Published in final edited form as:

Skin Res Technol. 2003 May ; 9(2): 147–156.

Colour histogram analysis for melanoma discrimination in clinical images

Yunus Faziloglu¹, R. Joe Stanley², Randy H. Moss¹, William Van Stoecker³, and Rob P. McLean⁴

¹University of Missouri-Rolla, Department of Electrical and Computer Engineering, 229 Emerson Electric Co. Hall, USA

²127 Emerson Electric Co. Hall, Rolla, MO 65409-0040, USA

³Stoecker & Associates, 1702 East 10th Street, Rolla, MO 65401-4600, USA

⁴3820 16th Avenue South Minneapolis, MN 55407, USA

Abstract

Background—Malignant melanoma, the most deadly form of skin cancer, has a good prognosis if treated in the curable early stages. Colour provides critical discriminating information for the diagnosis of malignant melanoma.

Methods—This research introduces a three-dimensional relative colour histogram analysis technique to identify colours characteristic of melanomas and then applies these ‘melanoma colours’ to differentiate benign skin lesions from melanomas. The relative colour of a skin lesion is determined based on subtracting a representative colour of the surrounding skin from each lesion pixel. A colour mapping for ‘melanoma colours’ is determined using a training set of images. A percent melanoma colour feature, defined as the percentage of the lesion pixels that are melanoma colours, is used for discriminating melanomas from benign lesions. The technique is evaluated using a clinical image data set of 129 malignant melanomas and 129 benign lesions consisting of 40 seborrheic keratoses and 89 nevocellular nevi.

Results—Using the percent melanoma colour feature for discrimination, experimental results yield correct melanoma and benign lesion discrimination rates of 84.3 and 83.0%, respectively.

Conclusions—The results presented in this work suggest that lesion colour in clinical images is strongly related to the presence of melanoma in that lesion. However, colour information should be combined with other information in order to further reduce the false negative and false positive rates.

Keywords

image processing; dermatology; colour; malignant melanoma; histogram

Physician's low accuracy in diagnosing pigmented lesions has been well documented (1–4). In three previous studies of diagnostic accuracy, sensitivity for malignant melanoma detection by dermatologists was 89, 77, and 81% (2, 3). Concern over the rising numbers of malignant melanomas has been paralleled by an increasing awareness of the specific signs of early malignant melanoma. Current melanoma detection techniques in most clinics in US

rely on guidelines such as asymmetry, border irregularity, colour variegation and diameter of the skin lesion, called the ABCD criteria (5). The clinical training for recognition of malignant melanoma that dermatology residents receive, aside from learning by examples of malignant melanoma in the clinic and in slide sets includes the same ABCD rule that the general public is taught. Conventional educational programs can improve diagnostic accuracy, but some discriminations such as melanoma versus dysplastic nevi remain difficult (6).

Numerous computer-based techniques have been applied in the past to pigmented lesion images for investigating features to detect malignant melanoma (7–25). In the research presented here, skin lesion colour analysis is examined. Clinical examination of colour provides discriminating information in the diagnosis of malignant melanoma (5). There are colours generally associated with melanocytic lesions, including shades of tan, brown or black and occasional patches of red, white or blue. Colours characteristic of melanoma are represented using ranges of colours in the red, green, and blue (RGB) colour space. Several colour descriptors have been applied to melanoma detection or discrimination, including variation of hues (8), analytical colour techniques for detecting colour variegation (17) and RGB colour channel statistical parameters (23–25). Ercal et al. (11) utilized relative chromaticity, spherical colour coordinates and (L , a^* , b^*) colour coordinate features as part of an overall neural network approach for melanoma detection. Ganster et al. (14) also used a neural network-based melanoma detection scheme, which utilized absolute, unnormalised colour-based percent melanoma colour features for melanoma colour discrimination in dermatoscopy images. The colour features included the percentages of absolute colour shades of reddish, bluish, greyish, and blackish areas found within the skin lesion as well as the number of these colours present within the skin lesion. Colour quantization was performed using the median cut colour quantization algorithm (26).

In this research, a data-driven relative colour histogram analysis technique is presented for determining colours characteristic of melanomas, i.e., ‘melanoma colours’. The use of relative colour (17), in which the background skin colour is subtracted from the lesion colour, has been proposed as a technique to avoid colour distortion in the imaging process. The skin lesion images used in this research were digitised with a Nikon Coolscan® from 35 mm slides collected from three different sources. Accordingly, there was no way to directly control colour variations from the different light sources, film types and film processing used for the acquired 35 mm slides. The relative colour technique attempts to address some of variations in the acquired photographs and also helps in correcting for different skin colouration (pigmentation) and for the Tyndall effect (preferential reflection of short (blue) wavelengths in the turbid medium) for different patients.

The new colour analysis technique presented here is based on the automated histogram mapping of colours representative of melanomas within clinical images from a training set. These melanoma colours are applied to differentiating melanoma and benign lesions solely based on the percentage of the melanoma colours found within the lesion. The outline remainder of this paper contains the colour histogram analysis technique in section II, experiments performed, experimental results and discussion in Section III and conclusions in Section IV.

Methods

The steps for skin lesion image colour analysis are:

1. obtain a training set of melanoma and benign skin lesion images,
2. determine the lesion boundary for each image,

3. determine a representative or average surrounding skin colour value for normalizing the skin lesion colour for each image,
4. perform colour histogram analysis for each image, generating a cumulative histogram over the training images,
5. identify colours characteristic of melanomas and benign lesions from the cumulative histogram,
6. compute the colour feature percentage of each skin lesion within the training images that has colours characteristic of melanoma for analysis, and
7. determine a threshold or cut-off for discriminating melanomas or benign skin lesions based on the percent melanoma colour feature found over the training images.

The following sections present the algorithm above in greater detail.

Skin lesion boundary determination

In this research, skin lesion borders within clinical images were found using a semi-automated procedure. Points along the border were manually chosen and joined with a least mean squares distance second order spline curve. A dermatologist working with a group of students determined all borders. Every attempt was made to determine accurate borders. There is no known gold standard for border accuracy.

Relative colour and colour histogram analysis

In this research, malignant melanoma discrimination is explored for skin lesions within RGB colour clinical images. Let I denote an M row $\times$ N column RGB skin lesion image with each pixel (x, y) for $1 \leq x \leq N$ and $1 \leq y \leq M$ represented as a column vector with red, green, and blue components, denoted as: $\vec{C} = \text{Col.}(r \ g \ b)$. The colour of a lesion pixel is defined using 'relative colour', given as

$$\vec{R} = \begin{bmatrix} r_{\text{rel}} \\ g_{\text{rel}} \\ b_{\text{rel}} \end{bmatrix} = \vec{C}_{\text{les}} - \vec{C}_{\text{skin}} = \begin{bmatrix} r_{\text{les}} - r_{\text{skin}} \\ g_{\text{les}} - g_{\text{skin}} \\ b_{\text{les}} - b_{\text{skin}} \end{bmatrix}.$$

The absolute colour and surrounding skin colour vectors are represented by $\vec{C}_{\text{les}}$ and $\vec{C}_{\text{skin}}$, respectively, where $\vec{C}_{\text{les}} = \text{Col.}(r_{\text{les}} \ g_{\text{les}} \ b_{\text{les}})$ and $\vec{C}_{\text{skin}} = \text{Col.}(r_{\text{skin}} \ g_{\text{skin}} \ b_{\text{skin}})$ represents the average colour of a surrounding area of normal skin. The procedure for finding $\vec{C}_{\text{skin}}$ is presented in Section 'Surrounding skin determination'. r_{rel} , r_{les} , and r_{skin} are the relative colour of the red component of a lesion pixel, the red component of a lesion pixel and the average value of red for a surrounding area of skin, respectively. Green and blue colour components are similarly defined. The relative colour components are bounded as $-255 \leq r_{\text{rel}} \leq 255$, $-255 \leq g_{\text{rel}} \leq 255$, and $-255 \leq b_{\text{rel}} \leq 255$, for systems using eight bits each for red, green, and blue, respectively.

The relative colour values are mapped into histograms bins using a simple procedure. The relative colour value for each pixel is translated by subtracting the minimum possible relative colour, -255 , resulting in translated relative colour values in the range 0–510. The primary difficulty in mapping the translated relative colours to histogram bins is that there are $511 \times 511 \times 511$ bins in the three-dimensional relative colour histogram, and there are a limited number of images available to populate those histogram bins. The relative RGB

histogram bins are requantized to represent discrete ranges of colours that are characteristic of melanoma. The requantization approach used in this research divides each relative colour value by 4 and then truncates or, equivalently, shifts right the relative RGB values by two positions. The requantized histogram is $128 \times 128 \times 128$, containing 64 relative colours per bin to represent ranges of colours within a single bin. Each requantized bin is referenced based on red, green and blue relative colours. A given lesion pixel with relative colour

$\vec{R} = \text{Col.}(r_{\text{rel}} \ g_{\text{rel}} \ b_{\text{rel}})$ maps into histogram bin $(r_{\text{bin}}, g_{\text{bin}}, b_{\text{bin}})$, where $r_{\text{bin}} = \text{trunc}\left(\frac{r_{\text{rel}} + 255}{4}\right)$ with the range of values $-0 \leq r_{\text{bin}} \leq 127$, and g_{bin} and b_{bin} are determined similarly.

Division by 4 was chosen to reduce the range of relative RGB values based on analysing bin populations of different quantization levels. Division by 4 (on each axis) combines 64 bins of the original relative colour histogram into a single bin. Even with this, the resulting histogram was sparsely populated when all available training set images were used. Combining fewer original histogram bins would result in higher colour resolution at the expense of an even more sparsely populated histogram. The number 4 (per axis, or 64 total bins to be combined) was chosen as a trade-off between colour resolution and scarcity of histogram population, given the number of images available for populating the histogram.

Surrounding skin determination

The surrounding skin colour is assumed to be a constant value and is approximated by the average colour of pixels within a specified region surrounding the skin lesion. The circular region neighbours the skin lesion with origin at the centroid of the skin lesion. The surrounding skin region size was found as an empirically optimised function of the skin lesion size. The circular area to encompass the surrounding skin region was computed as $A_C = A_L + A_S$, with centre at the lesion centroid. A_C denotes the number of pixels required to obtain an adequate sample of surrounding skin. A_L is the lesion area found by counting the pixels inside the lesion boundary. A_S is determined based on the lesion area A_L and is given as:

$$A_S = \begin{cases} 4A_L & \text{if } A_L < 4,700 \\ 3A_L & \text{if } 4,700 \leq A_L < 10,000 \\ 20,000 & \text{if } 10,000 \leq A_L \end{cases}.$$

where all areas are given as pixel counts. These figures resulted from circle sizes optimised over all lesions in the data set to yield a surrounding skin region $A_S > A_L$ for small lesions, and $A_S = A_L$ for large lesions. The circle radius for finding surrounding skin is approximated as $r_c = \sqrt{\frac{A_C}{\pi}}$, where r_c is the distance from the lesion centroid to the outermost surrounding skin pixel.

Using the surrounding skin region, non-skin coloured pixels such as from clothing, hair, teeth, direct camera flash reflections, shadows, etc., must be removed in order to achieve a reliable sample of the surrounding skin. From prior research for skin lesion border segmentation, it has been found that normal skin colour falls within a certain region of the RGB colour space (22). Figure 1 presents a flowchart for the non-skin finder algorithm used to eliminate non-skin coloured pixels from the average surrounding skin colour calculation. An iterative process of increasing A_C is performed for increasing the number of surrounding skin pixels if a minimum number of skin-coloured pixels is not found (22).

Melanoma colour mapping and relative colour bin mapping

A training set of images is used to populate the relative colour histogram bins. For populating the relative colour histogram, 70% of the melanoma images and 70% of the

benign images are randomly chosen from the image data set. For each of these training lesions, the histogram bins are populated. To avoid noise artefacts, a minimum percentage of the lesion pixels are required to fall within a given bin to consider that bin populated. Empirical analysis yielded the optimum threshold $K = 0.00125A_L$ for a lesion image with area A_L . For higher K , too few bins were populated. For lower K , the amount of noise increased. Thus, if a histogram bin for a lesion image contains more pixels than 0.125% of the total skin lesion pixels in that image, then the bin is considered as 'populated' for that image. This comparison is done for all melanoma images, and the 'populated' bins are found. For each histogram bin over the training set of images, the number of images with skin lesions that populated that histogram bin is determined. The image count for populating each histogram bin is computed for the melanoma training set of images. The process is repeated to generate image counts for populating each histogram bin over the benign training set of images using the same criteria.

The image counts for populating each corresponding relative colour histogram bin over the training sets of melanoma and benign images are combined for bin labelling as follows. Let $\vec{B}$ the set of all 128^3 relative colour histogram bins. Each bin $\vec{B}_k \in \vec{B}$ has a melanoma colour bin probability and a benign colour bin probability. The melanoma colour bin probability, denoted $P(\vec{B}_k|\text{mel})$, is defined as the number of melanoma images that populate $\vec{B}_k$, satisfying the bin membership threshold, divided by the total number of melanoma training images. Likewise, the benign colour bin probability, denoted $P(\vec{B}_k|\text{benign})$, is defined as the number of benign images that populate $\vec{B}_k$, satisfying the bin membership threshold, divided by the total number of benign training images.

An arbitrary bin $\vec{B}_i \in \vec{B}$ at $\vec{B}_i = \text{Col.}(r_{\text{bin}} \ g_{\text{bin}} \ b_{\text{bin}})$ is considered a melanoma colour bin if $P(\vec{B}_i|\text{mel}) > P(\vec{B}_i|\text{benign})$, where the relative colours $(r_{\text{rel}} \ g_{\text{rel}} \ b_{\text{rel}})$ contained within $\vec{B}_i$ are specified using the expressions previously presented and $0 \leq r_{\text{bin}} \leq 127$, $0 \leq g_{\text{bin}} \leq 127$ and $0 \leq b_{\text{bin}} \leq 127$. A similarly defined bin $\vec{B}_i$ is considered a benign colour bin if $P(\vec{B}_i|\text{mel}) < P(\vec{B}_i|\text{benign})$. For the case $P(\vec{B}_i|\text{mel}) = P(\vec{B}_i|\text{benign})$, the bin $\vec{B}_i$ is equally likely to be a benign or melanoma colour bin, referred to as an uncertain bin. By performing this comparison for all the bins in the relative colour space, a map of melanoma and benign colour bins is determined. This map, denoted $c(\vec{B}_k)$ for $\vec{B}_k \in \vec{B}$, is generated as

$$c(\vec{B}_k) = \begin{cases} \text{unpopulated} & \text{for } P(\vec{B}_k|\text{mel}) = P(\vec{B}_k|\text{benign}) = 0 \\ \text{melanoma} & \text{for } P(\vec{B}_k|\text{mel}) > P(\vec{B}_k|\text{benign}) \\ \text{uncertain} & \text{for } P(\vec{B}_k|\text{mel}) = P(\vec{B}_k|\text{benign}) \neq 0 \\ \text{benign} & \text{for } P(\vec{B}_k|\text{mel}) < P(\vec{B}_k|\text{benign}) \end{cases}.$$

Colour mapping extrapolation

The data set consists of 256 clinical images, with 128 benign lesions and 128 melanomas. Due to the size of the available training set, the resulting histograms are sparsely populated, as already mentioned. If a bin is completely unpopulated, then the probability of it being populated is 0 for both melanoma images and benign images. When a test image (not a

member of the set of images used to generate the original histogram) is used, then it is possible that some of the colours in this lesion will be mapped to unpopulated bins. To expand the number of bins labelled as melanoma or benign, a histogram bin labelling extrapolation technique is applied.

An iterative bin neighbour evaluation approach was used to re-label unpopulated and uncertain bins. The re-labelling process is applied first to unpopulated bins and then to uncertain bins, because the uncertain bins were felt to lie on the border between melanoma colours and benign colours. Each bin not on a boundary of the $128 \times 128 \times 128$ histogram has 26 neighbours from the three-dimensional relative colour histogram. Each bin and its neighbours are evaluated as follows:

1. Determine if the bin is unpopulated.
2. If the bin is unpopulated, count the number of neighbouring melanoma and benign bins.
3. Re-label the unpopulated bin using the following rules
 - a. If the number of melanoma bins outnumbers the benign bins by 3 or more or if there is at least one melanoma bin and no benign bins, re-label the bin as a melanoma bin.
 - b. If the number of benign bins outnumber the melanoma bins by 8 or more or if there is at least one benign bin and no melanoma bins, re-label the bin as a benign bin.
 - c. If *a* and *b* are not met, retain the bin as unpopulated.

Ten iterations of this process are performed to facilitate re-labelling convergence. There is a stricter criterion for re-labelling unpopulated bins as benign than as melanoma because of the need to avoid false negative lesion classifications.

After completing the re-labelling process for unpopulated bins, a similar iterative process is performed for re-labelling the uncertain bins. The updated histogram mapping created from re-labelling the unpopulated bins is processed to determine bin mappings from the uncertain bins over 10 iterations. The only difference in the re-labelling process for the unpopulated and uncertain bins is in step 3. In step 3 for the uncertain bins, the bin in question is labelled the same as the majority of the neighbours. If there are equal numbers of melanoma and benign bins surrounding an uncertain bin, the bin retains the uncertain label.

Percentage of melanoma coloured pixels and lesion classification

Using a training set of images with equal numbers of benign skin lesions and melanomas, the relative colour histogram bins are populated and labelled. Once the histogram bins are labelled, the percentage of melanoma colour within each training image lesion is computed.

Formally, the percentage of melanoma colour within a lesion is given as $C = \frac{100M}{A_L}$, where M is the total number of pixels within the lesion that are contained in histogram bins labelled as melanoma bins and A_L is the lesion area in pixels.

The percent of melanoma coloured pixels (C) is computed for all training images. Skin lesion classification is based on automatically determining a threshold (C_T) percentage of melanoma colour pixels. For each C_T from 0 to 100 in increments of 0.01, the true positive rate (correct benign lesion classification rate) and the true negative rate (correct melanoma classification rate) are determined from the training images. C_T is chosen such that the true positive and true negative rates are equal. C_T is applied to lesion classification in the test images.

Experiments performed, results and discussion

In order to evaluate the percent melanoma colour feature for clinical image lesions, lesion classification was examined in 256 clinical images, with 128 benign lesions and 128 melanomas. The benign lesions include 40 seborrheic keratoses and 88 nevocellular nevi. The melanomas included in the data set are invasive, moderately advanced lesions, which typically exhibit ‘melanoma colour’ to a greater degree than early lesions. No *in situ* melanomas were included in the data set. Figure 2 presents clinical image examples of (a) a melanoma and (b) a benign skin lesion (seborrheic keratosis). From this data set, 70% of the images were used in the training set (89 benign lesions and 89 melanomas), with the remaining 30% of the images comprising the test set (39 benign lesions and 39 melanomas). Eighteen randomly chosen training and test sets were used for evaluating the percent melanoma colour feature using this data set. For each training and test set, the training images were used for determining colour bin assignments and the percent melanoma colour threshold C_T based on the approach described in the previous section. For each training and testing set, the corresponding threshold C_T was applied. If the percent melanoma colour computed for a lesion was greater than C_T , the lesion was classified as a melanoma. Otherwise, the lesion was classified as benign.

The following experimental results are presented from the 18 randomly generated training and test sets using the algorithm previously described. Figure 3 shows a three-dimensional representation of the resulting relative colour bin labels using the colour mapping procedure presented in Section II for a randomly generated training set of images. The light grey boxes correspond to the melanoma-labelled bins, and the darker grey boxes represent the benign-labelled bins. When this is viewed as a three-dimensional plot on a system that allows the viewing angle to be changed via a trackball, it can be seen that the melanoma- and benign-labelled bins have non-overlapping regions in the relative colour histogram, providing skin lesion discrimination information. There are also overlapping regions in the relative colour histogram, where it is difficult to clearly distinguish melanoma colours from benign colours.

Figure 4 presents a plot of the training set melanoma and benign lesion classification rates for a representative randomly chosen training/testing set over the lesion percent melanoma colour threshold C range, 0–100. The horizontal axis is the percent melanoma coloured pixels threshold C . The vertical axis provides the percent correct lesion classification rates as a function of C . The melanoma and benign lesion classification rates are shown separately. Based on the approach presented earlier for determining the threshold C_T , $C_T = 39.7$ is the point at which the melanoma and benign lesion classification rates are equal (88.8%) for the training data. Figure 5 shows the plots of the corresponding test set melanoma and benign lesion recognition rates over the lesion percent melanoma colour range 0–100. For $C_T = 39.7\%$ melanoma coloured pixels, the correct corresponding test set melanoma and benign lesion correct recognition rates are 87.2 and 82.1%, respectively.

Table 1 presents the classification results over the 18 randomly chosen training/test sets. The first column shows the random trial number. Column 2 provides the percent melanoma colour pixel threshold C_T determined from the training data for the corresponding trial number. Column 3 displays the training set melanoma and benign lesion correct classification rates for the corresponding trial number. Columns 4 and 5 present the corresponding testing set melanoma and benign lesion correct classification rates, respectively, for the corresponding trial number based on the threshold C_T from column 2. Note that the training melanoma and benign classification rates in column 3 are identical based on the threshold C_T selection process. Also note that trial number 12 corresponds to the representative case shown in Figs 4 and 5.

The results from Figs 3 and 4 and Table 1 lead to several observations. First, the histogram bin re-labelling (presented in section E. Colour mapping extrapolation) yielded a small improvement in diagnostic accuracy. Without re-labelling, some melanomas had colours in uncertain bins surrounded by melanoma bins. The re-labelling process decreases the number of diagnostically uncertain colours. The same pattern was seen for benign colours, but less frequently, because the number of benign bins is smaller, as seen from Fig. 3. Second, the percent melanoma colour pixel feature performs well in differentiating melanomas from benign lesions in the clinical image data set examined. The average test set melanoma and benign lesion classification results from Table 1 are 84.6 and 83.0%, respectively. For comparative purposes, Ercal et al. (11) presented experimental results based on neural network detection between melanomas and benign lesions using lesions from the same data set. In that study, 14 features were used, including irregularity index (border irregularity measure), percent asymmetry (of the lesion), RGB colour variances over the lesion, relative chromaticity, spherical colour coordinates and L^* , a^* , b^* colour coordinates. Average experimental test results ranged from 74.2–86.0% for melanomas and 83.2–86.3% for benign lesions. In this study, a single feature, the percent melanoma colour feature, produces average melanoma and benign lesion success rates of 84.6 and 83.0%, respectively. These results are comparable to the neural network approach using 14 features, including several colour indices.

A dermatologist performed the truthing for the skin lesions used in this research. All melanomas and nevocellular nevi diagnoses were confirmed with biopsies. Some clinically typical seborrheic keratoses were not confirmed with biopsies, but were confirmed with follow-up examinations. The experimental results obtained from this single feature are similar to diagnostic rates of dermatologists (1–3). Third, this research introduces a novel data-driven technique for identifying colours that are characteristic of melanomas. Using relative colour for colour histogram analysis provides a robust approach for quantifying colours typically found in melanomas and for classifying lesions based on their colour distribution. Fourth, the percent melanoma colour pixels threshold C_T is inconsistent for the 18 random training/testing sets, with maximum, minimum and average C_T values of 54.4, 36.4, and 44.7, respectively. The experimental results show that on average close to half of the lesion must contain melanoma colour pixels for the lesion to be considered a melanoma. Fifth, the threshold C_T selection procedure is performed to facilitate automated image analysis and lesion classification. The consequences of a false negative misdiagnosis are much worse than the consequences of a false positive misdiagnosis. Accordingly, a dermatologist may accept a higher false positive classification rate to facilitate a higher melanoma discrimination rate. If a 20% false positive rate is allowable for the random training set from Fig. 4 (trial #12 in Table 1), 95.5% of melanomas are correctly classified for the training data. From the corresponding test set in Fig. 5, 94.9 and 71.8% of the melanoma and benign lesions are recognised, respectively. Choosing the threshold C_T based on a 20% false positive rate for each training/test set yielded average test melanoma and benign lesion recognition rates of 94.7 and 69.80%, respectively. Finally, the test results from Table 1 show that the image composition in the training set influence classification capability using colours characteristic of melanomas. This is reflected in the testing set results ranging from 71.8 to 92.3% for melanoma discrimination and from 74.4 to 92.3% for benign lesion discrimination.

Conclusion

In this research, a three-dimensional relative colour histogram analysis technique is introduced for identifying colours characteristic of melanomas and applying these ‘melanoma colours’ to differentiate benign skin lesions from melanomas. Experimental results using the percent melanoma colour pixel feature for 256 clinical images yielded on

average a correct melanoma discrimination rate of 84.6% with a corresponding 83.0% benign lesion discrimination rate. Although the results from this work were quite reasonable in diagnosing melanomas, the percent melanoma colour feature alone does not contain enough information to make a definitive melanoma diagnosis. Instead, the colour information should be combined with other information in order to further reduce the false negative and false positive rates. However, the results presented in this work suggest that lesion colour in clinical images is strongly related to the presence of melanoma in that lesion.

Acknowledgments

This research was supported by NIH-SBIR grant CA 60294–03.

References

1. Cassileth BR, Clark WH Jr, Lusk EJ, Frederick BE, Thompson CJ, Walsh P. How well do physicians recognize melanoma and other problem lesions? *J Am Acad Dermatol.* 1986; 14:555–560. [PubMed: 3958271]
2. Kopf A, Mintzis M, Bar R. Diagnostic accuracy in malignant melanoma. *Arch Dermatol.* 1975; 111:1291–1292. [PubMed: 1190800]
3. Grin C, Kopf AW, Welkovich B, Bar R, Levenstein M. Accuracy in the clinical diagnosis of melanoma. *Arch Dermatol.* 1990; 126:763–766. [PubMed: 2189362]
4. Lindelof B, Hedblad MA. Accuracy in the clinical diagnosis and pattern of malignant melanoma at a dermatological clinic. *J Dermatol.* 1994; 21:461–464. [PubMed: 8089364]
5. Friedman RJ, Rigel DS, Kopf AW. Early detection of malignant melanoma: the role of physician examination and self-examination of the skin. *Ca-A Cancer J for Clinicians.* 1985; 35:130–151.
6. Marks R, Jolley D, McCormack C, Dorevitch AP. Who removes pigmented skin lesion? *J Am Acad Dermatol.* 1997; 36:721–726. [PubMed: 9146533]
7. Dhawan AP. An expert system for the early detection of melanoma using knowledge-based image analysis. *Anal Quant Cytol Histol.* 1989; 10:405–416. [PubMed: 3064762]
8. Landau M, Matz H, Tur E, Dvir M, Brenner S. Computerized system to enhance the clinical diagnosis of pigmented cutaneous malignancies. *Int J Dermatol.* 1999; 38:443–446. [PubMed: 10397584]
9. Schindewolf T, Stolz W, Albert R, Abmayr R, Abmayr W, Harms H. Classification of melanocytic lesions with colour and texture analysis using digital image processing. *Anal Quant Cytol Histol.* 1993; 15:101–111. [PubMed: 8391265]
10. Andreassi L, Perotti R, Burroni M, Dell'Eva G, Biagioli M. Computerized image analysis of pigmented lesions. *Chronica Dermatol.* 1995; 1:11–24.
11. Ercal F, Chawla A, Stoecker WV, Lee HC, Moss RH. Neural network diagnosis of malignant melanoma from colour images. *IEEE Trans Biomed Eng.* 1994; 41:837–845. [PubMed: 7959811]
12. Stoecker WV, Li WW, Moss RH. Automatic detection of asymmetry in skin tumours. *Comput Med Imag Graph.* 1992; 16:191–197.
13. Golston JE, Stoecker WV, Moss RH, Dhillon IPS. Automatic detection of irregular borders in melanoma and other skin tumours. *Comput Med Imag Graph.* 1992; 16:199–203.
14. Ganster H, Pinz A, Rohrer R, Wilding E, Binder M, Kittler H. Automated melanoma recognition. *IEEE Trans Med Imag.* 2001; 20:233–238.
15. Xu L, Jackowski M, Goshtasby A, Roseman D, Bines S, Yu C, Dhawan A, Huntley A. Segmentation of skin cancer images. *Image Vis Computing.* 1999; 17:65–74.
16. Sober AJ, Burstein JM. Computerized digital image analysis: an aid for melanoma diagnosis—preliminary investigations and brief review. *J Dermatol.* 1994; 21:885–890. [PubMed: 7852652]
17. Umbaugh SE, Moss RH, Stoecker WV. Automatic colour segmentation of images with application to detection of variegated colouring in skin tumours. *IEEE Eng Med Biol.* 1989; 8:43–52.

18. Claridge E, Hall PN, Keefe M, Allen JP. Shape analysis for classification of malignant melanoma. *J Biomed Engineering*. 1992; 14(3):229–234.
19. Moss RH, Stoecker WV, Lin S-J, et al. Skin cancer recognition by computer vision. *Comput Med Imag Graphics*. 1989; 13:31–36.
20. Donohoe, GW.; Nemeth, S.; Soliz, P. *IEEE Symposium Computer-Based Medical Systems*. IEEE Comp Soc; Los Alamitos CA: 1998. ART-based image analysis for pigmented lesions of the skin.; p. 293-298.
21. Lee, T.; Ng, V.; McLean, D.; Coldman, A.; Gallagher, R.; Sale, J. *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers, and Signal Processing*. IEEE; Piscataway, NJ: 1995. A multi-stage segmentation method for images of skin lesions.; p. 602-605.
22. Hance GA, Umbaugh SE, Moss RH, Stoecker WV. Unsupervised colour image segmentation with application to skin tumour borders. *IEEE Eng Med Biol*. 1996; 15:104–111.
23. Green A, Martin N, Pfitzner J, O'Rourke M, Knight N. Computer image analysis in the diagnosis of melanoma. *J Am Acad Dermatol*. 1994; 31:958–964. [PubMed: 7962777]
24. Seidenari S, Burroni M, Dell'Eva G, Pepe P, Belletti B. Computerized evaluation of pigmented skin lesion images recorded by a video microscope: comparison between polarizing mode observation and oil/slide mode observation. *Skin Res Technol*. 1995; 1:187–191.
25. Aitken JF, Pfitzner J, Battistutta D, O'Rourke PK, Green AC, Martin NG. Reliability of computer image analysis of pigmented skin lesions of Australian adolescents. *Cancer*. 1996; 78:252–257. [PubMed: 8674000]
26. Heckbert, P. *Proceedings of the SIGGRAPH '82*. Vol. 16. Comput Graph; 1982. Colour image quantization for frame buffer display.; p. 297-307.

Fig. 1.
Algorithm for determining non-skin coloured pixels.

Fig. 2. Clinical image examples of benign and melanoma skin lesions, (a) Melanoma image. (b) Benign lesion image (seborrheic keratosis).

Fig. 3. Example of a three-dimensional representation of the resulting relative colour bin labels for a randomly generated training set of images. The light grey regions are melanoma-labelled bins. The dark grey regions are benign-labelled bins.

Fig. 4. Representative training set melanoma and benign lesion classification rates for a randomly chosen training/test set. The arrow points to the threshold C_T selected.

Fig. 5.
Corresponding test set recognition rates for melanoma and benign lesions.

TABLE 1

Experimental results for 18 randomly selected training/testing sets

Trial number	Threshold C_T	Training set (% Correct) Melanoma/Benign	Testing set (% Correct) Melanoma	Benign
1	46.4	87.6	89.7	79.5
2	42.8	91.0	92.3	82.1
3	47.0	93.3	76.9	76.9
4	42.7	87.6	89.7	84.6
5	54.4	87.6	79.5	89.7
6	36.4	89.9	92.3	82.1
7	42.7	91.0	79.5	79.5
8	44.5	89.9	74.4	74.4
9	53.0	86.5	71.8	92.3
10	49.1	88.8	84.6	79.5
11	46.5	84.3	87.2	89.7
12	39.7	88.8	87.2	82.1
13	37.5	89.9	89.7	84.6
14	45.3	88.8	89.7	84.6
15	41.8	86.5	87.2	87.2
16	47.2	87.6	87.2	89.7
17	44.0	91.0	79.5	82.1
18	43.4	89.9	84.6	74.4
Average	44.7	88.9	84.6	83.0

Column 1 gives the trial number; Column 2 shows the threshold C_T ; Column 3 provides the training melanoma/benign lesion discrimination rates; Columns 4 and 5 display the testing melanoma and benign lesion classification rates, respectively.