

Computer Aided Analysis of Epi-illumination and Transillumination Images of Skin Lesions for Diagnosis of Skin Cancers

Brian D'Alessandro, Atam P. Dhawan, *Fellow, IEEE*, and Nizar Mullani

Abstract—Skin lesion pigmentation area from surface, or, epi-illumination (ELM) images and blood volume area from transillumination (TLM) images are useful features to aid a dermatologist in the diagnosis of melanoma and other skin cancers in early curable stages. However, segmentation of these areas is difficult. In this work, we present an automatic segmentation tool for ELM and TLM images that also provides additional choices for user selection and interaction with adaptive learning. Our tool uses a combination of k-means clustering, wavelet analysis, and morphological operations to segment the lesion and blood volume, and then presents the user with six segmentation suggestions for both ELM and TLM images. The final selection of segmentation boundary may then be iteratively improved through scoring by multiple users. The ratio of TLM to ELM segmented areas is an indicator of dysplasia in skin lesions for detection of skin cancers, and this ratio is found to show a statistically significant trend in association with lesion dysplasia on a set of 81 pathologically validated lesions ($p = 0.0058$). We then present a support vector machine classifier using the results from the interactive segmentation method along with ratio, color, texture, and shape features to characterize skin lesions into three degrees of dysplasia with promising accuracy.

I. INTRODUCTION

IF detected early, the survival rate of skin cancer is quite high. The Nevoscope is a novel device developed for the non-invasive imaging of suspicious skin lesions through surface reflectance based epi-illumination (ELM) and diffuse reflectance based transillumination (TLM). In TLM imaging, surface illumination is blocked. Light photons enter the skin at an angle through a fiber-optic directed illumination ring source placed against the skin. Transillumination imaging using the Nevoscope is advantageous because light is directed beneath the skin, which allows the visualization of subsurface structures in the diffuse backscattered image. Figure 1 shows the difference between the pathways of light in TLM and ELM imaging.

TLM images of skin lesions are able to show the vascular blood flow information [1]. Cancerous lesions tend to require more blood to supply the cancerous cells with nutrients, and therefore, the capillary bed beneath the lesion will be larger and more saturated with blood than the

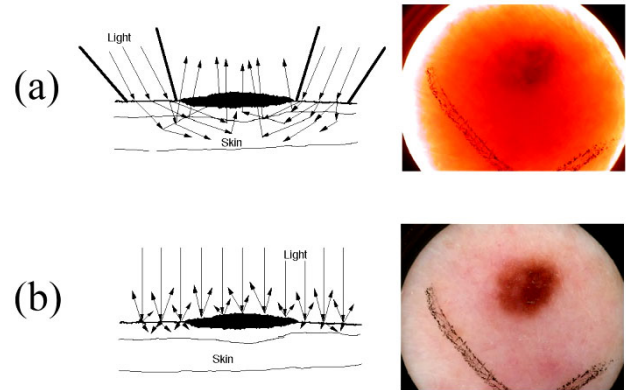


Figure 1: (a) Transillumination Imaging. (b) Epi-illumination Imaging

surrounding healthy tissue. Thus, the size of the subsurface blood volume can serve as an indicator of cancerous or pre-cancerous lesions in addition to the conventional method of observing the size and shape of the lesion on the surface, visible through ELM imaging.

If the area of the lesion blood volume and the area of the lesion surface pigmentation can be segmented out, then these areas may serve as useful features to help classify various types of lesions, such as melanoma. These features may also help to classify the various stages of cancer progression, from mild to moderate and severe. However, a completely automatic computer based segmentation which is reliably accurate for many different types of lesions and sets of images is difficult to construct. Instead, we present an interactive segmentation tool which shows the user six choices for both the ELM and TLM images. The user may then select which segmentation boundary best fits the lesion, and these boundaries are stored for later analysis and classification using a Support Vector Machine.

II. METHODS

To test our interactive algorithm, we acquired a set of ELM and TLM images for skin lesions marked out by a dermatologist. The objective of our interactive algorithm is thus two-fold: 1) to segment out the lesion boundary in the cross-polarized, surface illuminated (ELM) image; and 2) to segment out the subsurface blood volume boundary of the lesion in the transilluminated (TLM) image. Segmentation methods include k-means clustering and wavelet analysis. The use of *a priori* knowledge is also essential for reliable segmentation.

Manuscript received April 15, 2011.

Brian D'Alessandro (email: bmd5@njit.edu) and Atam P. Dhawan (email: dhawan@njit.edu) are with the Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ 07102 USA.

Nizar Mullani is with Translite LLC, Sugarland, TX.




Figure 2: Left, original ELM image. Right, image with black marks segmented in green.

A. Artifact Removal

Like many real world problems, the images are not perfect. Most notably, a large percentage of the images contain one or more “black marks,” that is, dark markings made by a pen or marker on the skin where the dermatologist indicated which areas should be imaged. Due to their dark nature, such marks may easily be mistaken for lesions by the segmentation algorithm. Hence, these marks must be detected and segmented themselves first, and then removed from any subsequent lesion boundary analysis.

The black mark areas are detected through wavelet analysis. Wavelet analysis is a powerful multi-resolution tool used to decompose an image into variable spatial and frequency resolutions [2]. Decomposition using the Daubechies D4 is performed in both spatial dimensions of a digital image. It was found that the marks were easily segmented in the ELM images by a thresholding and morphological processing of the high-high wavelet coefficients (high frequency in the x-dimension, and high frequency in the y-dimension) of the saturation channel in the Hue, Saturation, Intensity (HSV) representation of the ELM image (see Figure 2 for an example). Likewise, the black marks may be reliably segmented in the TLM image through similar thresholding and morphological processing of the high-high wavelet coefficients of the red channel of the TLM image [3].

B. Background and Color Correction

The original TLM images are largely saturated red in color, and this makes it difficult to distinguish subtle intensity features which may be of value to determine the blood volume. In addition, because of the physical nature of the Nevoscope as a ring source around the lesion, regions close to the light appear brighter in the image, while the center of the image (furthest from the light source) appears darker. This dark field effect is dependent on the optical properties of the patient as well as the contact and positioning of the Nevoscope on the individual’s skin. As a result, the background characteristics of the TLM image can understandably vary significantly between images and patients. Ideally, a background image of skin without the lesion should be taken on each patient, but this was not available in our image set. Thus, the background must be estimated mathematically for correction.




Figure 3: Left, original TLM image. Right, background and color corrected TLM image, more clearly showing the underlying vasculature.

We performed an initial rough segmentation of the mark-removed TLM image into two classes: a background class and a foreground class. Next, we take the pixels in the background class and fit a two dimensional, second-order polynomial curve to this data. Given this background estimate, we can then correct the TLM image with respect to this background, thus flattening out the intensity of the skin surrounding the lesion and heightening the contrast between the blood volume and the surrounding tissue. This procedure is performed on each color channel separately. If I_R and I_G are the original red and green channels of the TLM image respectively, and BG_R and BG_G are the estimated red and green channel background curves, then the background and color corrected images \hat{I}_R and \hat{I}_G are defined as:

$$\hat{I}_R = \left(\frac{I_R}{BG_R} - 0.16 \right)^{1.6} \quad \hat{I}_G = \left(\frac{I_G}{3 \cdot BG_G + 0.02} \right)^{0.452} \quad (1)$$

Division of the image by the background estimate provides the needed background correction while the constants present in these equations ensure that the corrected image values are kept within a valid and visible range. The exponents were selected heuristically based on two example images. This gamma correction aids in contrast and color enhancement, so that the final corrected TLM image can be more easily segmented by the k-means algorithm and is visually informative to a user when displayed on a typical RGB computer monitor.

C. Segmentation Options

Three of the six segmentation options for ELM and five of the six for TLM images are based on clusters generated by the k-means algorithm. K-means clustering on a set of data starts by choosing a set of cluster centroids randomly from the data. Each data element (in our case, a pixel value) is then assigned to the cluster associated with the closest centroid, based on a distance measure between each point and each centroid. The k-means algorithm then re-computes the centroid location of each cluster, and then updates the cluster assignments of the data based on the new centroid locations. The algorithm iteratively updates the centroid of each cluster and the membership of each cluster until a local minimum solution has been found. For the ELM image, the algorithm is run on pixels in the image’s saturation component of its HSV representation. For the TLM image, the algorithm is run on the G channel of the RGB image.

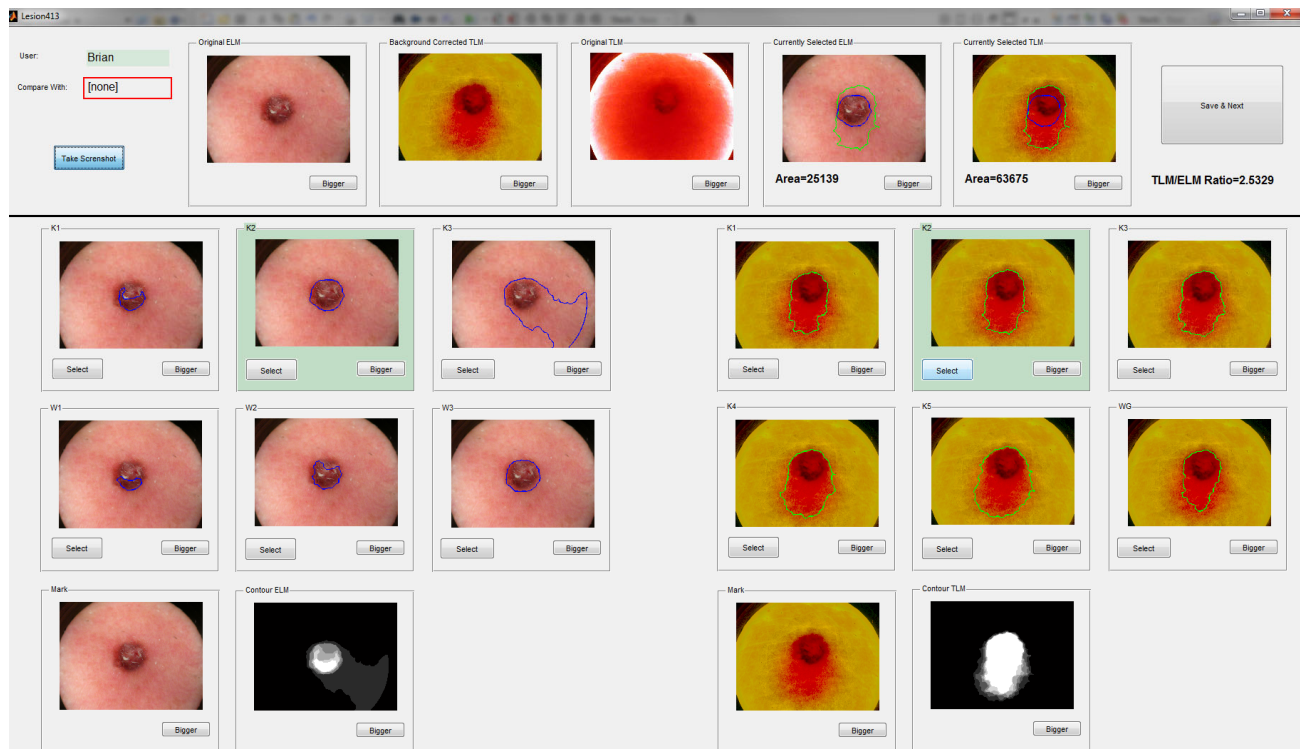


Figure 4: The Interactive Segmentation Interface

Wavelet analysis is also used to produce three additional segmentation options for ELM and one additional for TLM images. For the ELM image, another k-means algorithm is run on the high-high decomposition of the saturation channel to produce three levels of segmentation options. For TLM, the additional option includes a thresholded region where the green channel high-high wavelet component is near zero. Morphological operations are used to clean up the segmentations when necessary. Segmentation is also constrained through *a priori* knowledge that the lesion is likely centered in the image, and the TLM and ELM images must overlap.

D. The Interactive Segmentation Interface

Figure 4 shows the interactive segmentation selection interface. Based on the algorithms described above, six choices are given for both the ELM and TLM images for a lesion. In the top row, the user can view the original ELM and TLM images, and can also view the currently selected ELM and TLM boundaries superimposed on the same image for spatial comparison of the boundaries. The bottom row shows contour maps of both ELM and TLM segmentation options, providing a visual cue as to the similarity between the different selection options. Once the appropriate selections are made, the user can click “Next” to save the segmentation boundaries and then load the next lesion.

Another beneficial aspect of our interactive system is the support for multiple users. Different analysts may score the same lesions, so the selections may be compared and improved with feedback from multiple individuals. The

selections of a previous user may also be displayed on screen for comparison while a new user makes his or her selections. In this way, the additional input can refine the final selection of segmentation boundaries to approach the best solution. The thresholds used in the analysis algorithms may also be updated based on the selection feedback, thus using the advantages of adaptive learning to improve the segmentation.

Once the ELM and TLM images are segmented with respect to the pigmentation area and the blood volume area, the ratio between the two areas is computed. This ratio serves as a metric in an adaptive learning algorithm to differentiate the pathological classes of lesions.

E. Additional Color and Texture Features

Once the ELM and TLM areas were segmented using the interactive interface, an additional set of color, texture, and shape features were extracted from the pixels within the segmented areas to assist in the classification task.

A total of 80 features were generated from the second order histogram, or, gray-level co-occurrence matrix of the R and G channels of the original TLM image, and the S and V channels of the ELM image in HSV space. Such features include contrast, correlation, energy, and homogeneity, evaluated over five offset distances [4]. Principal Component Analysis (PCA) was then used to reduce the number of these features to 16 so as to reduce dimensionality and minimize overfitting of the data, while still keeping relevant feature information. An additional 160 features were calculated from the mean and variance of wavelet coefficients of the high-high and low-low wavelet packet decompositions for two

Table I: Mean (μ) and Standard Deviation (σ) of Lesion Classes

Class	N	μ	σ
Mild	21	1.13	0.20
Moderate/Severe	52	1.26	0.39
Malignant Melanoma	8	1.63	0.50

levels in the multichannel ELM and original TLM images. PCA was used again to reduce the number of wavelet features to only ten. Morphological features (eccentricity, solidity, extent, perimeter, and circularity) were also extracted from both the ELM and TLM region boundaries, providing ten additional features.

The GLCM, wavelet, and region boundary features were joined with the ELM area and the TLM/ELM ratio, for a total set of 38 extracted features. Still, the classification ability of good features may be clouded by noisy features with poor or no classification ability. Some subset of features is thus desirable for improved classification. Hence, we utilized a genetic algorithm to quickly find an optimal subset of the 38 features. Chromosomes in the GA were defined to be binary ‘flags’ which indicate if an individual feature should be included or ignored in classification.

III. RESULTS

A total of 81 lesions were processed using our novel segmentation interface to find the TLM/ELM ratio. Segmentation and selection took only a matter of seconds for each lesion. These lesions had been pathologically evaluated according to the severity of the dysplasia, and we grouped them into three classes: Mild, Moderate/Severe, and Malignant Melanoma. Lesion classes included both compound melanocytic nevi and junctional melanocytic nevi.

The mean and standard deviation of the ratios within each of the three classes were calculated (see Table I). There exists a clear, gradually increasing trend in the average ratio from mild to malignant lesions, thus confirming the relationship between lesion vascularity and dysplasia, and indicating that our methodology, image processing and segmentation algorithms can successfully detect and highlight this trend. We performed an ANOVA test on the three classes to statistically evaluate the differences in means observed. The p value indicates that the trend is significant ($p = 0.0058$). Our statistical significance and the separability of the classes shows a good improvement over previous work where the significance was lower and the means between classes were much closer together [1].

A. Support Vector Machine Classification

A learning based Support Vector Machine (SVM) classifier was constructed to classify lesions into the different grades of dysplasia based on the color corrected images and segmentations obtained using our interactive interface. A 60% random selection of lesions was used for training the SVM, while the other 40% was used for testing the classifier.

This procedure was performed 10 times to find the average result, each time with a different random training and test set. The fitness function for the GA search for an optimal feature subset utilized a simple linear kernel for the SVM to determine how well the features flagged by each chromosome classify the data.

After the GA had finished running, a total of 13 features were selected out of the original 38: the TLM/ELM ratio, six wavelet PCA components, four GLCM PCA components, TLM solidity, and TLM perimeter. Independent runs of the GA with different random seeds confirmed this same selection of features. It is clear that lesions with a large TLM/ELM ratio, high wavelet and GLCM activity, and large TLM perimeter are more likely to be classified as severe or malignant lesions. Smoother, smaller, and homogeneous lesions are more likely to be classified as mild.

Classification accuracy can be improved slightly through use of an RBF kernel $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$ in the SVM with $\gamma=10^{-5}$ and cost of error $C=10^{4.5}$. These parameters were used along with the 13 selected best features from the GA for a final SVM classification on the data. The classification was repeated 10 times, to find an average test accuracy of 81.0% and average training accuracy of 90.4%.

IV. CONCLUSIONS

Our proposed algorithm shows much promise in the ability to classify the different grades of skin lesion dysplasia. Our TLM background and color correction algorithm along with our lesion segmentation algorithm and interactive interface are clearly able to highlight the increase in vascularity present in increasingly dysplastic lesions. The trend we observe in the TLM/ELM ratio is highly significant. Classification of lesions using SVM shows good promise in grading the specific severity of lesion dysplasia, with an aim towards grouping lesions into classes where appropriate action can be taken by a dermatologist. These tools and methods would be useful to such a dermatologist as additional information to assist in the decision to biopsy. Our goal in the future is to improve our classification results through the intelligent selection of other distinguishing features using multispectral imaging to reliably classify our three levels of dysplastic skin lesions.

REFERENCES

- [1] V. Terushkin, S. W. Dusza, N. A. Mullani, M. Weinstock, R. Drugge, A. Dhawan, M. Duvic, and A. A. Marghoob, "Transillumination as a means to differentiate melanocytic lesions based upon their vascularity," *Archives In Dermatology*, vol. 145, pp. 1060-1062, 2009.
- [2] S. Mallat, *A Wavelet Tour of Signal Processing*: Academic Press, 1999.
- [3] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*: Pearson Prentice Hall, 2008.
- [4] R. M. Haralick, K. Shanmugam, and I. H. Dinstein, "Textural Features for Image Classification," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 3, pp. 610-621, November 1973.