pdfs using an unsupervised modified EM (Expectation-Maximization) algorithm to identify finite Gaussian mixture models as proposed by FIGUEIREDO; JAIN (2002);

- Each sample L that has been previously classified as benign is re-tested, and it is re-classified by a Maximum Likelihood Classifier as malignant if its ψ features indicate a higher probability of being malignant (see Eq. 3.33).

We tested this two-stage discrimination scheme with different ψ sets of features, and the obtained results are presented in Section 3.4.6.

3.4.5 Method Overview

In the last three Sections we presented all the steps of the proposed two-stage discrimination scheme, and Fig. 3.17 presents the overall system workflow used to identify a melanocytic skin lesion as benign or malignant. We start by preprocessing the input image, i.e. attenuating the shading effects, and segmenting the skin lesion. After that, the first discrimination stage is performed, using the 52 features based on the ABCD rule and a KNN classifier, assigning a benign/malignant label to the imaged lesion.

The second discrimination stage tries to enhance as much as possible the melanocytic skin lesion classification sensitivity. All images that have been initially classified as benign are re-classified now using the melanin variation features extracted from inside the

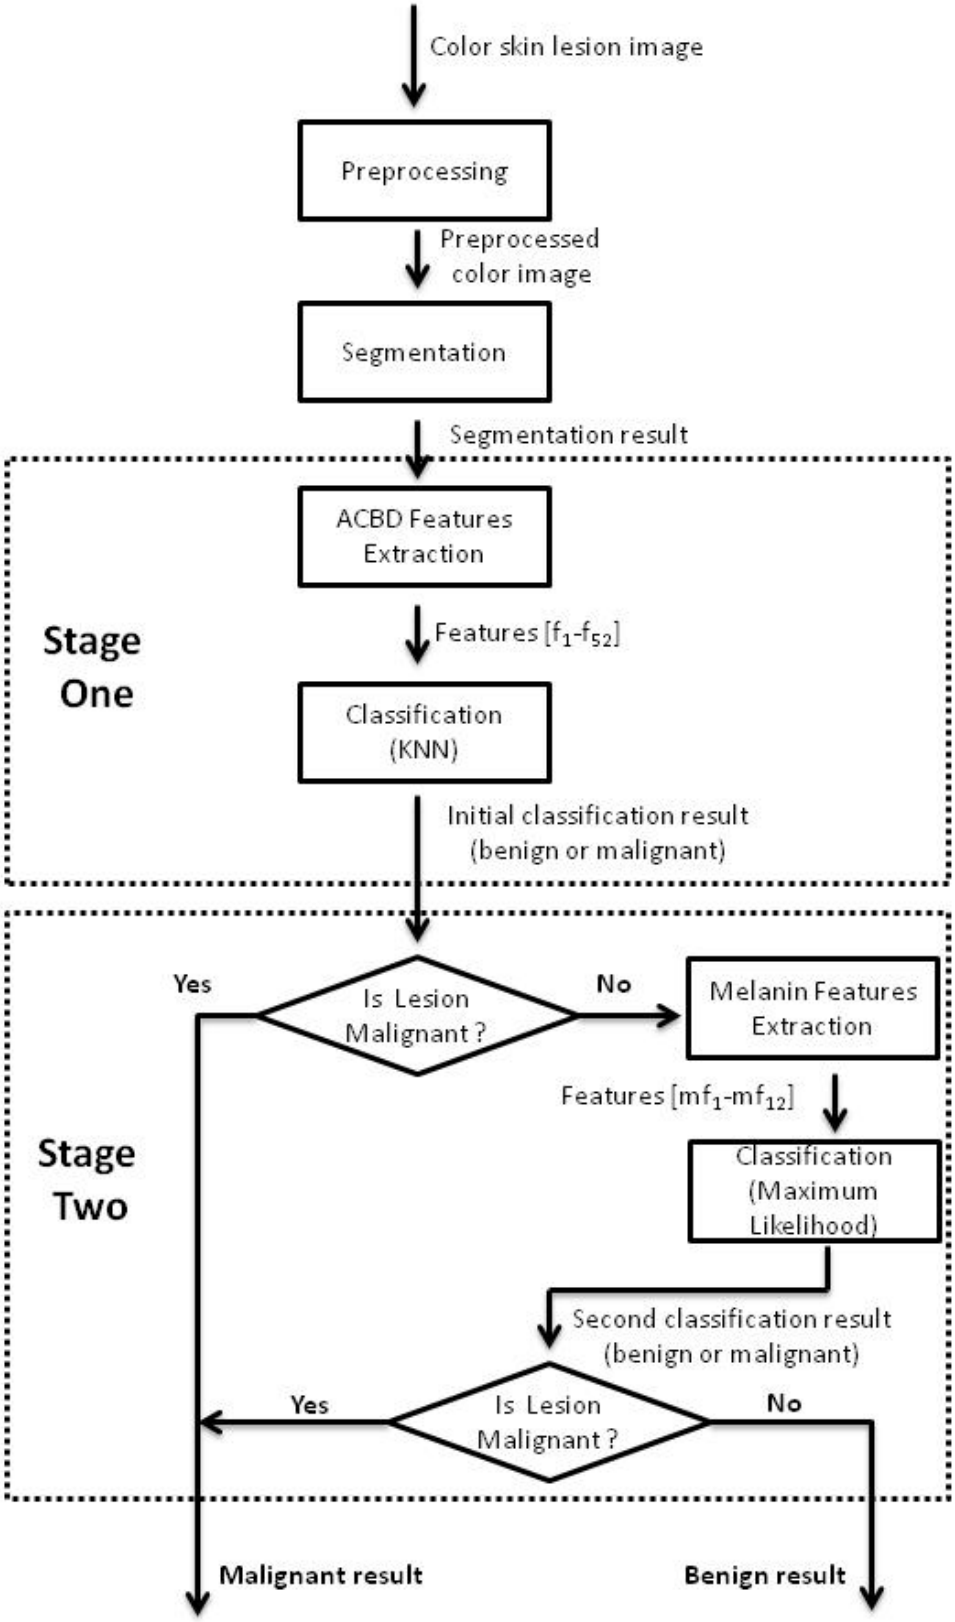

Figure 3.17: Overall system workflow.

lesion and from the healthy skin area surrounding the lesion. The final classification result is provided by a Maximum Likelihood Classifier.

3.4.6 Experimental Results

We performed our experiments on the same image dataset that have been used in the experiments of ALCON et al. (2009). They collected an image set containing 45 benign atypical nevi (or Clark nevi) and 107 melanomas from the Dermnet dataset (DERMNET, 2010).

Considering that training our proposed two-stage discrimination scheme requires a random resampling step, we repeated the whole process 50 times, seeking to obtain a training set that characterizes well the two classes, i.e., the benign and malignant cases. Furthermore, we tested different ψ sets of melanin variation features (see Section 3.4.4.2), trying to identify the feature set that obtains less false negatives. We started the selection of the ψ feature sets by taking each mf_i feature individually, and then increased the set size by adding one more feature at a time, until we reached 6 features per set, when we observed that the obtained results were decreasing in quality. We tested all the possible feature set combinations: all the 12 features individually, the 66 combinations of 2 features, the 220 combinations of 3 features, the 495 combinations of 4 features, the 792 combinations of 5 features, and 924 combinations of 6 features. In Table 3.8 we present the ψ feature sets considering these number of features that resulted in the best classification results, i.e. the sets that obtained 100% of sensitivity associated with high specificity, and using the number of times that results was obtained as a tiebreaker. We also tested the 6 eumelanin features, the 6 pheomelanin features, and the combination of all 12 melanin variation features as three feature sets, and the corresponding results also are presented in Table 3.8.

Considering that for each feature set ψ we generate randomly 50 training and testing sets by Smoothed Bootstrap Resampling, we only present the average results, the best obtained results in terms of accuracy and the number of times that these results were obtained (when testing with the 50 different sets). However, although we use randomly generated training sets to improve the chance of characterizing well the lesion classes, it shall be observed that we only use the original data (and not the synthetic data) for testing.

It is important to observe that in the experiments of ALCON et al. (2009) with the same dataset, but not including bootstrapping in their training step, the best obtained results were: sensitivity of 89.72% and specificity of 75.56%. Later, CAVALCANTI; SCHARCANSKI (2011) included the bootstrapping step, and presented a system even more accurate in this dataset, achieving a sensitivity of 96.2% and a specificity of 97.7%. So, comparing the results of these approaches with our results shown in Table 3.8, we may conclude that our proposed discrimination scheme potentially can enhance the sensitivity (i.e., reduce the number of false negatives), while maintaining or even increasing the accuracy, although the specificity may decrease slightly. The set of features containing all 12 melanin variation features obtained very good results in average, and achieved 100% of sensitivity in 11 of the 50 test sets. However, the best classification results have been obtained using the feature set $\psi = \{mf_6, mf_8, mf_{10}, mf_{12}\}$. This set containing one feature based on the eumelanin concentration and three based on the pheomelanin concentration achieved 100% of sensitivity in 24 of the 50 test sets, combined with high specificity and, consequently, high accuracy.

Important to observe the presence of mf_6 and mf_{12} (i.e., the ratio between the eumelanin concentration inside the lesion and in the healthy skin, and the ratio between the pheomelanin concentration inside the lesion and in the healthy skin) in the feature set that obtained the best classification results. That indicates that although the concentration of melanin inside the lesion is an important feature to discriminate malignant and benign

Table 3.8: Discrimination results in 50 trials.

		Sensitivity	Specificity	Accuracy	Number of Times the result was obtained
$\psi = \{mf_{12}\}$	Average	100%	73.02%	92.01%	-
	Best result	100%	80%	94.08%	1
$\psi = \{mf_{10}, mf_{12}\}$	Average	99.83%	85.87%	95.70%	-
	Best result	100%	95.56%	98.68%	2
$\psi = \{mf_6, mf_{10}, mf_{12}\}$	Average	99.68%	90.76%	97.04%	-
	Best result	100%	97.78%	99.34%	6
$\psi = \{mf_6, mf_8, mf_{10}, mf_{12}\}$	**Average**	**99.70%**	**96.18%**	**98.66%**	-
	Best result	**100%**	**97.78%**	**99.34%**	**24**
$\psi = \{mf_6, mf_8, mf_{10}, mf_{11}, mf_{12}\}$	Average	99.57%	95.96%	98.50%	-
	Best result	100%	97.78%	99.34%	16
$\psi = \{mf_3, mf_6, mf_8, mf_{10}, mf_{11}, mf_{12}\}$	Average	99.64%	94.31%	98.07%	-
	Best result	100%	97.78%	99.34%	9
$\psi = \{mf_1, \dots, mf_6\}$	Average	99.35%	90.18%	96.63%	-
	Best result	99.07%	97.78%	98.68%	8
$\psi = \{mf_7, \dots, mf_{12}\}$	Average	99.66%	92.98%	97.68%	-
	Best result	100%	97.78%	99.34%	4
$\psi = \{mf_1, \dots, mf_{12}\}$	Average	99.48%	96.31%	98.54%	-
	Best result	100%	97.78%	99.34%	11

cases, it is also important to consider the characteristics of the healthy skin to identify if the lesion pigmentation variation is comparable to a melanoma or not.

We also analyzed the performance of each classification stage of our approach separately. We observed that when our system achieves its best result (i.e., sensitivity of 100%, specificity of 97.78% and accuracy of 99.34%), the first stage which uses the ABCD features only generated a sensitivity of 96.26% and specificity of 97.78%, and a total accuracy of 96.71%. Considered independently, the second stage which is based on melanin features, obtained a lower sensitivity to 94.39%, high specificity of 100%, and 96.05% of accuracy. However, it is important to recall that we only use this second stage for the samples that have been previously classified as benign (as a verification stage). Consequently, our two-stage approach achieves 100% of sensitivity and a very high level of accuracy (i.e. 99.34%). Although the second stage individually obtains worse results in terms of sensitivity and accuracy, our experiments demonstrate that the melanin features potentially can help identify melanomas that have been erroneously classified as benign initially.

It shall be observed that our results were obtained using the 1065 different skin colors, that have been modeled using the average characteristics of human skin. Our classification results suggest that our approach potentially could estimate the melanin variation in different types of skin. Despite of the fact that our image set does not retain any information about patients, the variability of skin characteristics is noticeable in this set. However, this dataset does not include melanocytic lesions occurring in dark (black) skin, and computing a different color table may be necessary for these skin tones. The very low incidence of melanomas (melanocytic lesions) in dark (black) skin, and the focus on the most common cases, justifies the absence of such examples of in our dataset. In fact, the incidence of melanomas in dark (black) skin is a rare phenomenon, and recent statistics show this incidence to be more than 20 times lower than in other skin tones (HOWLADER et al., 2012)).

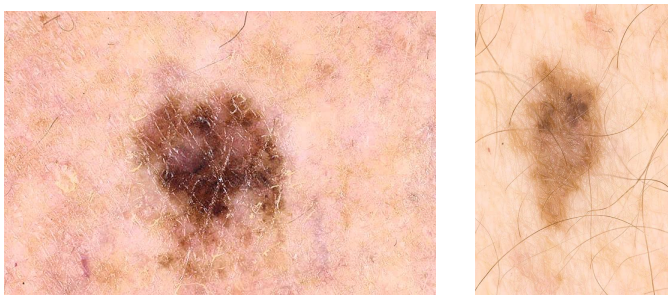

(a) (b)

Figure 3.18: Examples of lesions that have been previously identified as benign and then were correctly re-classified as malignant. These images are displayed with their illumination corrected for better visualization.

In Fig. 3.18, we illustrate two examples of melanocytic skin lesions that have been incorrectly identified as benign nevi in the first classification stage, and then were correctly classified as malignant melanomas in the second classification stage using the proposed melanin variation features.

3.4.7 Conclusions

According to recent studies (MARCHESINI; BONO; CARRARA, 2009; ZONIOS et al., 2010; MATTHEWS et al., 2011), the concentration of the two main classes of melanin present in human skin, eumelanin and pheomelanin, can potentially be used to identify melanocytic skin lesions as benign or malignant if employed in conjunction with current diagnostic techniques. Based on these observations, we presented in this paper a new method that uses only a standard camera image (a simple photograph) to compute classification features based on the estimation of eumelanin and pheomelanin contents in melanocytic skin lesions, and employs these melanin features in a two-stage classification scheme.

In medical applications, it is especially desirable that a prescreening system achieves 100% of sensitivity, i.e., all malignant cases are correctly classified and no false negatives occur. However, as far as we know, this goal still poses a challenge for researches. With this objective in mind, we proposed a scheme that starts with a discrimination stage based on the ABCD rule, which are well known medical criteria that prescreening systems usually try to implement. Next, a second discrimination stage based on melanin variation features is used to help improve the discrimination sensitivity and accuracy. Our preliminary experimental results suggest that the proposed framework which combines ABCD rule-based and melanin variation features potentially can be very accurate and obtain high sensitivity values (near 100% in our experiments). In addition, we remark that although the proposed framework is designed for pre-screening purposes, which ideally should involve melanocytic skin lesions (e.g., melanomas) in their initial stages, our experiments were performed on images depicting melanomas in different stages of development. Hence, the high level accuracy of our preliminary results also illustrates the robustness of the proposed framework with respect to different input data (skin images).

As future work, we plan to test our framework more extensively in clinical trials, and make it part of a complete teledermatological system. As more biophysical data on melanomas becomes available, we also intend to further enhance its predictive capabilities, notably with respect of melanomas in advanced stages, by incorporating into our simulations other factors affecting the spectral signatures of melanocytic lesions such as abnormalities in the distribution and morphology of the melanosomes. Nevertheless, other features could improve the reproduction of the ABCD rule and we intend to develop and test it according to necessity. For example, fuzzy counters to represent the color presence (i.e., assigning weight distances to each reference color) and textural descriptors trying to identify the dermatological structures may possibly contribute for the differentiation of malignant and benign cases.

3.4.8 Appendix A: BioSpec: A Spectral Model of Light Interaction with Human Skin

For completeness, we summarize the main characteristics of the BioSpec model in this section. We remark that detailed descriptions of BioSpec can be found in the related publications (KRISHNASWAMY; BARANOSKI, 2004; BARANOSKI; KRISHNASWAMY, 2010) and its source code is openly available online (Natural Phenomena Simulation Group, 2011).

The BioSpec model uses Monte Carlo based algorithms to simulate the processes of light propagation (surface reflection, subsurface reflection and transmission) and absorption in the skin tissues. It considers the stratification of skin into four semi-infinite main

layers: stratum corneum, epidermis, papillary dermis, and reticular dermis. The model parameter space includes: the refractive index and thickness of each layer, the refractive index and the diameter of collagen fibrils, the extinction coefficient, concentration, and volume fraction of the main chromophores present in the skin tissues (e.g., eumelanin, phaeomelanin, oxyhemoglobin, deoxyhemoglobin, β-carotene and bilirubin) and the aspect ratio of the stratum corneum folds.

The propagation of light in the skin tissues is simulated by the BioSpec model as a random walk process (that rely on the generation of random numbers ξ_j, for $j = 1, 2, ..., 9$, uniformly distributed in the interval $[0, 1]$) using ray optics. In this random walk process, the transition probabilities are associated with Fresnel coefficients computed at each interface between the layers, and the termination probabilities are determined by the ray free path length.

Once a ray hits a skin specimen at the air/stratum corneum interface, it can be reflected back or refracted into the stratum corneum. In the former case, the BioSpec model computes the distribution of the reflected light taking into account the aspect ratio (or oblateness) of the stratum corneum folds. In the BioSpec model, these mesostructures are represented as ellipsoids. The aspect ratio, $\varsigma \in [0, 1]$, of these ellipsoids is defined as the quotient of the length of the vertical axis by the length of the horizontal axis, which are parallel and perpendicular to the specimen's normal respectively. As the folds become flatter (lower ς), the reflected light becomes less diffuse. In order to account for this effect, the reflected rays are perturbed using angular displacements obtained from the surface-structure function proposed by TROWBRIDGE; REITZ (1975), which represents rough air-material interfaces using microareas randomly curved. These displacements are given in terms of the polar perturbation angle:

$$\theta_s = \arccos\left[\left(\left(\frac{\varsigma^2}{\sqrt{\varsigma^4 - \varsigma^4\xi_1 + \xi_1}} - 1\right) b\right)^{\frac{1}{2}}\right], \tag{3.34}$$

where b corresponds to $1/(\varsigma^2 - 1)$, and the azimuthal perturbation angle, ϕ_s, is given by $2\pi\xi_2$.

If the ray penetrates the skin specimen, then it can be reflected and refracted multiple times within the skin layers before it is either absorbed or propagated back to the environment through the air/stratum corneum interface. Since the subcutaneous tissue is a highly reflective medium, for body areas characterized by the presence of hypodermis, the BioSpec model assumes total reflection at the reticular dermis/hypodermis interface. In the epidermis and stratum corneum, scattering is simulated using angular displacements measured by BRULS; LEUN (1984).

Every ray entering one of the dermal layers is initially tested for Rayleigh scattering. If the test fails or the ray has already been bounced off one of the dermal interfaces, then the ray is randomized around the normal direction using a warping function based on a cosine distribution in which the polar perturbation angle, α_c, and the azimuthal perturbation angle, β_c, are given by:

$$(\alpha_c, \beta_c) = (\arccos((1 - \xi_3)^{\frac{1}{2}}), 2\pi\xi_4). \tag{3.35}$$

In order to perform the Rayleigh scattering test, the spectral Rayleigh scattering amount, $\mathcal{S}(\lambda)$, is computed using the appropriate expression for Rayleigh scattering involving particles (MCCARTNEY, 1976). Next, a random number ξ_5 is generated. If $\xi_5 < 1 - e^{-\mathcal{S}(\lambda)}$,

then the ray is scattered using an azimuthal perturbation angle, β_R, given by $2\pi\xi_8$, and a polar perturbation angle, α_R, obtained using the following rejection sampling algorithm based on the Rayleigh phase function (MCCARTNEY, 1976):

$$\begin{aligned} &do \\ &\qquad \alpha_R = \pi\xi_6 \\ &\qquad \chi = \tfrac{3}{2}\xi_7 \\ &while \quad (\chi > \tfrac{3\sqrt{6}}{8}(1+\cos^2\alpha_R)\sin\alpha_R) \end{aligned}$$

Once a ray has been scattered, it is tested for absorption. The absorption testing is performed probabilistically every time a ray starts a run in a given layer, and it consists in estimating the ray free path length using an expression based on Beer's law (TUCHIN, 2000), and considering the total absorption coefficient, $\mu_{ai}(\lambda)$, of a given layer i, which takes into account the extinction coefficient and the concentration of each pigment present in this layer. This test consists in estimating the ray free path length, $p(\lambda)$, through the following expression:

$$p(\lambda) = -\frac{1}{\mu_{ai}(\lambda)}\ln(\xi_9)\cos\theta, \tag{3.36}$$

where θ corresponds to the angle between the ray and the specimen's normal. If $p(\lambda)$ is greater than the thickness of the layer, then the ray is propagated, otherwise it is absorbed.

3.4.9 Appendix B: Converting the Modeled Spectral Reflectance Values to Skin Colors

Unfortunately, we cannot directly access the original reflectance values associated with the pixels of standard camera images. Therefore, we first transform the modeled reflectance functions $R(\lambda)$ (in our experiments, obtained using BioSpec) to colors that can be matched straightforwardly with the lesion colors appearing in a standard image. The color tristimulus values XYZ can be calculated by the additive law of color matching (LEE, 2005):

$$X = N\sum_{\lambda} R(\lambda)\, S(\lambda)\, \overline{x}(\lambda)\, \Delta\lambda, \tag{3.37}$$

$$Y = N\sum_{\lambda} R(\lambda)\, S(\lambda)\, \overline{y}(\lambda)\, \Delta\lambda, \tag{3.38}$$

$$Z = N\sum_{\lambda} R(\lambda)\, S(\lambda)\, \overline{z}(\lambda)\, \Delta\lambda, \tag{3.39}$$

where $S(\lambda)$ is the relative spectral power distribution of the illuminant, $\overline{x}(\lambda)$, $\overline{y}(\lambda)$ and $\overline{z}(\lambda)$ are the spectral sensitivity functions, and $\Delta\lambda$ represents the wavelength intervals. In our experiments, we used the CIE Standard Illuminant D65 (Commission Internationale de l'Eclairage, 2011), the CIE 1931 sensitivity functions of the standard observer for 2°, and $\Delta\lambda = 5nm$. The constant N is defined as follows:

$$N = \sum_{\lambda} S(\lambda)\, \overline{y}(\lambda)\, \Delta\lambda. \tag{3.40}$$

Obtained the color tristimulus XYZ, we convert XYZ to the RGB space using the sRGB standard (LEE, 2005):

$$R' = 3.2410X - 1.5374Y - 0.4986Z, \tag{3.41}$$
$$G' = -0.9692X + 1.8760Y + 0.0416Z, \tag{3.42}$$
$$B' = 0.0556X - 0.2040Y + 1.0570Z. \tag{3.43}$$

As shown in (LEE, 2005), to obtain the RGB colors in the correct range, an additional non-linear transform is required. If $R', G', B' \leq 0.003040$, then we obtain the corrected $\overline{RGB}$ color representation by:

$$\overline{R} = 12.92R', \tag{3.44}$$
$$\overline{G} = 12.92G', \tag{3.45}$$
$$\overline{B} = 12.92B', \tag{3.46}$$

otherwise,

$$\overline{R} = 1.055R'^{(1/2.4)} - 0.055, \tag{3.47}$$
$$\overline{G} = 1.055G'^{(1/2.4)} - 0.055, \tag{3.48}$$
$$\overline{B} = 1.055B'^{(1/2.4)} - 0.055. \tag{3.49}$$

Finally, the obtained $\overline{RGB}$ are normalized, i.e., are transformed to the range $[0, 1]$, and then the correct eight-bit values RGB are obtained by multiplying $\overline{RGB}$ by 255.

References

See the unified bibliography of the thesis.

3.5 Other Publications

Beyond those presented in the last Chapters, the authors of this Thesis also had other publications related to the automated analysis of melanocytic skin lesions. They are described next:

- *Cavalcanti, P. G. ; Yari, Y. ; Scharcanski, J. . Pigmented Skin Lesion Segmentation on Macroscopic Images. In: 25th IEEE International Conference on Image and Vision Computing New Zealand, 2010, Queenstown. Proceedings of the 25th IEEE International Conference on Image and Vision Computing New Zealand, 2010.*
  This paper presents a very simple segmentation technique, based on the Otsu's thresholding method of the Red channel (R of the RGB color space). Although very simple, this paper demonstrates the importance of segmentation methods developed specifically for standard camera images. In our experiments, it achieves lower average segmentation error in comparison to more complex algorithms that were proposed for dermoscopy images.

- *Cavalcanti, Pablo G. ; Scharcanski, Jacob . Macroscopic Pigmented Skin Lesion Segmentation and Its Influence on the Lesion Classification and Diagnosis. In: M. Emre Celebi; Gerald Shaefer. (Org.). Color Medical Image Analysis. 1ed. Heidelberg: Springer, 2012, v. 1, p. 15-39.*
  This book chapter compares different segmentation methods proposed for melanocytic skin lesions in standard camera images, including the methods presented in Sections 3.2 and 3.3. Then, the set of 52 features proposed in Section 3.2 is extracted from each different segmentation mask, and using the Nearest Neighbor Classifier we classify each image. Experiments indicate that those features can obtain good classification results independently of the previous segmentation step. As already mentioned, this book chapter has been included as Appendix A in this Thesis.

- *Cavalcanti, Pablo G. ; Scharcanski, Jacob . Shading Correction in Human Skin Color Images. In: A. Mishra. (Org.). An Introductory Guide to the Emerging Areas of Digital Image Processing. 1ed. Sunnybank Hills: iConcept Press, 2012, v. 1, p. 1-16.*
  Due to the paper presented in Section 3.1, authors have been invited to present their shading attenuation method also as a book chapter. Beyond the experiments with melanocytic skin lesions and face images, this publication also includes experiments with hand images (usually utilized in gesture recognition systems).

The author of this Thesis also was part of other research projects while working towards the Doctoral Degree, collaborating with other colleagues of the research group. These collaborations resulted in three publications:

- *Yari, Y. ; Cavalcanti, P. G. ; Scharcanski, J. . Estimation of the Head Pose Based on Monocular Images. In: 25th IEEE International Conference on Image and Vision Computing New Zealand, 2010, Queenstown. Proceedings of the 25th IEEE International Conference on Image and Vision Computing New Zealand, 2010.*

- *Scharcanski, J. ; Oliveira, A. B. ; Cavalcanti, P. G. ; Yari, Y. . A Particle Filtering Approach for Vehicular Tracking Adaptive to Occlusions. IEEE Transactions on Vehicular Technology, v. 60, p. 381-389, 2011.*

- *Welfer, D. ; Scharcanski, Jacob ; Cavalcanti, Pablo G. ; Marinho, Diane R. ; Ludwig, Laura W.B. ; Kitamura, Cleyson. M. ; Dal Pizzol, Melissa M. . A Computer-Aided Diagnosis Scheme for Grading Diabetic Macular Edema Using Color Eye Fundus Images. In: M. Emre Celebi; Gerald Shaefer. (Org.). Color Medical Image Analysis. 1ed. Heidelberg: Springer, 2012, v. 1, p. 109-128.*

4 DISCUSSION

The approaches presented in Chapter 3 provide promising results for the automatic processing and analyzing of standard camera images of melanocytic skin lesions.

The proposed shading attenuation method is potentially useful in many teledermatology situations. Ideally the photographer should acquire the lesion image the best he/she can, i.e. without the presence of shading effects, but this is a very difficult task. Due to the typical small size of melanocytic lesions, the camera usually should be positioned very close to the human body and shading is a consequence. Therefore, this shading artifact may cause difficulties for the correct image segmentation. However, as we have already demonstrated, the use of the shading attenuation method before the segmentation step decreases this difficulty.

Moreover, the shading attenuation technique is not useful only for automated screening systems. Let use assume a situation when a physician wants to receive an opinion about a suspicious lesion from an expert, and sends a photograph of that lesion by email or any other Internet service. If the shading attenuation method is applied to the lesion image before or during the image transfer, it will make easier for the expert to provide his opinion about the lesion. As as indication of the usefulness of this technique, it has already attracted interest in the literature (KRUPINSKI et al., 2012; GLAISTER; WONG; CLAUSI, 2012; PICCININI et al., 2012; KOROTKOV; GARCIA, 2012; GLAISTER et al., 2013).

The segmentation procedure is the usual next step in an automated screening system, and also has been the focus of the approaches proposed in this Thesis. We proposed techniques that minimize the errors when differentiating healthy and unhealthy skin pixels. We also performed experiments comparing segmentation techniques and the influence of the segmentation result on the final classification, and our results indicate that the features usually extracted from the segmented lesion area based on the ABCD rule do not improve classification performance when the segmentation is improved (see Appendix A). However, it is important to observe that the segmentation step still is a very important task. The automatically lesion boundary determination can also be useful to help experts in remote screening since the lesion borders are not always easily recognizable. Moreover, several groups are working on melanoma diagnosis nowadays (KOROTKOV; GARCIA, 2012), and potentially new lesion features might be proposed, taking advantage of more precise lesion border location.

In terms of lesion features, we also made some research contributions. Although it requires more experiments to affirm that it can be used in practical situations, our set of features based on the ABCD rule demonstrated to be very promising, achieving high accurate levels using a simple classification algorithm, i.e. the KNN or the Nearest Neighbor Classifier. Furthermore, we also proposed a combination of this classifier with a decision tree, aiming to enhance the sensitivity (and consequently the accuracy) as much as

possible, and preventing the system to misclassify malignant lesions as benign.

Besides the set of features based on the ABCD rule, we also proposed a novel set of features based on the quantity of melanin in a lesion. The color of human skin is largely determined by the presence of these pigments, and recently studies suggest that malignant lesions may present larger melanin concentrations than in benign cases. We presented a technique to automatically estimate the melanin variation inside the lesion segments, and a two-stage framework combining the ABCD features and these melanin features. Our preliminary results indicate high levels of accuracy and sensitivity. However, as already pointed out in the discussion of the prescreening system based on the ABCD rule, we should be careful when analyzing these results. The experiments were realized in a limited number of images, and more experiments are necessary. Melanocytic skin lesions are very dangerous and refinements may be required before using these tools in real situations.

5 CONCLUSIONS

Before starting this Doctoral research project that resulted in this Thesis, the use of standard camera images to analyze melanocytic skin lesions did not receive much attention in the literature. Although the quality of these images was increasing quickly, making possible to acquire high-quality photographs with hand-held cameras and lots of other portable devices such as smartphones, most of approaches proposed in the literature in the 90s and in the early 2000s focused mostly on dermoscopic images.

However, in the recent years the Internet became an almost natural part of life, being accessible anywhere, by many different types of devices, and reaching isolated areas of the world. Consequently, medicine experts started taking advantage of this technology that enables telemedicine, making it feasible to discuss cases and obtain diagnosis from experts remotely through the exchange of data and images. Since dermatology is probably the most visual medical specialty, teledermatology has become a reality.

In addition to the benefits of teledermatology for early remote diagnosis, it has also been observed that preprocessing or an automated screening of the transferred images would be really important. An automated system like that could possibly filter obviously benign cases, and also detect severe malignant cases, creating priority attendance queues. Consequently, different research groups around the world started developing tools to help the screening of melanocytic skin lesions focusing on standard camera images, since this is usually the only available imaging device.

The tools presented in this Thesis, including two complete automated screening systems for melanocytic skin lesions, and both achieved promising results in our experiments. We proposed a technique to attenuate the shading effects commonly present in the images, and segmentation techniques that resulted in more accurate results than comparable state-of-the-art methods proposed in literature. The set of features based on the ABCD rule achieved high levels of accuracy with adequate training and using simple classifiers. Moreover, to the best of our knowledge the proposed melanin features were the first proposed to estimate important melanocytic lesions characteristic using only standard camera images.

As future work, it would be important to test all the proposed tools more extensively in clinical trials. Although our preliminary results indicate high level of accuracy, we are dealing with a very dangerous kind of cancer and before using these tools in real situations caution is necessary.

Besides more experimental work, other future works can also be mentioned. For instance, the Pattern Analysis methodology (PEHAMBERGER; STEINER; WOLFF, 1987) suggests using a type of global texture pattern (reticular, globular, homogeneous, etc) as an important feature for the lesion diagnosis. However, the texture analysis approach has not been sufficiently explored for lesion images in the literature.

Still considering the feature extraction step, the color variation representation also did not receive sufficient attention. One of the reasons is that color variation description is quite of subjective in the ABCD rule. For instance, it defines that melanocytic lesion may contain six distinct hues: white, red, light and dark brown, blue-gray, and black. However, to differentiate light and dark brown (and maybe even black) is a very subjective issue. Moreover, there is no standard when acquiring the lesion images, and consequently the ambient light may significantly influence the resulting colors. Consequently, color calibration can be an important research topic to help melanoma identification.

Finally, the classification step may also be a potential focus of research. Machine learning and pattern recognition are research areas in constant development, and many feature selectors, resampling and classifier algorithms are being proposed. It is possible that some of these untested techniques capture well the universe of the melanocytic skin lesions.

REFERENCES

AKSOY, S.; HARALICK, R. M. Probabilistic vs. geometric similarity measures for image retrieval. In: IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, 2000. **Proceedings...** [S.l.: s.n.], 2000. v.2, p.357–362.

ALCON, J. F.; CIUHU, C.; KATE, W. ten; HEINRICH, A.; UZUNBAJAKAVA, N.; KREKELS, G.; SIEM, D.; HAAN, G. de. Automatic Imaging System With Decision Support for Inspection of Pigmented Skin Lesions and Melanoma Diagnosis. **IEEE Journal of Selected Topics in Signal Processing**, [S.l.], v.3, n.1, p.14–25, Feb. 2009.

ALPAYDIN, E. **Introduction to Machine Learning (Adaptive Computation and Machine Learning)**. [S.l.]: The MIT Press, 2004.

ARGENZIANO, G.; ZALAUDEK, I.; SOYER, H. Which is the most reliable method for teaching dermoscopy for melanoma diagnosis to residents in dermatology? **British Journal of Dermatology**, [S.l.], v.151, p.512–3, 2004.

BAKOS, L.; WAGNER, M.; BAKOS, R. M.; LEITE, C. S. M.; SPERHACKE, C. L.; DZEKANIAK, K. S.; GLEISNER, A. L. M. Sunburn, sunscreens, and phenotypes: some risk factors for cutaneous melanoma in southern brazil. **International Journal of Dermatology**, [S.l.], v.41, n.9, p.557–562, 2002.

BALLERINI, L.; FISHER, R. B.; ALDRIDGE, B.; REES, J. A Color and Texture Based Hierarchical K-NN Approach to the Classification of Non-Melanoma Skin Lesions. In: CELEBI, M. E.; SCHAEFER, G. (Ed.). **Color Medical Image Analysis**. [S.l.]: Springer, 2012. (Lecture Notes in Computational Vision and Biomechanics, v.6).

BALLERINI, L.; FISHER, R. B.; ALDRIDGE, B.; REES, J. Non-melanoma skin lesion classification using colour image data in a hierarchical K-NN classifier. In: ISBI, 2012. **Anais...** [S.l.: s.n.], 2012. p.358–361.

BARANOSKI, G.; KRISHNASWAMY, A. **Light & Skin Interactions**: simulations for computer graphics applications. [S.l.]: Morgan Kaufmann, 2010.

BAXTER, A. J.; HUGHES, M. C.; KVASKOFF, M.; SISKIND, V.; SHEKAR, S.; AITKEN, J. F.; GREEN, A. C.; DUFFY, D. L.; HAYWARD, N. K.; MARTIN, N. G.; WHITEMAN, D. C. The Queensland Study of Melanoma: environmental and genetic associations (q-mega); study design, baseline characteristics, and repeatability of phenotype and sun exposure measures. **Twin Res Hum Genet**, [S.l.], v.11, p.1832–4274, 2008.

BISHOP, C. M. **Pattern Recognition and Machine Learning (Information Science and Statistics)**. 1.ed. [S.l.]: Springer, 2007.

BLUM, A.; ZALAUDEK, I.; ARGENZIANO, G. Digital image analysis for diagnosis of skin tumors. **Semin Cutan Med Surg**, [S.l.], v.27, n.1, p.11–15, Mar 2008.

BONFÁ, R.; BONAMIGO, R. R.; BONFÁ, R.; DURO, K. M.; FURIAN, R. D.; ZELMANOWICZ, A. d. M. A precocidade diagnóstica do melanoma cutâneo: uma observação no sul do brasil. **Anais Brasileiros de Dermatologia**, [S.l.], v.86, p.215 – 221, 04 2011.

BREIMAN, L.; FRIEDMAN, J. H.; OLSHEN, R. A.; STONE, C. J. **Classification and Regression Trees**. [S.l.]: Chapman & Hall, New York, NY, 1984.

BRULS, W.; LEUN, J. van der. Forward Scattering Properties of Human Epidermal Layers. **Photochem. Photobiol.**, [S.l.], v.40, p.231–242, 1984.

BURRONI, M.; CORONA, R.; DELL'EVA, G.; SERA, F.; BONO, R.; PUDDU, P.; PEROTTI, R.; NOBILE, F.; ANDREASSI, L.; RUBEGNI, P. Melanoma Computer-Aided Diagnosis: reliability and feasibility study. **Clinical Cancer Research**, [S.l.], v.10, p.1881–1886, 2004.

CAVALCANTI, P. G.; SCHARCANSKI, J. Automated prescreening of pigmented skin lesions using standard cameras. **Computerized Medical Imaging and Graphics**, [S.l.], v.35, n.6, p.481 – 491, 2011.

CAVALCANTI, P. G.; SCHARCANSKI, J.; PERSIA, L. E. D.; MILONE, D. H. **An ICA-Based Method for the Segmentation of Pigmented Skin Lesions in Macroscopic Images**. 2011. -p.

CAVALCANTI, P.; SCHARCANSKI, J.; LOPES, C. Shading Attenuation in Human Skin Color Images. In: BEBIS, G.; BOYLE, R.; PARVIN, B.; KORACIN, D.; CHUNG, R.; HAMMOUD, R.; HUSSAIN, M.; KAR-HAN, T.; CRAWFIS, R.; THALMANN, D.; KAO, D.; AVILA, L. (Ed.). **Advances in Visual Computing**. [S.l.]: Springer Berlin Heidelberg, 2010. p.190–198. (Lecture Notes in Computer Science, v.6453).

CAVALCANTI, P.; YARI, Y.; SCHARCANSKI, J. Pigmented skin lesion segmentation on macroscopic images. In: IMAGE AND VISION COMPUTING NEW ZEALAND (IVCNZ), 2010 25TH INTERNATIONAL CONFERENCE OF, 2010. **Anais...** [S.l.: s.n.], 2010. p.1 –7.

CELEBI, M. E.; ASL, Y. A.; STOECKER, W. V.; IYATOMI, H.; OKA, H.; CHEN, X. Unsupervised border detection in dermoscopy images. **Skin Res. Technol.**, [S.l.], 2007.

CELEBI, M. E.; KINGRAVI, H. A.; IYATOMI, H.; ASLANDOGAN, Y. A.; STOECKER, W. V.; MOSS, R. H.; MALTERS, J. M.; GRICHNIK, J. M.; MARGHOOB, A. A.; RABINOVITZ, H. S.; MENZIES, S. W. Border detection in dermoscopy images using statistical region merging. **Skin Res. Technol.**, [S.l.], v.14, n.3, p.347–353, Aug 2008.

CELEBI, M. E.; KINGRAVI, H. A.; UDDIN, B.; IYATOMI, H.; ASLANDOGAN, Y. A.; STOECKER, W. V.; MOSS, R. H. A methodological approach to the classification of dermoscopy images. **Computerized Medical Imaging and Graphics**, [S.l.], v.31, n.6, p.362 – 373, 2007.

CELEBI, M.; IYATOMI, H.; SCHAEFER, G.; STOECKER, W. V. Lesion border detection in dermoscopy images. **Computerized Medical Imaging and Graphics**, [S.l.], v.33, n.2, p.148 – 153, 2009.

CHAN, T. F.; SANDBERG, B. Y.; VESE, L. A. Active Contours without Edges for Vector-Valued Images. **Journal of Visual Communication and Image Representation**, [S.l.], v.11, n.2, p.130 – 141, 2000.

Commission Internationale de l'Eclairage. **Selected Colorimetric Tables**. Disponível em: <http://www.cie.co.at>. Acesso em: Novembro 2011.

CUST, A. E.; SCHMID, H.; MASKIELL, J. A.; JETANN, J.; FERGUSON, M.; HOLLAND, E. A.; AGHA-HAMILTON, C.; JENKINS, M. A.; KELLY, J.; KEFFORD, R. F.; GILES, G. G.; ARMSTRONG, B. K.; AITKEN, J. F.; HOPPER, J. L.; MANN, G. J. Population-based, case-control-family design to investigate genetic and environmental influences on melanoma risk: australian melanoma family study. **Am J Epidemiol**, [S.l.], v.170, p.1541–54, 2009.

DATASUS. **Taxa de incidência de neoplasias malignas**. Disponível em: <http://tabnet.datasus.gov.br/cgi/idb2009/d05.htm>. Acesso em: Dezembro 2011.

DERMNET. **Skin Disease Image Atlas**. Disponível em: <http://www.dermnet.com>. Acesso em: Março 2010.

DERMQUEST.COM. **The Art, Science and Practice of Dermatology**. Disponível em: <http://www.dermquest.com>. Acesso em: Abril 2010.

DIMITROW, E.; RIEMANN, I.; EHLERS, A.; KOEHLER, M.; NORGAUER, J.; ELSNER, P.; KÖNIG, K.; KAATZ, M. Spectral fluorescence lifetime detection and selective melanin imaging by multiphoton laser tomography for melanoma diagnosis. **Experimental Dermatology**, [S.l.], v.18, p.509–515, 2009.

DREISEITL, S.; OHNO-MACHADO, L.; KITTLER, H.; VINTERBO, S.; BILLHARDT, H.; BINDER, M. A comparison of machine learning methods for the diagnosis of pigmented skin lesions. **J Biomed Inform**, [S.l.], v.34, n.1, p.28–36, Feb 2001.

EHRSAM, E. **Dermoscopy**. Disponível: <http://dermoscopic.blogspot.com>. Acesso em: Julho 2010.

FIGUEIREDO, M. A. F.; JAIN, A. K. Unsupervised learning of finite mixture models. **IEEE Trans. on Pattern Analysis and Machine Intelligence**, [S.l.], v.24, n.3, p.381–396, Mar. 2002.

FIKRLE, T.; PIZINGER, K. Digital computer analysis of dermatoscopical images of 260 melanocytic skin lesions; perimeter/area ratio for the differentiation between malignant melanomas and melanocytic nevi. **J Eur Acad Dermatol Venereol**, [S.l.], v.21, n.1, p.48–55, Jan 2007.

FREUND, Y.; SCHAPIRE, R. E. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. **Journal of Computer and System Sciences**, [S.l.], v.55, n.1, p.119–139, 1997.

FRIEDMAN, R. J.; RIGEL, D. S.; KOPF, A. W. Early detection of malignant melanoma: the role of physician examination and self-examination of the skin. **CA Cancer J. Clin.**, [S.l.], v.35, n.3, p.130–51, May-Ju 1985.

GANSTER, H.; PINZ, A.; RöHRER, R.; WILDLING, E.; BINDER, M.; KITTLER, H. Automated Melanoma Recognition. **IEEE Transactions on Medical Imaging**, [S.l.], v.20, p.233–239, 2001.

GARNAVI, R.; ALDEEN, M.; CELEBI, M. E.; VARIGOS, G.; FINCH, S. Border detection in dermoscopy images using hybrid thresholding on optimized color channels. **Computerized Medical Imaging and Graphics**, [S.l.], v.35, n.2, p.105 – 115, 2011. Advances in Skin Cancer Image Analysis.

GLAISTER, J.; AMELARD, R.; WONG, A.; CLAUSI, D. MSIM: multi-stage illumination modeling of dermatological photographs for illumination-corrected skin lesion analysis. **Biomedical Engineering, IEEE Transactions on**, [S.l.], v.PP, n.99, p.1, 2013.

GLAISTER, J.; WONG, A.; CLAUSI, D. A. Illumination correction in dermatological photographs using multi-stage illumination modeling for skin lesion analysis. In: ENGINEERING IN MEDICINE AND BIOLOGY SOCIETY (EMBC), 2012 ANNUAL INTERNATIONAL CONFERENCE OF THE IEEE, 2012. **Anais...** [S.l.: s.n.], 2012. p.102–105.

GOMEZ, D. D.; BUTAKOFF, C.; ERSBOLL, B. K.; STOECKER, W. Independent Histogram Pursuit for Segmentation of Skin Lesions. **IEEE Transactions on Biomedical Engineering**, [S.l.], v.55, n.1, p.157–161, Jan. 2008.

GOURIER, N.; HALL, D.; CROWLEY, J. L. Estimating Face Orientation from Robust Detection of Salient Facial Features. In: POINTING 2004, ICPR, INTERNATIONAL WORKSHOP ON VISUAL OBSERVATION OF DEICTIC GESTURES, 2004, Cambridge, UK. **Proceedings...** [S.l.: s.n.], 2004.

GUTKOWICZ-KRUSIN, D.; RABINOVITZ, H. Multispectral Image Analysis. In: SOYER, H. P.; ARGENZIANO, G.; HOFMANN-WELLENHOF, R.; JOHR, R. H. (Ed.). **Color Atlas of Melanocytic Lesions of the Skin**. [S.l.]: Springer Berlin Heidelberg, 2007. p.52–56.

HOWLADER, N.; NOONE, A.; KRAPCHO, M.; NEYMAN, N.; AMINOU, R.; ALTEKRUSE, S.; KOSARY, C.; RUHL, J.; TATALOVICH, Z.; CHO, H.; MARIOTTO, A.; EISNER, M.; LEWIS, D.; CHEN, H.; FEUER, E.; CRONIN, K. **SEER Cancer Statistics Review, 1975-2009 (Vintage 2009 Populations), National Cancer Institute**. Disponível em: <http://seer.cancer.gov/csr/1975_2009_pops09/>. Acesso em: Julho 2012.

HYVÄRINEN, A. Fast and robust fixed-point algorithms for independent component analysis. **IEEE Transactions on Neural Networks**, [S.l.], v.10, n.3, p.626–634, 1999.

HYVäRINEN, A.; KARHUNEN, J.; OJA, E. **Independent Component Analysis**. [S.l.]: John Wiley & Sons, Inc., 2001.

IYATOMI, H.; OKA, H.; CELEBI, M.; HASHIMOTO, M.; HAGIWARA, M.; TANAKA, M.; OGAWA, K. An improved Internet-based melanoma screening system with dermatologist-like tumor area extraction algorithm. **Computerized Medical Imaging and Graphics**, [S.l.], v.32, n.7, p.566 – 579, 2008.

JIMBOW, K.; RESZKA, K.; SCHMITZ, S.; SALOPEK, T.; THOMAS, P. Distribution of ue- and pheomelanins in human skin and melanocytic tumors, and their photoprotective vs phototoxic properties. In: MELANIN: ITS ROLE IN HUMAN PHOTOPROTECTION, 1995, Overland Park, Kansas, USA. **Anais...** Valdenmar Publishing Company, 1995. p.165–175.

JOHR, R. H. Dermoscopy: alternative melanocytic algorithms–the abcd rule of dermatoscopy, menzies scoring method, and 7-point checklist. **Clinics in Dermatology**, [S.l.], v.20, n.3, p.240 – 247, 2002.

KOPF, A. W. Clinical Examination of Melanocytic Neoplasms Including ABCDE Criteria. In: SOYER, H. P.; ARGENZIANO, G.; HOFMANN-WELLENHOF, R.; JOHR, R. H. (Ed.). **Color Atlas of Melanocytic Lesions of the Skin**. [S.l.]: Springer Berlin Heidelberg, 2007. p.3–6.

KOROTKOV, K.; GARCIA, R. Computerized analysis of pigmented skin lesions: a review. **Artificial Intelligence in Medicine**, [S.l.], v.56, n.2, p.69 – 90, 2012.

KRISHNASWAMY, A.; BARANOSKI, G. A Biophysically-Based Spectral Model of Light Interaction with Human Skin. **Computer Graphics Forum**, [S.l.], v.23, n.3, p.331–340, 2004.

KRUPINSKI, E. A.; SILVERSTEIN, L. D.; HASHMI, S. F.; GRAHAM, A. R.; WEINSTEIN, R. S.; ROEHRIG, H. Observer Performance Using Virtual Pathology Slides: impact of lcd color reproduction accuracy. **Journal of Digital Imaging**, [S.l.], p.1–6, 2012.

LAZOVA, R.; PAWELEK, J. Why do melanomas get so dark? **Experimental Dermatology**, [S.l.], v.18, p.934–938, 2009.

LEE, H.-C. **Introduction to Color Imaging Science**. New York, NY, USA: Cambridge University Press, 2005.

MAGLOGIANNIS, I.; DOUKAS, C. Overview of Advanced Computer Vision Systems for Skin Lesions Characterization. **Information Technology in Biomedicine, IEEE Transactions on**, [S.l.], v.13, n.5, p.721 –733, sep. 2009.

MANOUSAKI, A. G.; MANIOS, A. G.; TSOMPANAKI, E. I.; PANAYIOTIDES, J. G.; TSIFTSIS, D. D.; KOSTAKI, A. K.; TOSCA, A. D. A simple digital image processing system to aid in melanoma diagnosis in an everyday melanocytic skin lesion unit: a preliminary report. **Int J Dermatol**, [S.l.], v.45, n.4, p.402–410, Apr 2006.

MARCHESINI, R.; BONO, A.; CARRARA, M. In vivo characterization of melanin in melanocytic lesions: spectroscopic study on 1671 pigmented skin lesions. **Journal of Biomedical Optics**, [S.l.], v.14, n.1, p.014027, 2009.

MASSONE, C.; WURM, E. M. T.; HOFMANN-WELLENHOF, R.; SOYER, H. P. Teledermatology: an update. **Semin. Cutan. Med. Surg.**, [S.l.], v.27, n.1, p.101–105, Mar 2008.

MATTHEWS, T. E.; PILETIC, I. R.; SELIM, M. A.; SIMPSON, M. J.; WARREN, W. S. Pump-Probe Imaging Differentiates Melanoma from Melanocytic Nevi. **Science Translational Medicine**, [S.l.], v.3, n.71, p.71ra15, 2011.

MAYER, J. Systematic review of the diagnostic accuracy of dermatoscopy in detecting malignant melanoma. **The Medical Journal of Australia**, [S.l.], v.167, p.206–210, 1997.

MCCARTNEY, E. **Optics of the Atmosphere**: scattering by molecules and particles. [S.l.]: John Wiley & Sons Inc., 1976.

Melanoma Research Project. **http**://www.melresproj.com. 2010.

MELLI, R.; GRANA, C.; CUCCHIARA, R. **Comparison of color clustering algorithms for segmentation of dermatological images**. [S.l.]: SPIE, 2006. v.6144, n.1.

NACHBAR, F.; STOLZ, W.; MERKLE, T.; COGNETTA, A. B.; VOGT, T.; LANDTHALER, M.; BILEK, P.; BRAUN-FALCO, O.; PLEWIG, G. The ABCD rule of dermatoscopy: high prospective value in the diagnosis of doubtful melanocytic skin lesions. **Journal of the American Academy of Dermatology**, [S.l.], v.30, n.4, p.551 – 559, 1994.

Natural Phenomena Simulation Group. **Run BioSpec Online**. Disponível em: <http://www.npsg.uwaterloo.ca/models/biospec.php>. Acesso em: Julho 2011.

OTSU, N. A threshold selection method from gray-level histograms. **IEEE Transactions on Systems, Man and Cybernetics**, [S.l.], v.9, n.1, p.62–66, January 1979.

PAROLIN, A.; HERZER, E.; JUNG, C. R. Semi-automated Diagnosis of Melanoma through the Analysis of Dermatological Images. In: VALDENMAR PUBLISHING COMPANY, 2010, Los Alamitos, CA, USA. **Anais...** IEEE Computer Society, 2010. v.0, p.71–78.

PEHAMBERGER, H.; STEINER, A.; WOLFF, K. In vivo epiluminescence microscopy of pigmented skin lesions. I. Pattern analysis of pigmented skin lesions. **Journal of the American Academy of Dermatology**, [S.l.], v.17, p.571–583, 1987.

PEREZ-GOMEZ, B.; POLLï£¡N, M.; GUSTAVSSON, P.; PLATO, N.; ARAGONï£¡S, N.; Lï£¡PEZ-ABENTE, G. Cutaneous melanoma: hints from occupational risks by anatomic site in swedish men. **Occup Environ Med**, [S.l.], v.61, p.117–26, 2004.

PICCININI, F.; LUCARELLI, E.; GHERARDI, A.; BEVILACQUA, A. Multi-image based method to correct vignetting effect in light microscopy images. **Journal of Microscopy**, [S.l.], 2012.

RAO, B. K.; MARGHOOB, A. A.; STOLZ, W.; KOPF, A. W.; SLADE, J.; WASTI, Q.; SCHOENBACH, S. P.; DE-DAVID, M.; BART, R. S. Can early malignant melanoma be differentiated from atypical melanocytic nevi by in vivo techniques? **Skin Research and Technology**, [S.l.], v.3, n.1, p.8–14, 1997.

RUIZ, D.; BERENGUER, V. J.; SORIANO, A.; MARTIN, J. A cooperative approach for the diagnosis of the melanoma. **Conf Proc IEEE Eng Med Biol Soc**, [S.l.], v.2008, p.5144–5147, 2008.

RUIZ, D.; BERENGUER, V.; SORIANO, A.; SÃąNCHEZ, B. A decision support system for the diagnosis of melanoma: a comparative approach. **Expert Systems with Applications**, [S.l.], v.38, n.12, p.15217 – 15223, 2011.

SCHINDEWOLF, T.; SCHIFFNER, R.; STOLZ, W.; ALBERT, R.; ABMAYR, W.; HARMS, H. Evaluation of different image acquisition techniques for a computer vision system in the diagnosis of malignant melanoma. **Journal of the American Academy of Dermatology**, [S.l.], v.31, n.1, p.33 – 41, 1994.

SEIDENARI, S.; PELLACANI, G.; PEPE, P. Digital videomicroscopy improves diagnostic accuracy for melanoma. **Journal of the American Academy of Dermatology**, [S.l.], v.39, n.2, p.175 – 181, 1998.

SHAPIRO, L.; STOCKMAN, G. **Computer Vision**. [S.l.]: Prentice Hall, 2001. 468-72p.

SILVEIRA, M.; NASCIMENTO, J.; MARQUES, J.; MARCAL, A.; MENDONCA, T.; YAMAUCHI, S.; MAEDA, J.; ROZEIRA, J. Comparison of Segmentation Methods for Melanoma Diagnosis in Dermoscopy Images. **Selected Topics in Signal Processing, IEEE Journal of**, [S.l.], v.3, n.1, p.35 –45, feb. 2009.

SKVARA, H.; TEBAN, L.; FIEBIGER, M.; BINDER, M.; KITTLER, H. Limitations of Dermoscopy in the Recognition of Melanoma. **Arch. Dermatol.**, [S.l.], v.141, p.155–160, 2005.

SMITH, A. R. Color gamut transform pairs. **SIGGRAPH Comput. Graph.**, New York, NY, USA, v.12, n.3, p.12–19, 1978.

SMOLLE, J. Automatic Diagnosis. In: SOYER, H. P.; ARGENZIANO, G.; HOFMANN-WELLENHOF, R.; JOHR, R. H. (Ed.). **Color Atlas of Melanocytic Lesions of the Skin**. [S.l.]: Springer Berlin Heidelberg, 2007. p.47–51.

SOILLE, P. Morphological Operators. In: JÄHNE, B.; HAUSSECKER, H.; GEISSLER, P. (Ed.). **Handbook of Computer Vision and Applications**. San Diego: Academic Press, 1999. v.2, p.627–682.

TABATABAIE, K.; ESTEKI, A.; TOOSSI, P. Extraction of skin lesion texture features based on independent component analysis. **Skin Res Technol**, [S.l.], v.15, n.4, p.433–439, Nov 2009.

TAN, X.; TRIGGS, B. Enhanced Local Texture Feature Sets for Face Recognition Under Difficult Lighting Conditions. **Image Processing, IEEE Transactions on**, [S.l.], v.19, n.6, p.1635 –1650, june 2010.

TANG, J. A multi-direction GVF snake for the segmentation of skin cancer images. **Pattern Recogn.**, New York, NY, USA, v.42, p.1172–1179, June 2009.

The Skin Site. **Atypical Moles**. Disponível em: <http://www.skinsite.com/info_atypical_moles.htm>. Acesso em: Julho 2010.

TROWBRIDGE, T.; REITZ, K. Average Irregularity representation of a rough surface for ray reflection. **Journal of the Optical Society of America**, [S.l.], v.65, n.5, p.531–536, May 1975.

TUCHIN, V. **Tissue Optics Light Scattering Methods and Instruments for Medical Diagnosis**. Bellingham, WA, USA: The International Society for Optical Engineering, 2000.

VASSILI, V. V.; SAZONOV, V.; ANDREEVA, A. A Survey on Pixel-Based Skin Color Detection Techniques. In: GRAPHICON-2003, 2003. **Proceedings...** [S.l.: s.n.], 2003. p.85–92.

WADHAWAN, T.; SITU, N.; RUI, H.; LANCASTER, K.; YUAN, X.; ZOURIDAKIS, G. **Implementation of the 7-Point Checklist for Melanoma Detection on Smart Handheld Devices**. 2011.

WANG, H.; MOSS, R. H.; CHEN, X.; STANLEY, R. J.; STOECKER, W. V.; CELEBI, M. E.; MALTERS, J. M.; GRICHNIK, J. M.; MARGHOOB, A. A.; RABINOVITZ, H. S.; MENZIES, S. W.; SZALAPSKI, T. M. Modified watershed technique and post-processing for segmentation of skin lesions in dermoscopy images. **Computerized Medical Imaging and Graphics**, [S.l.], v.35, n.2, p.116 – 120, 2011. <ce:title>Advances in Skin Cancer Image Analysis</ce:title>.

WHITED, J. D. Teledermatology research review. **Int. J. Dermatol.**, [S.l.], v.45, n.3, p.220–229, Mar 2006.

WIGHTON, P.; LEE, T. K.; LUI, H.; MCLEAN, D. I.; ATKINS, M. S. Generalizing common tasks in automated skin lesion diagnosis. **IEEE Transactions on Information Technology in BioMedicine**, [S.l.], p.622–629, 2011.

WIGHTON, P.; SADEGHI, M.; LEE, T. K.; ATKINS, M. S. A Fully Automatic Random Walker Segmentation for Skin Lesions in a Supervised Setting. In: INTERNATIONAL CONFERENCE ON MEDICAL IMAGE COMPUTING AND COMPUTER-ASSISTED INTERVENTION: PART II, 12., 2009, Berlin, Heidelberg. **Proceedings...** Springer-Verlag, 2009. p.1108–1115. (MICCAI '09).

WONG, A.; SCHARCANSKI, J.; FIEGUTH, P. Automatic Skin Lesion Segmentation via Iterative Stochastic Region Merging. **Information Technology in Biomedicine, IEEE Transactions on**, [S.l.], v.15, n.6, p.929 –936, nov. 2011.

World Health Organization. **How commom is skin cancer?** Disponível em: <http://www.who.int/uv/faq/skincancer/en/index1.html>. Acesso em: Novembro 2011.

WURM, E. M.; SOYER, H. P. Scanning for melanoma. **Australian Prescriber**, [S.l.], v.33, p.150–5, 2010.

XU, C.; PRINCE, J. L. Snakes, Shapes, and Gradient Vector Flow. **IEEE TRANSACTIONS ON IMAGE PROCESSING**, [S.l.], v.7, n.3, p.359–369, 1998.

YAN, F.; ZHANG, H.; KUBE, C. R. A multistage adaptive thresholding method. **Pattern Recogn. Lett.**, New York, NY, USA, v.26, p.1183–1191, June 2005.

YOUNG, G. A. Alternative Smoothed Bootstraps. **Journal of the Royal Statistical Society. Series B (Methodological)**, [S.l.], v.52, n.3, p.477–484, 1990.

ZHOU, H.; MILLER, P.; ZHANG, J. Age classification using Radon transform and entropy based scaling SVM. In: BRITISH MACHINE VISION CONFERENCE, 2011. **Proceedings...** BMVA Press, 2011. p.28.1–28.12. http://dx.doi.org/10.5244/C.25.28.

ZHOU, H.; SCHAEFER, G.; CELEBI, M. E.; LIN, F.; LIU, T. Gradient vector flow with mean shift for skin lesion segmentation. **Computerized Medical Imaging and Graphics**, [S.l.], v.35, n.2, p.121 – 127, 2011. <ce:title>Advances in Skin Cancer Image Analysis</ce:title>.

ZHOU, H.; SCHAEFER, G.; SADKA, A.; CELEBI, M. Anisotropic Mean Shift Based Fuzzy C-Means Segmentation of Dermoscopy Images. **Selected Topics in Signal Processing, IEEE Journal of**, [S.l.], v.3, n.1, p.26 –34, feb. 2009.

ZONIOS, G.; DIMOU, A.; BASSUKAS, I.; GALARIS, D.; TSOLAKIDIS, A.; KAXIRAS, E. Melanin absorption spectroscopy: new method for noninvasive skin investigation and melanoma detection. **Journal of Biomedical Optics**, [S.l.], v.13, n.1, p.014017–1–014017–8, 2008.

ZONIOS, G.; DIMOU, A.; CARRARA, M.; MARCHESINI, R. In Vivo Optical Properties of Melanocytic Skin Lesions: common nevi, dysplastic nevi and malignant melanoma. **Photochemistry and Photobiology**, [S.l.], v.86, n.1, p.236–240, 2010.

APPENDIX A MACROSCOPIC PIGMENTED SKIN LESION SEGMENTATION AND ITS INFLUENCE ON THE LESION CLASSIFICATION AND DIAGNOSIS

Cavalcanti, Pablo G. ; Scharcanski, Jacob . Macroscopic Pigmented Skin Lesion Segmentation and Its Influence on the Lesion Classification and Diagnosis. In: M. Emre Celebi; Gerald Shaefer. (Org.). Color Medical Image Analysis. 1ed. Heidelberg: Springer, 2012, v. 1, p. 15-39.

A.1 Abstract

Melanoma is a type of malignant pigmented skin lesion, and currently is among the most dangerous existing cancers. However, differentiating malignant and benign cases is a hard task even for experienced specialists, and a computer-aided diagnosis system can be an useful tool. Usually, the system starts by pre-processing the image, i.e. removing undesired artifacts such as hair, freckles or shading effects. Next, the system performs a segmentation step to identify the lesion boundaries. Finally, based on the image area identified as lesion, several features are computed and a classification is provided. In this chapter we describe all these steps, giving special attention to segmentation approaches for pigmented skin lesions, proposed for standard camera images (i.e. simple color photographs). Next, we compare the segmentation results to identify which techniques have more accurate results, and discuss how these results may influence in the following steps: the feature extraction and the final lesion classification.

A.2 Introduction

Pigmented skin lesions include both, benign and malignant forms. According to World Health Organization (World Health Organization, 2011), about 132000 melanoma cases, a dangerous kind of malignant pigmented skin lesion, occur globally each year. The early diagnosis of melanomas is very important for the patient prognosis, since most malignant skin lesion cases can be treated successfully in their early stages. However, research work has shown that discriminating benign from malignant skin lesions is a challenging task (RAO et al., 1997; FIKRLE; PIZINGER, 2007).

To help diagnosing pigmented skin lesions, physicians often use dermoscopy, which is a non-invasive technique that magnifies submacroscopic structures with the help of an optical lens (a dermoscope) and liquid immersion. According to Mayer (MAYER, 1997), the use of dermoscopy can increase the diagnosis sensitivity in 10-27% with respect to the clinical diagnosis. Also, several automatic segmentation and classification methods

have been proposed to help obtain a diagnosis with a dermoscopy image (CELEBI et al., 2007, 2008; GOMEZ et al., 2008; IYATOMI et al., 2008; ZHOU et al., 2009, 2011). However, even with the help of dermoscopy, differentiating malignant and benign lesions is a challenging task. In fact, specialists affirm that in the early evolution stages of malignant lesions, dermoscopy may not be helpful since it often does not improve the diagnosis accuracy (SKVARA et al., 2005).

Still considering early stage cases, there are practical situations where a non-specialist (e.g. a physician not trained on Dermatology) wishes to have a qualified opinion about a suspect skin lesion, but only standard camera imaging is available on site. In such situations, telemedicine is justifiable, and the non-specialist can capture a macroscopic pigmented skin lesion (MPSL) image of the suspect skin lesion and send it to a specialist, who can analyze it in higher detail. In this particular situation, a teledermatology consultation brings benefits, like the easier access to health care and faster clinical results (MASSONE et al., 2008). Besides, comparing the physical (face-to-face) patient diagnosis with the remote diagnosis by teledermatology, recent results suggest that teledermatology also tends to be effective and reliable (WHITED, 2006).

In the last decades, several segmentation techniques have been proposed to facilitate the remote diagnosis of MPSL images. Since there is no standardized protocol for acquiring these images, often they contain artifacts like hair, shading and other disturbances that make the remote diagnosis by specialists more difficult. With the help of the automatic segmentation, this task may be facilitated. Moreover, the segmentation is an initial step for computer-aided diagnosis systems. Starting from the lesion area identification, lesion features can be extracted and an automatic classification/diagnosis can be provided.

However, since these MPSL images may present characteristics that could make the remote diagnosis more difficult, the automatic processing and analysis also poses some challenges for the researchers in this field. Most of the MPSL image segmentation techniques proposed in the literature convert the original color image to a monochromatic image, and use a thresholding algorithm to identify the lesion area (MANOUSAKI et al., 2006; RUIZ et al., 2008; TABATABAIE; ESTEKI; TOOSSI, 2009; ALCON et al., 2009). Even more complex segmentation approaches, such as active contours techniques (PAROLIN; HERZER; JUNG, 2010; TANG, 2009), process grayscale images. Nevertheless, the discriminating lesion and healthy skin areas may be more difficult on a monochromatic images, since the chromatic aspect is lacking in them.

In the following Sections we will present segmentation algorithms that process multichannel MPSL images, based on thresholding (CAVALCANTI; SCHARCANSKI, 2011) and on level-sets (CAVALCANTI et al., 2011). After that, we show that such methods working on multichannel MPSL images are more efficient than methods working on monochromatic images. Also, we will discuss how features that are relevant from the medical point of view can be extracted, and how the final classification/diagnosis of the acquired lesion is affected by the segmentation quality.

A.3 Pre-processing

MPSL images usually contain artifacts that make the segmentation process more difficult. Skin characteristics, such as freckles, are easily detected by these algorithms based on color or size. However, most methods available try to identify the lesion area assuming that pigmented skin lesions correspond to locally darker skin discolorations. Consequently, artifacts such as hair and shading, that usually also are darker than healthy skin

may be mistaken as lesions during the segmentation process.

Although some of the approaches available eliminate hair as a pre-processing step (RUIZ et al., 2008; TABATABAIE; ESTEKI; TOOSSI, 2009; TANG, 2009), this task can be performed as a post-processing step, after the image segmentation. Hair is thin, and its shape is quite distinct from the lesion shapes, and consequently it is easy to eliminated hair from MPSL images by morphological operations or other methodologies.

On the other hand, the presence of shading requires pre-processing in advance to segmentation. The shading areas assume any shape, and require a treatment that is different from that given to artifacts like hair. Moreover, if the shading attenuation is well performed, it also contributes for the enhance the contrast between healthy and unhealthy skin.

A.3.1 Shading attenuation

Alcón et al. (ALCON et al., 2009) proposed to correct the uneven illumination by removing the low frequency spatial component of the image. Although this method can be efficient for some images, it requires specific parameters. It is very difficult to obtain a specific value that can be used for any input image, and the authors (ALCON et al., 2009) do not detail how to obtain this value automatically.

Face images also are skin images and can be affected by shading effects. Tan and Triggs (TAN; TRIGGS, 2010) and Zhou et al. (ZHOU; MILLER; ZHANG, 2011) proposed to use Difference of Gaussians (DoG) filtering to correct the shading artifacts. However, this methodology needs specific parameters (e.g., definition of a window size and the filters standard deviations) which may require adjustments for different types of shading effects in the skin lesion images. Moreover, the authors observed that DoG filtering may generate strong-edges in hair areas (TAN; TRIGGS, 2010), which could affect negatively the overall segmentation process.

Therefore, Cavalcanti et al. (CAVALCANTI et al., 2011) proposed a shading attenuation method that is adaptive to the MPSL image data. Their method assumes that images are acquired in a way that the lesion appears in the image center, and it does not touch the image outer borders. The first step of the method is to convert the image from the original RGB color space to the HSV color space, and retain the Value channel V. This is justified by the fact that this channel presents the higher visibility of the shading effects. A region of 20×20 pixels is extracted from each V corner, and the union of these four sets define the pixel set S. This pixel set is used to adjust the following quadric function $z(x, y)$:

$$z(x, y) = P_1x^2 + P_2y^2 + P_3xy + P_4x + P_5y + P_6, \tag{A.1}$$

where the six quadric function parameters P_i $(i = 1, ..., 6)$ are chosen to minimize the error ϵ:

$$\epsilon = \sum_{j=1}^{N_s} \left[V(S_{j,x}, S_{j,y}) - z(S_{j,x}, S_{j,y})\right]^2, \tag{A.2}$$

where, $S_{j,x}$ and $S_{j,y}$ are the x and y coordinates of the jth element of the set S, respectively, and N_s is the total number of pixels of the four corners (in our case, $N_s = 1600$).

Calculating the quadric function $z(x, y)$ for each image spatial location (x, y), we have an estimate $z(x, y)$ of the local illumination intensity in the image $V(x, y)$. Dividing the original $V(x, y)$ channel by $z(x, y)$, we obtain a new Value channel where the shading effects have been attenuated. The final step is to replace the original Value channel by

this new Value channel, and convert the image from the HSV color space to the original RGB color space. In Fig. A.1, an example of applying this method to a skin lesion image is presented. The result is a color image easier to be segmented.

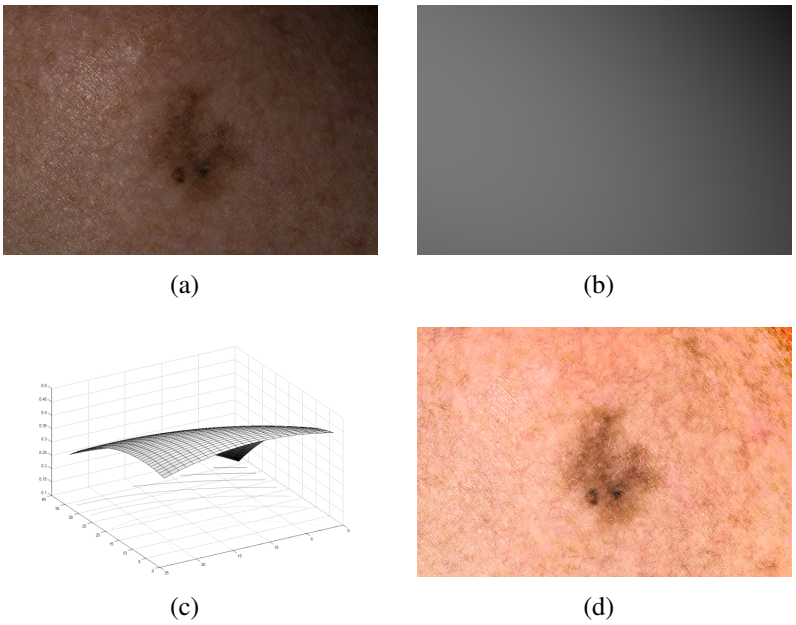

(a) (b) (c) (d)

Figure A.1: Shading attenuation example. (a) Input image; (b) Obtained quadric model using the corners of the input Value channel; (c) Obtained quadric model in 3D; (d) Result obtained by the division of the Value channel by the obtained quadric model.

A.4 Segmentation

As already mentioned in Section A.2, several segmentation techniques have been proposed for MPSL images in the last decades. We outline some representative recent methods and their characteristics in the following subsections. Also, at the end of this section, we present and discuss the performance of these segmentation techniques for a MPSL image database.

A.4.1 Grayscale-Based Methods

Based on the principle that a pigmented skin lesion is a depigmentation of the skin, and to reduce the computation cost, many segmentation methods start by converting the input image from color to grayscale. After that, most algorithms try to distinguish between healthy to unhealthy pixels. The following techniques illustrate the algorithms that have recently been used for this purpose.

A.4.1.1 Thresholding-Based Methods

Otsu's Thresholding method (OTSU, 1979) has been widely used in grayscale images (MANOUSAKI et al., 2006; RUIZ et al., 2008; TABATABAIE; ESTEKI; TOOSSI, 2009). Furthermore, Cavalcanti et al. (CAVALCANTI; YARI; SCHARCANSKI, 2010) also employed this thresholding scheme to the Red channel (R of the RGB color space), trying to take advantage of the fact that healthy skin usually has a reddish tone. This method assumes two pixel classes, namely healthy and unhealthy skin pixels, and searches

exhaustively for the threshold th that minimizes the total intra-class variance $\sigma_w^2(th)$, defined as the weighted sum of variances of the two classes:

$$\sigma_w^2(th) = \omega_1(th)\sigma_1^2(th) + \omega_2(th)\sigma_2^2(th), \tag{A.3}$$

where ω_i are the a priori probabilities of the two classes separated by the threshold th, and σ_i^2 are their intra-class variances. Minimizing the intra-class variance is equivalent to maximizing the inter-class variance $\sigma_b^2(th)$:

$$\begin{aligned} \sigma_b^2(th) &= \sigma^2 - \sigma_w^2(th) \\ &= \omega_1(th)\omega_2(th)\left[\mu_1(th) - \mu_2(th)\right]^2, \end{aligned} \tag{A.4}$$

where σ^2 is the image pixels variance, and μ_i are the class means. Computed the th threshold, the lesion pixels correspond to the pixels with values lower than th.

The Otsu's method usually is followed by a post-process step, constituted by successive morphological operations, to eliminate other regions that may be thresholded (besides the lesion). Cavalcanti et al. (CAVALCANTI; YARI; SCHARCANSKI, 2010) suggest the following procedures: select the largest threshold area, perform a hole filling operation, and a dilation with a disk with 5 pixels of radius.

However, Alcón et al. (ALCON et al., 2009) recently suggested that Otsu's method may over-segment the lesion area. So, they proposed a new thresholding method specific for MPSL images. They observed that, although the lesion intensities distribution $f_l(x)$ is unknown, the distribution $f_s(x)$ of the skin correspond to a Gaussian-like distribution:

$$f_s(x) = A\, e^{\frac{-(x-\mu_s)^2}{2\sigma_s^2}}, \tag{A.5}$$

where, μ_s is the mean value of healthy skin pixel intensities. Being f_{l+s} the distribution of grayscale intensities of the whole image, μ_s is determined by the corresponding intensity value of the highest peak of f_{l+s}. Since $f_{l+s} = f_l + f_s$, and μ_l (the mean value of lesion pixels) always is lower than μ_s, this distribution can be approximated as:

$$f_{l+s}(x) = \begin{cases} f_s(x) & , \quad x \geq \mu_s \\ f_l(x) + f_s(x) & , \quad x < \mu_s \end{cases}. \tag{A.6}$$

Therefore, based on this assumption, the skin pixels distribution can be estimated as :

$$\widetilde{f}_s(x) = \begin{cases} f_{l+s}(2\mu_s - x) & , \quad x < \mu_s \\ f_{l+s}(x) & , \quad x \geq \mu_s \end{cases}, \tag{A.7}$$

and, consequently, the lesion pixels distribution as :

$$\widetilde{f}_l(x) = f_{l+s}(x) - \widetilde{f}_s(x). \tag{A.8}$$

Finally, the means $\widetilde{E}(X_s)$ and $\widetilde{E}(X_l)$ of the distributions $\widetilde{f}_s(x)$ and $\widetilde{f}_l(x)$, respectively, are used for the computation of the threshold T as follows :

$$T = \frac{\widetilde{E}(X_s) + \widetilde{E}(X_l)}{2}, \tag{A.9}$$

and, as in the Otsu's method, the pixels with values lower than the computed threshold, are segmented as lesion pixels.

Figure A.2 illustrates the performance of the above mentioned thresholding-based segmentation methods. As can be observed, these low-complexity algorithms are able to determine the lesion area, suffering from boundary definition inaccuracies caused by hair. Also, it is important to observe that the method proposed by Alcon et al. is negatively affected by the lack of a post-processing method to eliminate undesired thresholded areas. Finally, Figure A.2(i) shows the difference between Otsu and Alcon et al. results, obtained by thresholding different regions of the histogram.

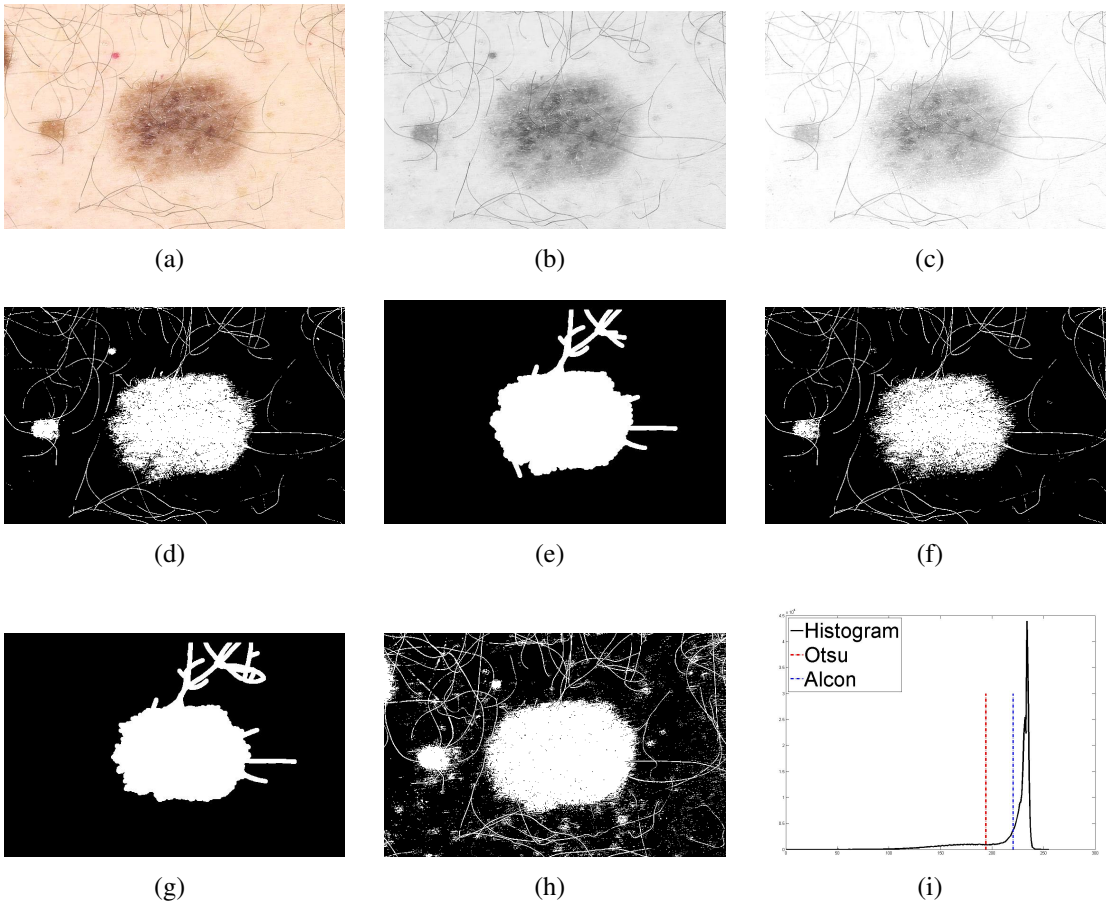

Figure A.2: MPSL image segmentation using thresholding-based methods. (a) The RGB image after pre-processing. (b) Figure (a) after conversion to grayscale. (c) The Red channel of (a). (d) The resultant binary mask of Otsu's method applied to figure (b). (e) The result of applying morphological operation in (d). (f) The resultant binary mask of Otsu's method applied to figure (c). (g) The result of applying morphological operation in (f). (h) The resultant binary mask of Alcon et al. method applied to figure (b). (i) The plot of histogram of figure (b), and the computed Otsu and Alcon thresholds.

A.4.1.2 Multi-Direction GVF Snake Method

Some researchers have proposed to use Snakes (or Active-Contours) methods to segmented MPSL images, instead of a thresholding method (TANG, 2009). These methods usually start by smoothing the image, and as an illustration of such methods we refer to the method proposed by Tang (TANG, 2009). In this method, the MPSL image is initially smoothed by adaptive anisotropic diffusion filtering. Tang modified the traditional anisotropic diffusion to make it more robust to noise, and more details can be found

in (TANG, 2009). Next, a modification of the traditional GVF (Gradient Vector Flow) snake (XU; PRINCE, 1998) is used to determine the lesion boundary. The original GVF (u, v) can be determined by the minimization of the following energy function :

$$E_{GVF}(u,v) = \begin{array}{l} \frac{1}{2}\int\int g(|\nabla f|)(u_x^2 + u_y^2 + v_x^2 + v_y^2) \\ +(1 - g(|\nabla f|))((u - f_x)^2 + (v - f_y)^2)dxdy, \end{array} \tag{A.10}$$

where, f is an edge map derived from the image, g is an edge-force magnitude :

$$g(|\nabla f|) = exp(-(\frac{|\nabla f|}{K})), \tag{A.11}$$

and K is a non-negative smoothing parameter for the field (u, v).

Tang uses a Multi-Directional GVF (MDGVF) in order to create a force-field (u, v) that enforces the snake to converge to the lesion area, and not to spurious image edges. Based on an initialization mask (containing a rough segmentation of the lesion), the author computes the lesion center $(\overline{x}, \overline{y})$ and the direction vector $\vec{d}(x, y) = (d_x, d_y)$ at each image pixel (x, y), pointing to the lesion center :

$$d_x = \frac{\overline{x} - x}{\sqrt{(\overline{x} - x)^2 + (\overline{y} - y)^2}}\,, \tag{A.12}$$

$$d_y = \frac{\overline{y} - y}{\sqrt{(\overline{x} - x)^2 + (\overline{y} - y)^2}}\,. \tag{A.13}$$

After that, the author determines the unitary vector $\vec{v}(x, y) = (v_x, v_y)$ for each pixel closer to $\vec{d}(x, y)$ by cosine similarity, i.e. $\vec{v}(x, y)$ is one of the nine vectors (-1,-1), (-1,0), (-1,1), (0,-1), (0,0), (0,1), (1,-1), (1,0) or (1,1).

Being $I(x, y)$ the grayscale intensity of a pixel (x, y), its respective directional gradient DI can be determined as follows :

$$DI(x, y) = I(x + v_x, y + v_y) - I(x, y). \tag{A.14}$$

Since often lesions are darker than healthy skin, the negative values of this gradient can be used to determine a new edge-map :

$$F(x,y) = \begin{cases} DI(x,y), & if \;\; DI(x,y) < 0 \\ 0 & otherwise \end{cases}\,. \tag{A.15}$$

Replacing f by F in Eq. A.10, a new energy function is obtained to force the snake to converge specifically along the direction of the lesion. To initialize this process, Tang suggests using the Multistage Adaptive Thresholding method (YAN; ZHANG; KUBE, 2005) to segment the image roughly. In Figure A.3, we present a typical lesion segmentation obtained using this method. The reader may observe that this algorithm handles better artifacts like hair than thresholding-based methods, but unfortunately the lesion boundary is not well determined.

A.4.2 Multichannel-Based Methods

Although grayscale images have widely being used for segmenting MPSL images, some approaches rely on multichannel images as described in the following Sections.

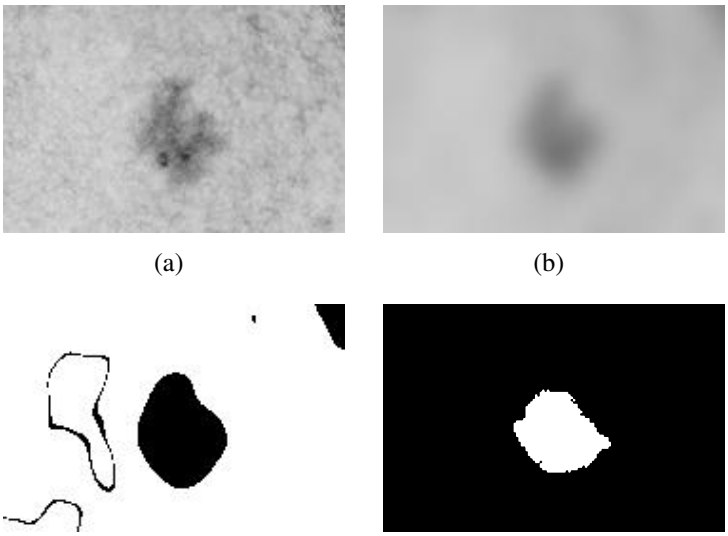

(a) (b) (c) (d)

Figure A.3: MPSL image segmentation using multi-direction GVF snake method. (a) An example of image after pre-processing and grayscale conversion. (b) Image (a) after the application of anisotropic diffusion filter. (c) The rough segmentation of (b). (d) The final segmentation after the snake convergence.

A.4.2.1 Thresholding-Based Methods

In a similar way that thresholding is used in grayscale MPSL images, we can threshold color images. In this case, thresholds are computed and a binary mask is generated for each channel, and masks are combined to form the final color MPSL image segmentation. However, working with color images adds new concerns. For example, pigmented skin lesions are not easily discriminated from healthy skin on the Green channel (G of the RGB color space), and the color information may disturb instead of benefiting the final results.

So, to facilitate the use of thresholding methods, Cavalcanti and Scharcanski (CAVALCANTI; SCHARCANSKI, 2011) recently proposed a multichannel image representation for MPSL images that maximizes the discrimination between healthy and unhealthy skin regions. The idea is to create a new 3-channel image $\bar{I}_i^N$ based on the normalization of the RGB channels $\bar{I}_i^C$ of the input image, and then use a thresholding algorithm based on Otsu's method to segment it.

The first channel is a representation of the image darkness, relying on the fact that lesion areas are depigmented skin regions. Each pixel is defined as $\bar{I}_1^N(x,y) = 1 - \bar{I}_i^C(x,y)$, i.e. the complement of the normalized Red channel.

The second channel is a texture representation, since local textural variability usually is higher in lesions than in healthy skin areas. Being $\bar{L}$ a normalized Luminance image defined by the average of the three $\bar{I}_i^C$ channels, the textural variability in $\bar{L}(x,y)$ can be quantified by computing $\tau(x,y,\sigma)$ as follows :

$$\tau(x,y,\sigma) = \bar{L}(x,y)\frac{\widetilde{S}(x,y,\sigma)}{S(x,y,\sigma)}, \tag{A.16}$$

where, $S(x,y,\sigma) = \bar{L}(x,y) * G(\sigma)$ (i.e., the Luminance image $\bar{L}$ is smoothed by a Gaussian filter with standard deviation σ), and $\widetilde{S}(x,y,\sigma)$ represents its complement. In this

way, if an image region is dark, its textural information is emphasized; if the region is bright, its textural information is de-empshasized. However, a single Gaussian filter may not be sufficient to capture the textural variability, so $\tau(x, y, \sigma)$ is calculated for different σ values and we select its maximum value at each pixel:

$$T(x, y) = \max_{\sigma}[\tau(x, y, \sigma)],\ \sigma \in \{\sigma_1, \sigma_2, ..., \sigma_N\}\,. \tag{A.17}$$

Finally, the texture variation channel T is normalized, obtaining $\bar{I}_2^N(x, y)$ as follows :

$$\bar{I}_2^N(x, y) = (T(x, y) - \min(T)) \,/\, (\max(T) - \min(T)). \tag{A.18}$$

The third channel $\bar{I}_3^N(x, y)$ of the representation describes the local color variation, assuming that healthy and unhealthy skin regions present different color distributions. The Principal Component Analysis (PCA) method is applied on the normalized colors of the image $\bar{I}_i^C(x, y)$, and the first component is retained (i.e. the component that maximizes the local data variance). In this representation, lesion pixels usually have higher variability values than healthy skin pixels, and to detect these lesion pixels in this channel (corresponding to those detected in channels $\bar{I}_1^N$ and $\bar{I}_2^N$), the following PCA property described next is utilized. Since the input data is centered around the mean, and healthy skin pixels often are more frequent in the MPSL image, the projections of the healthy skin pixels on the PCA space tend to generate values nearer to zero than the lesion pixels (i.e., the projected lesion pixels tend to have larger magnitudes, i.e. positive or negative). Therefore, the color variability information C is represented by the pixel projection magnitudes, and the normalization of C generates the $\bar{I}_3^N$ channel:

$$\bar{I}_3^N(x, y) = (C(x, y) - \min(C)) \,/\, (\max(C) - \min(C)). \tag{A.19}$$

Also, the $\bar{I}_3^N$ channel is filtered with a 5×5 median filter to reduce the noise. Obtained this multichannel representation, the Otsu's thresholding method (see Section A.4.1.1) is used to segment the image. Three thresholds th_i are computed, one for each channel $\bar{I}_i^N(x, y)$, and a pixel (x, y) is defined as part of a lesion region (i.e., $\phi(x, y) = 1$) if its value is higher than the threshold th_i in at least two of the three channels :

$$\phi(x, y) = \begin{cases} 1 \ , \ \text{if } (\bar{I}_1^N(x, y) > th_1 \ \wedge \ \bar{I}_2^N(x, y) > th_2), \\ 1 \ , \ \text{if } (\bar{I}_2^N(x, y) > th_2 \ \wedge \ \bar{I}_3^N(x, y) > th_3), \\ 1 \ , \ \text{if } (\bar{I}_1^N(x, y) > th_1 \ \wedge \ \bar{I}_3^N(x, y) > th_3), \\ 0 \ , \ \text{otherwise.} \end{cases} \tag{A.20}$$

As mentioned before (see Section A.3), after thresholding the remaining skin artifacts (such as freckles and hair) are eliminated more easily. These artifacts usually occur in isolated regions that differ in area and perimeter from skin lesions, since lesions often have larger areas and more irregular boundaries. Therefore, the perimeter and the area of all thresholded connected pixel sets (i.e. where $\phi(x, y) == 1$) are computed, and then this set of regions is partitioned in two clusters. All regions in the cluster with smaller areas (in average) are eliminated, and their correspondent mask pixels are set to $\phi(x, y) = 0$. At the end, the resultant mask is filtered by a 5×5 median filter, eliminating any possible remaining artifacts that may originate rim imperfections.

In Figure A.4, we present the results for all steps of this method, including the multichannel representation generation, the thresholding and post-processing steps. The reader may observe that the lesion boundary is determined with higher precision in comparison to the methods presented previously.

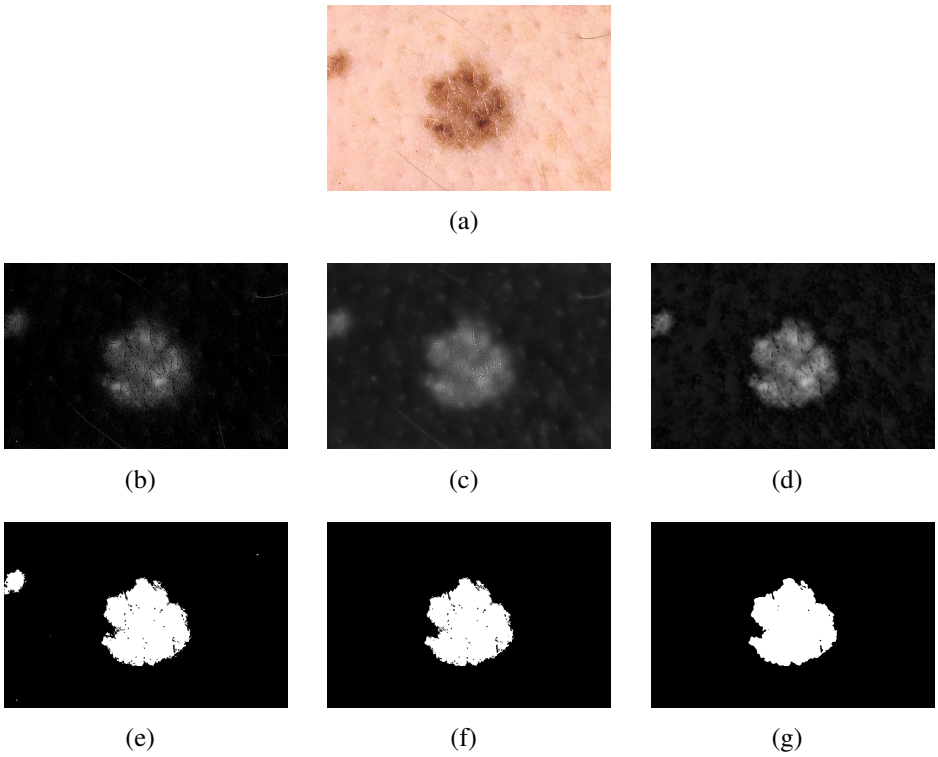

(a)

(b) (c) (d)

(e) (f) (g)

Figure A.4: Segmentation process for the image shown in (a) using thresholding method on a multichannel representation. In (b)-(d), respectively, the $\bar{I}_i^N$ channels representing darkness, texture variation and color variation. In (e)-(g), respectively, binary masks after thresholding, artifacts elimination and filtering.

A.4.2.2 ICA-Based Active-Contours Method

Instead of creating a multichannel representation for a MPSL color image, Cavalcanti et al. (CAVALCANTI et al., 2011) recently proposed a segmentation method to be used on the image original color channels. They proposed to use a classical active-contours method (Chan-Vese (CHAN; SANDBERG; VESE, 2000)), followed by morphological operations as a post-processing step.

Active-contours methods have already been used to segment pigmented skin lesion images. However, as we seen in Section A.4.1.2, usually a conversion to a grayscale image precedes the processing stages. Other approaches have been proposed using color images (SILVEIRA et al., 2009), but these algorithms were designed for dermoscopic images, which have different characteristics, and a common drawback of such methods is the difficulty to determine a convenient way to initialize the active-contours algorithm. If the initialization do not indicate the lesion regions with some accuracy, the final segmentation may be incorrect and include healthy skin regions.

Cavalcanti et al. (CAVALCANTI et al., 2011) proposed to use independent component analysis (ICA) to generate a reliable binary mask for initializing the active-contours algorithm. They observed that when ICA is applied to a MPSL image, one of the resultant ICA components corresponds mainly to the lesion area, the second component to the healthy skin, and the third component corresponds to noise artifacts. Nevertheless, there is an ordering indeterminacy inherent to the ICA method, and it is not possible to know in advance which component will show the lesion more clearly. However, due to the lesion variability, the histogram of the component that shows more clearly the lesion often

has a non-Gaussian histogram (frequently multimodal). The noise artifacts component histogram tends to be non-Gaussian, and the component that shows healthy skin more clearly tends to have a Gaussian histogram. Thus, they estimate the non-Gaussianity of the ICA histogram components with differential entropy, i.e. $\mathcal{J}(X) = |H(X) - H(X_g)|$, where X_g is a Gaussian distributed random variable with the same variance as X. The component that produces the largest differential entropy (i.e., contains the highest non-Gaussianity estimate) is identified as the one containing the lesion information more clearly, and the smallest differential entropy component carries basically healthy skin.

After reordering the channels, the lesion region is best represented in the first channel. Next, the component values are normalized in the range $[0, 1]$, and the Otsu's thresholding method is used (see Section A.4.1.1) to segment the skin lesion in this channel. Given the ICA results, the lesion information can be emphasized (closer to value 1) or deemphasized (closer to value 0) in this channel, and consequently the thresholded area may correspond to either, the lesion or the background. To guarantee that the lesion is captured in the thresholded area, the corner pixels (used in the shading attenuation step, that are known to correspond to healthy skin) are tested to check if they are thresholded as '1's or '0's. If most corner pixels are thresholded as '1's, the thresholded area corresponds to healthy skin, and the logical complement is used to obtain the lesion localization mask. In this way, a rough approximation of the lesion area is obtained. Next, the lesion boundary is better approximated and possible artifacts are eliminated with a morphological opening (i.e. the structuring element is a disk with a radius of 3 pixels).

Given this initialization binary mask, Cavalcanti et al. (CAVALCANTI et al., 2011) proposed determining the lesion boundary more precisely using the Chan-Vese Active-contours method for vector-valued images (CHAN; SANDBERG; VESE, 2000). Their method assumes that the color image I_i is formed by two regions of approximately constant intensities c_1 and c_2, separated by a curve C. The lesion localization mask is used as an initialization, indicating approximately the region to be segmented. Afterwards, the active-contours method iteratively tries to minimize the energy function $F(c_1, c_2, C)$ in the color image I_i:

$$
\begin{aligned}
F(c_1, c_2, C) = \mu\, length(C) + \\
\lambda_1 \int_{inside(C)} \frac{1}{3} \sum_{i=1}^{3} |I_i(x,y) - c_{1,i}|^2 \, dxdy + \\
\lambda_2 \int_{outside(C)} \frac{1}{3} \sum_{i=1}^{3} |I_i(x,y) - c_{2,i}|^2 \, dxdy,
\end{aligned}
$$

where μ, λ_1 and λ_2 are weighting parameters (λ_1=λ_2=1, as suggested in (CHAN; SANDBERG; VESE, 2000), and μ=0.2). Using the Level-set formulation, it is possible to minimize the energy function embedding the curve C, obtaining the zero level set $C(t) = \{(x,y)|\phi(t,x,y) = 0\}$ of a higher dimensional Level-set function $\phi(t,x,y)$. The evolution of $\phi(t,x,y)$ is given by the following motion Partial Differential Equation:

$$\frac{\partial \phi}{\partial t} = \delta_\epsilon(\phi) \left[\begin{array}{l} \mu \; div\left(\frac{\nabla \phi}{|\nabla \phi|}\right) - \\ \frac{1}{3}\sum_{i=1}^{3} \lambda_1 \left|I_i(x,y) - c_{1,i}\right|^2 + \\ \frac{1}{3}\sum_{i=1}^{3} \lambda_2 \left|I_i(x,y) - c_{2,i}\right|^2 \end{array} \right], \tag{A.21}$$

where $\delta_\epsilon(\phi)$ is the Dirac delta function, $c_{1,i}$ and $c_{2,i}$ are the averages inside and outside of the curve C in the i-th channel I_i, respectively.

It is possible that the final curve C contains regions beyond the lesion area. So, if the number of regions segmented by the Chan-Vese method is higher than one, local artifacts are eliminated. The area and the perimeter of each segmented region are computed, and these values are clustered with K-Means, where K=2. The regions in the cluster with the smaller areas (in average) are eliminated as artifacts and the other regions are kept. The regions kept are hole filled to improve their connectivity, forming the final segmentation mask. The final post-processing step is a morphological dilation (with a disk of 5 pixels of radius as the structuring element).

In Figure A.5, we present the results for all steps of this method. The reader may observe that the initialization mask already do not contain skin artifacts, forcing the active-contours method to a precise determination of the lesion area.

A.4.3 Comparison of Segmentation Methods Based on Experimental Results

In order to compare the performances of the six state-of-art segmentation methods for MPSL images presented above, we use the same image dataset as in Alcón et al. (ALCON et al., 2009), which contains 152 images that have been collected from the Dermnet dataset (DERMNET, 2010). This dataset consists of 107 melanomas and 45 Clark nevi (or atypical nevi), a benign kind of lesion that present similar characteristics to melanomas.

We implemented the shading attenuation step and all the presented algorithms, and processed all the 152 images with these implemented methods. Some examples of segmentation results can be seen in Figure A.6. We also measured the error for each resultant segmentation using the following criterion (CELEBI et al., 2008; GOMEZ et al., 2008; IYATOMI et al., 2008):

$$\epsilon = \frac{Area(Segmentation \oplus GroundTruth)}{Area(GroundTruth)} \times 100\%, \tag{A.22}$$

where, $Segmentation$ is the result of the method in test, $GroundTruth$ is the manual segmentation of the same lesion, $Area(S)$ denotes the number of pixels indicated as lesion in the segmentation result S, and $\oplus$ indicates the exclusive-OR, operation that gives the pixels for which the $Segmentation$ and $GroundTruth$ disagree.

The average error obtained by each segmentation method is presented in Table A.1. We also included a synopsis of the six segmentation methods tested in Table A.1. As can be seen, the only method that uses color information (ICA-Based Active-Contours Method) generates the smallest error segmentation, in average. Although the methods based on the Otsu's Thresholding Method are not computationally as intense as the method based on active-contours, those methods appear in the sequence, as those with the smallest segmentation errors (multichannel representation appears to be more effective than the methods based on grayscale images). The thresholding method proposed by Alcón

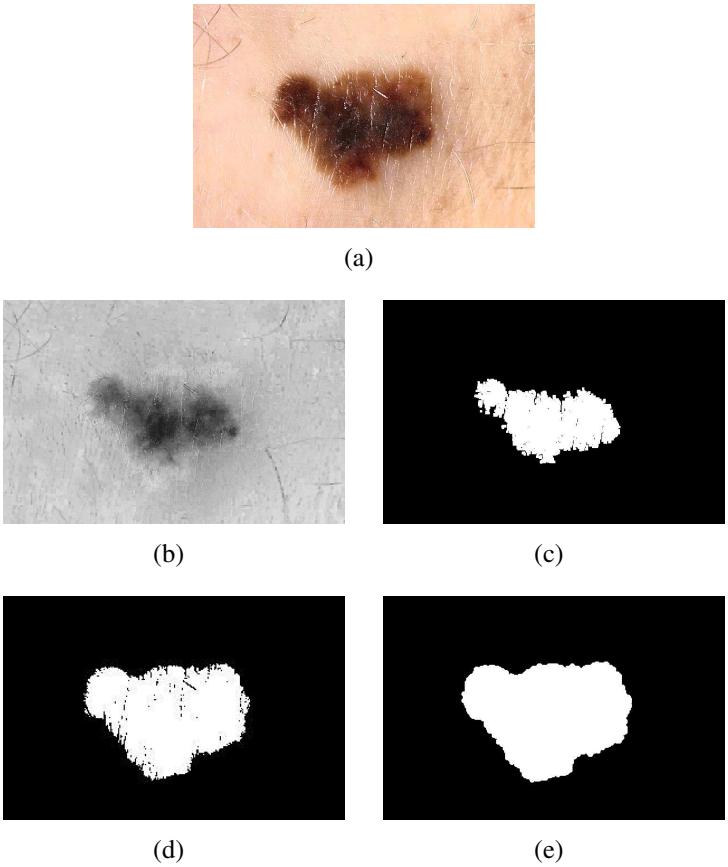

(a) (b) (c) (d) (e)

Figure A.5: Illustration of the segmentation using the ICA-Based Active-Contours method. (a) Color image after shading attenuation. (b) First re-ordered independent component/channel of image (a). (c) Lesion localization mask. (d) The active-contours segmentation result. (e) Final lesion segmentation, after post-processing image (d).

et al. and the Multi-Direction GVF Snake method were ranked last, since these methods presented higher segmentation errors. It shall be observed that these methods do not have a post-processing step. Analyzing the segmentation results visually (see Figure A.6), we may observe that the lack of a segmentation post-processing step to eliminate artifacts usually generate segmentations larger than the lesion itself.

Table A.1: Comparison of the six segmentation methods tested for the MPSL image database.

Approach	Image type	Post-processing type	Computational cost	ϵ in average
Otsu's Thresh. Method on Grayscale	Monochromatic.	M. Morphology	Low	42.33%
Otsu's Thresh. Method on the Red Channel	Monochromatic	M. Morphology	Low	38.58%
Alcón et al. Thresh. Method	Monochromatic	-	Low	165.31%
Multi-Direction GVF Snake Method	Monochromatic	-	High	59.60%
Thresh. Method on a Multichannel Repr.	Multichannel	M. Morphology	Medium	34.83%
ICA-Based Active-Contours Method	Multich. (Color)	M. Morphology	High	28.34%

Besides the average errors, we present in Table A.2 the percentage of images in the database that had lesion segmentation errors lower than 5%, 10%, 20%, 30% and 40%, respectively. As can be seen, the Otsu's thresholding method applied on grayscale images

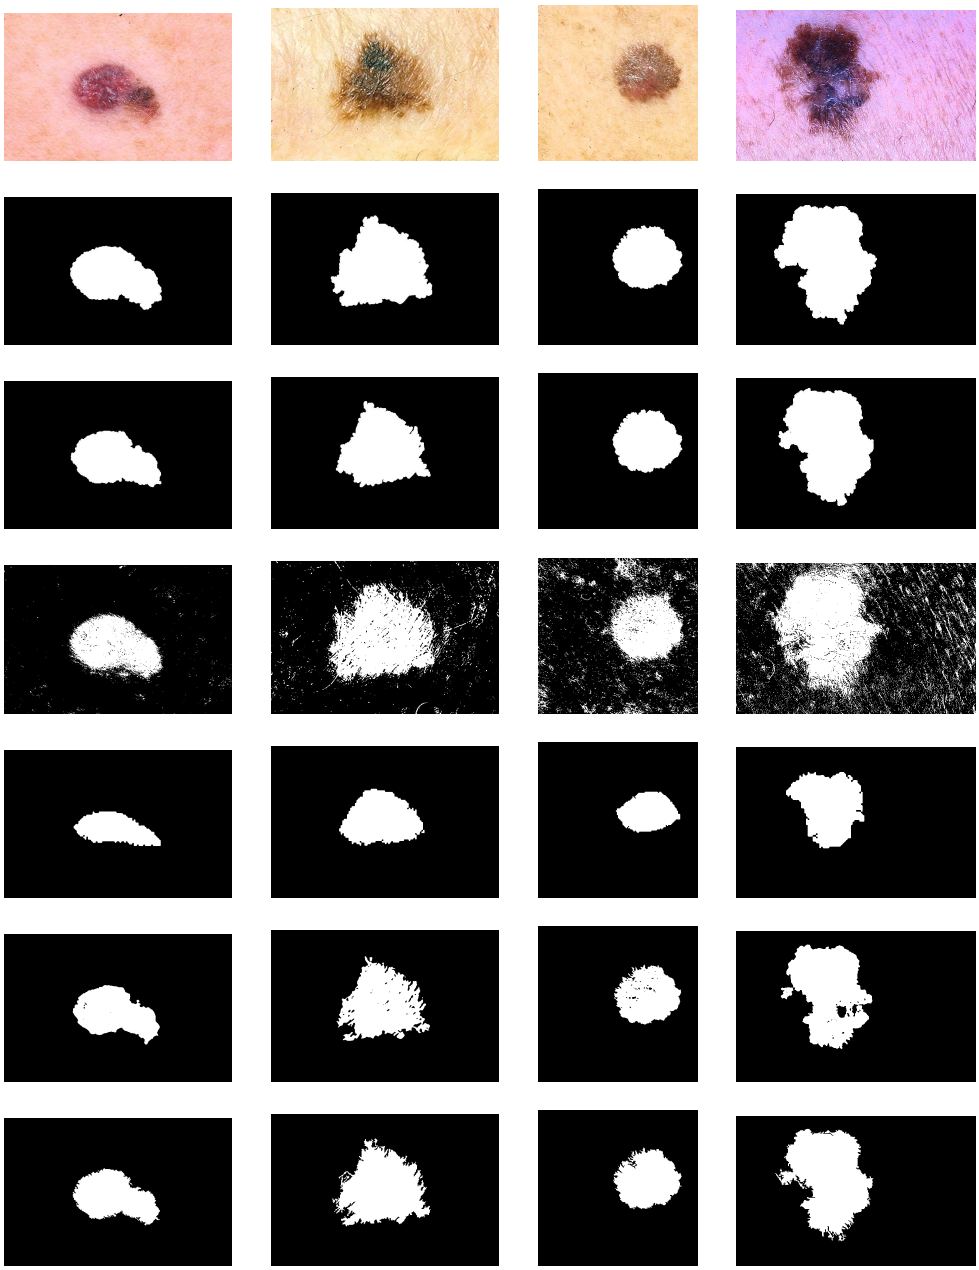

Figure A.6: Examples of MPSL images and their respective segmentation results for the six presented methods. In the 1st. line, the color images (after shading attenuation). From 2nd. to 7th. lines, respectively, results of Otsu's Thresholding Method on Grayscale, Otsu's Thresholding Method on the Red Channel, Alcón et al. Thresholding Method, Multi-Direction GVF Snake Method, Thresholding Method on a Multichannel Representation, and ICA-Based Active-contours Method.

tends to be more accurate, but potentially it also can generate larger segmentation errors than the ICA-Based Active-Contours method (the only method that uses color MPSL images). On the other hand, the ICA-Based Active-Contours method tends to be more reliable in the sense that it is less likely to produce large segmentation errors, and obtained experimentally the most accurate lesion segmentation results (in average).

Considering the results presented in Tables A.1 and A.2, we can infer that color can contribute to improve the segmentation of a MPSL image. Although methods based on

Table A.2: Segmentation errors in terms of error percentages.

Approach	$\epsilon < 5\%$	$\epsilon < 10\%$	$\epsilon < 20\%$	$\epsilon < 30\%$	$\epsilon < 40\%$
ICA-Based Active-Contours Method	3.29%	29.61%	73.68%	86.84%	94.08%
Otsu's Thresh. Method on Grayscale	3.29%	34.87%	75.66%	87.50%	93.42%
Otsu's Thresh. Method on the Red Channel	1.97%	27.63%	67.11%	83.55%	93.42%
Thresh. Method on a Multichannel Representation	1.32%	4.61%	43.42%	69.74%	81.58%
Alcón et al. Thresholding Method	0.00%	0.00%	8.55%	16.45%	25.00%
Multi-Direction GVF Snake Method	0.00%	0.00%	0.00%	5.92%	23.03%

one single channel may segment accurately some MPSL images, in average the color-based method provides more reliable segmentation results. In the next Sections, we discuss if the segmentation accuracy is correlated, or not, to the final classification/diagnosis accuracy.

A.5 Feature Extraction for Skin Lesion Discrimination

Given a MPSL image segmentation, we are able to obtain a classification (or diagnosis) of the acquired pigmented skin lesion. Before obtaining this classification, features representative of the skin lesion must be extracted. Computer-aided diagnosis systems often try to reproduce computationally the ABCD rule (NACHBAR et al., 1994), which is an acronym referring to the four criteria used in clinical diagnosis, namely: **A**symmetry, **B**order Irregularity, **C**olor Variation and **D**ifferential Structures. Several approaches have been proposed for describing quantitatively the first three criteria (FIKRLE; PIZINGER, 2007; ALCON et al., 2009; IYATOMI et al., 2008; CELEBI et al., 2007; GANSTER et al., 2001). These feature extraction techniques can be used jointly to represent benign and malignant cases better, and discriminate them more effectively. The fourth criterion (i.e. Differential structures) only is visible in dermoscopy images, but to describe quantitatively this lesion characteristic automatically still is challenging.

Since our ultimate goal is to evaluate the influence of image segmentation on the final classification/diagnosis, we will adopt the feature extraction and classification scheme proposed specifically for MPSL images by Cavalcanti and Scharcanski (CAVALCANTI; SCHARCANSKI, 2011). In the following Sections we present the 52 features that have been used in their ABCD rule implementation, also including features that were proposed by other authors. In terms of terminology, it is important to clarify that images $\bar{I}_i^C$ and $\bar{I}_i^N$ refer to a normalized color image and to the multichannel image representation presented in Section A.4.2.1, respectively.

A.5.1 Features Used for Lesion Asymmetry Characterization

The goal of these features is to quantify the lesion shape, in special the asymmetry of the lesion in relation to the principal axes. The major axis L_1 of the lesion is aligned with its longest diameter, passing through its center; the minor axis L_2 is orthogonal to L_1 and also passes through the shape center. The features utilized are the following :

f_1**:** Solidity: the ratio between the lesion area (A) and its convex hull area (ALCON et al., 2009);

f_2**:** Extent: the ratio between the lesion area and its bounding box area (ALCON et al., 2009);

f_3**:** Equivalent diameter: $4A/(L_1\pi)$ (ALCON et al., 2009; CELEBI et al., 2007);

f_4**:** Circularity: $4\pi A/(L_1 p)$, where p is the lesion perimeter (ALCON et al., 2009);

f_5**:** The ratio between the principal axes (L_2/L_1) (ALCON et al., 2009; CELEBI et al., 2007);

f_6**:** The ratio between sides of the lesion bounding box (ALCON et al., 2009);

f_7**:** The ratio between the lesion perimeter p and its area A (FIKRLE; PIZINGER, 2007);

f_8**:** $(B_1 - B_2)/A$, where, B_1 and B_2 are the areas in each side of axis L_1;

f_9**:** Similar to f_8, but makes use of the shorter axis L_2;

f_{10}**:** $B1/B2$ with respect to the axis L_1;

f_{11}**:** Similar to f_{10}, but makes use of the shorter axis L_2.

A.5.2 Features Used for Lesion Boundary Irregularity Characterization

The boundary sharpness is quantified by the magnitude of the gradient $\left|\vec{\nabla}\bar{I}_i^N\right|$ at each pixel using the Sobel operator (CAVALCANTI; SCHARCANSKI, 2011). However, instead of using pixels only at the lesion rim, we analyze pixels in an extended (dilated) rim[1] (ALCON et al., 2009). Consequently, lesions that have a smooth boundary (usually nevi) are better characterized. Also, the lesion boundary dilation makes the boundary representation more robust to the inaccuracies of the segmentation process. To characterize the lesion boundary irregularity, the following features are used (CAVALCANTI; SCHARCANSKI, 2011) :

f_{12}-f_{14}**:** Average gradient magnitude of the pixels in the lesion extended rim (ALCON et al., 2009), in each one of the three $\bar{I}_i^N$ channels;

f_{15}-f_{17}**:** Variance of the gradient magnitude of the pixels in the lesion extended rim (ALCON et al., 2009), in each one of the three $\bar{I}_i^N$ channels;

The lesion rim irregularity is characterized in the ABCD rule by dividing the rim in 8 symmetric regions (NACHBAR et al., 1994). In addition to the two principal axes L_1 and L_2, two additional axes are obtained by rotating by 45 degrees these orthogonal axes. Therefore, 8 symmetric regions $R = 1, ..., 8$ are generated. For each channel $\bar{I}_i^N$, the average gradient magnitudes of the extended rim pixels $\mu_{R,i}(R = 1, ..., 8)$ are computed. Therefore, 6 more features are calculated :

f_{18}-f_{20}**:** Average of the 8 $\mu_{R,i}$ values in each one of the three $\bar{I}_i^N$ channels;

f_{21}-f_{23}**:** Variance of the 8 $\mu_{R,i}$ values in each one of the three $\bar{I}_i^N$ channels;

[1]The rim is dilated by 2 pixels, producing a 5 pixels wide region centered at the lesion rim, as suggested in (ALCON et al., 2009).

Table A.3: Six Possible Colors of a Lesion on the RGB Color Space.

Color	Red	Green	Blue
White	1	1	1
Red	0.8	0.2	0.2
Light Brown	0.6	0.4	0
Dark Brown	0.2	0	0
Blue-Gray	0.2	0.6	0.6
Black	0	0	0

A.5.3 Features Used for Lesion Color Variation Characterization

These features quantify the color variation in the lesion, and the following measurements can be utilized for this purpose (CAVALCANTI; SCHARCANSKI, 2011) :

f_{24}-f_{27}**:** Maximum, minimum, mean and variance of the pixels intensities inside the lesion segment in the color variation channel $\bar{I}_3^N$;

f_{28}-f_{39}**:** Maximum, minimum, mean and variance of the pixels intensities inside the lesion segment in each one of three original $\bar{I}_i^C$ channels;

f_{40}-f_{42}**:** Ratios between mean values of the tree original $\bar{I}_i^C$ channels: $mean(\bar{I}_1^C)/mean(\bar{I}_2^C)$, $mean(\bar{I}_1^C)/mean(\bar{I}_3^C)$ and $mean(\bar{I}_2^C)/mean(\bar{I}_3^C)$, considering only pixels inside the lesion segment.

Physicians usually identify six distinct hues in skin lesions: white, red, light and dark brown, blue-gray, and black (NACHBAR et al., 1994). Lesions containing more of these hues are more likely to be malignant. The lesion color variability can be quantified by computing the occurrence of these typical hues within a lesion segment. Given a pixel in the lesion segment, the nearest reference color (associated with a typical hue, see Table A.3 (ALCON et al., 2009)) is found by the Euclidean distance to the pixel color in $\bar{I}_i^C$. A hue occurrence counter is created, one cell per typical hue. For each lesion pixel, the nearest typical hue counter is increased by 1. Finally, typical hues counters are normalized/divided by the lesion area A, and generate the 6 additional features f_{43}-f_{48}.

A.5.4 Features Used for Lesion Differential Structures Characterization

The lesion differential structures refer to submacroscopic morphologic and vascular structures only visible in dermoscope images. In an attempt to extract the characteristics of these structures also on a macroscopic image, differences between benign and malignant lesions can be measured using texture features. Cavalcanti and Scharcanski (CAVALCANTI; SCHARCANSKI, 2011) propose to extract the 4 features f_{49}-f_{52}, namely the maximum, minimum, mean and variance of the pixels intensities inside the lesion segment to represent the textural variation in the channel $\bar{I}_2^N$.

A.5.5 Feature Extraction Summary

Even for specialized physicians the discrimination of benign from malignant pigmented skin lesions may not be an easy task, and the development of techniques that facilitate this job is a current research topic. In the previous Sections, we presented 52 features that help in this task, but readers can probably find additional (or even alternative) features in literature, and probably new features will be proposed in the coming years as

this research area develops. The ultimate goal is to represent image lesion characteristics to could facilitate the early classification/diagnosis and reduce the number of deaths caused by these lesions.

Another important issue when dealing with the selection of feature sets is the adoption of automatic (or interactive) feature selection algorithms. Some authors of classification approaches for pigmented skins lesions (ALCON et al., 2009; CELEBI et al., 2007; RUIZ et al., 2008) suggest using a method to select the best features as those that help most the MPSL image class discrimination. However, this feature selection procedure can be tricky. Pigmented skin lesions have a large variability in terms of characteristics, and the discrimination between a malignant and a benign case can be determined by only 1 or 2 features. If these features have been eliminated because they did not seem to be important on a given training set, malignant cases can be incorrectly classified as benign, leading to costly false negatives.

A.6 Classification of Pigmented Skin Lesion Images

After extracting the features that characterize each MPSL image, we can use these data to obtain a final classification (or diagnosis) of the imaged lesion. We discuss how this information is processed In the following Sections.

A.6.1 Feature Normalization

The extracted features may have values in different ranges. Some of the proposed approaches do not perform feature normalization (ALCON et al., 2009; IYATOMI et al., 2008). However, classifiers tend to be more efficient if these feature values are normalized and represented in the same range (i.e., the feature values are scaled to fall in a specified range), therefore we adopt this normalization step in our work.

This feature value scaling is performed based on the mean and standard deviation of the captured feature values (CAVALCANTI; SCHARCANSKI, 2011; CELEBI et al., 2007; GANSTER et al., 2001). Among the possible feature normalization options (AKSOY; HARALICK, 2000; GANSTER et al., 2001), we chose to normalize the feature values using the well known z-score transformation (AKSOY; HARALICK, 2000):

$$Z_{i,j} = \frac{((v_{i,j} - \mu_j)/(3\sigma_j) + 1)}{2}, \tag{A.23}$$

where, $v_{i,j}$ is the value of the jth feature of the ith sample (image), μ_j and σ_j are the mean and standard deviation of the jth feature, respectively. After the z-score transformation, most of the $Z_{i,j}$ values are in the [0,1] range. The out-of-range values are saturated to either 0 or 1.

A.6.2 Defining Training Sets

Classification techniques are commonly used in machine learning. The validation test requires the definition of training sets, and two methodologies are usually applied, namely: cross-validation and holdout validation (BISHOP, 2007). The cross-validation method divides the samples in S portions, and $S - 1$ portions are used for training while the remaining portion is used for testing. This process is repeated until all the samples have been evaluated. In the holdout validation, part of the samples in each class (benign or malignant) are randomly selected and used for training, and the remaining samples are

used for testing. Often, up to half of the initial sample set is used for testing. The holdout validation method was used in our experiments described in Section A.7.

However, a common limitation of the public domain MPSL datasets is the relatively small number of cases, specially in the benign class (which often leads to unbalanced training sets). So, in addition to the selection of representative samples, we often need to balance and extend the training sets. One popular alternative is to add new training samples using the Smoothed Bootstrap Resampling method (CAVALCANTI; SCHARCANSKI, 2011; YOUNG, 1990). This method is used when there are not enough samples to guarantee the statistical significance of the data set. In this case, the original samples are randomly selected, and new ones are created by adding a small amount of zero-centered noise to their feature values, enlarging the data set (YOUNG, 1990). In our experiments (see Section A.7), we used zero mean Gaussian noise, with $\sigma = 0.1$, and made sure to obtain at least 2500 samples for each class (5000 samples in total) for a two class problem.

A.6.3 Classification Methods

After the feature normalization and the training samples selection steps, the generated data is the input used to train a classification method. For the classification of pigmented skin lesions images (macroscopic or dermoscopic), Support Vector Machines (SVM) is frequently utilized (TABATABAIE; ESTEKI; TOOSSI, 2009; CELEBI et al., 2007). However, due to the complex class shapes generated by the feature data, determining an adequate kernel and its parameters often is a difficult task. Artificial Neural Networks (RUIZ et al., 2008; DREISEITL et al., 2001) suffer from a similar limitation, determining the number of layers in multilayer perceptrons and their characteristics is critical, and may increase or decrease the final accuracy significantly. However, techniques such Decision and Regression Trees (ALCON et al., 2009) usually are less computationally intensive, and the parameters can be determined automatically. However, despite of their relatively simplicity, such classifier used alone hardly obtain the desirable accuracy levels by state-of-the-art of MPSL classification schemes.

In our experiments, where we try to relate segmentation and classification performance, we used a Nearest Neighbor Classifier (KNN) with K=1. This algorithm is very simple, where each sample/image is assigned to its neighbor class using the Euclidean Distance in feature space. This method was chosen for its simplicity and because it has been already used successfully in pigmented skin lesion image classification research (DREISEITL et al., 2001; BURRONI et al., 2004; CAVALCANTI; SCHARCANSKI, 2011).

A.7 Discussion of Experimental Evidences: Pigmented Skin Lesion Segmentation and Its Influence on the Lesion Classification and Diagnosis

In the previous Sections, we **discuss** the procedures utilized to obtain a classification/diagnosis for a MPSL image, from the pre-processing steps to the final classification. Now, we wish to evaluate the influence of the segmentation methods in the final lesion classification/diagnosis.

Recall that the Alcón et al. dataset used in our experiments (see Section A.4.3) contains 107 melanoma images (malignant cases) and 45 Clark nevi images (benign cases). In Tables A.4, A.5, A.6 and A.7 we present the classification results for these 152

images obtained by different state-of-the-art methods. In each one of these Tables, we present the Sensitivity (i.e, the percentage of MPSL images correctly classified in the malignant class), Specificity (i.e., the percentage of MPSL images correctly classified in the benign class) and Accuracy (i.e., the percentage of MPSL images correctly classified overall, considering all images).

Since the training sample selection process is random (see Section A.6.2), and may not assure that the selected samples represent well the characteristics of both classes, we considered as representative all the 50 training sets obtained in 50 random trails, as well as the 50 classification results corresponding to each one of these training sets. That is, we computed 50 times : (a) the training samples selection, (b) the Smoothed Bootstrap Resampling method, and (c) the Nearest Neighbor classifier results, so that our experimental results are statistically relevant. The final classification performance is measured based on these 50 trials.

We present the Accuracy, Specificity and Sensitivity averages of the 50 obtained results in Table A.4. In Tables A.5, A.6 and A.7, we present the best classification results for each segmentation method tested (in terms of Sensitivity, Specificity and Accuracy, respectively). As can be seen, these segmentation methods were ranked based on their segmentation (see Section A.4) and classification errors, but their segmentation and classification rankings differ. This indicates that the MPSL image feature set describes the lesions well, representing well their morphology even if the segmentation is not as precise as would be desirable. Additionally, these experimental results show that better feature extraction techniques that could take advantage of more accurate segmentations are needed, achieving classification results that are more accurate.

It shall be observed that our classification results indicate accuracies higher than 90% (in average), independent of the selected segmentation algorithm. These accuracies are higher than the diagnosis accuracies obtained by trained physicians in telemedicine, which range between 31% to 85% according to the literature (WHITED, 2006).

Table A.4: Comparison of classification results in average.

Approach	Sensitivity	Specificity	Accuracy
Alcón et al. Thresholding Method	92.56%	89.16%	91.55%
Otsu's Thresholding Method on Grayscale	93.89%	83.91%	90.93%
ICA-Based Active-Contours Method	92.95%	85.60%	90.78%
Thresholding Method on a Multichannel Representation	93.91%	83.29%	90.76%
Multi-Direction GVF Snake Method	92.34%	86.89%	90.72%
Otsu's Thresholding Method on the Red Channel	94.17%	80.36%	90.08%

Table A.5: Best classification results in terms of Sensitivity.

Approach	Sensitivity	Specificity	Accuracy
Otsu's Thresholding Method on the Red Channel	99.07%	75.56%	92.11%
Alcón et al. Thresholding Method	99.07%	75.56%	92.11%
Multi-Direction GVF Snake Method	98.13%	86.67%	94.74%
Thresholding Method on a Multichannel Representation	98.13%	86.67%	94.74%
ICA-Based Active-Contours Method	98.13%	84.44%	94.08%
Otsu's Thresholding Method on Grayscale	97.20%	80.00%	92.11%

Table A.6: Best classification results in terms of Specificity.

Approach	Sensitivity	Specificity	Accuracy
Alcón et al. Thresholding Method	92.52%	97.78%	94.08%
ICA-Based Active-Contours Method	92.52%	95.56%	93.42%
Multi-Direction GVF Snake Method	87.85%	93.33%	89.47%
Thresholding Method on a Multichannel Representation	93.46%	91.11%	92.76%
Otsu's Thresholding Method on Grayscale	94.39%	88.89%	92.76%
Otsu's Thresholding Method on the Red Channel	94.39%	88.89%	92.76%

Table A.7: Best classification results in terms of Accuracy.

Approach	Sensitivity	Specificity	Accuracy
Alcón et al. Thresholding Method	98.13%	88.89%	95.39%
Thresholding Method on a Multichannel Representation	98.13%	86.67%	94.74%
ICA-Based Active-Contours Method	97.20%	88.89%	94.74%
Multi-Direction GVF Snake Method	97.20%	88.89%	94.74%
Otsu's Thresholding Method on Grayscale	96.26%	88.89%	94.08%
Otsu's Thresholding Method on the Red Channel	94.39%	88.89%	92.76%

A.8 Summary and Future Trends

In this chapter, we reviewed the procedures used for classifying or diagnosing a pigmented skin lesion from a macroscopic image. Given the acquired color image, we showed how to eliminate shading effects, determine the lesion boundaries and some of the important lesion characteristics, and how to obtain the correct lesion classification as a (malignant) melanoma or a (benign) nevus.

The importance of using color also is outlined in this work, since the use of color can enhance the MPSL image segmentation precision. It shall be observed that the classification accuracies obtained in such segmentation and classification schemes already can be higher than the diagnosis accuracies obtained by trained physicians in telemedicine. Although the segmentation results do not correlate well with the final classification accuracies, we believe that in future new features can be developed to make better use of more precise segmentations, leading even higher classification accuracies.

In particular, we believe that such automatic MPSL image analysis schemes will contribute to increase the reliability of telemedicine. Consequently, the access to MPSL image prescreening systems shall be increase in the near future, which will contribute to improve the current early skin cancer detection rates, the skin cancer patient prognosis, and also it shall help increase the efficiency of the medical care systems.

References

See the unified bibliography of the thesis.