

A Colour Palette for Automatic Detection of Blue-White Veil

Ali Madooei, Mark S. Drew

School of Computing Science, Simon Fraser University

{amadooei, mark}@cs.sfu.ca

<http://www.cs.sfu.ca/~amadooei>

Abstract

Colour assessment of pigmented skin lesions are essential for the diagnosis of malignant melanoma. However, visual interpretation of colour is subjective and prone to error. Computer programs can provide support to clinicians to overcome this subjectivity. So far, methods for colour analysis of this nature have utilised statistical classification models. This paper puts forward an alternative framework: an effort to reproduce the experience of human observer. The proposed method introduces a perceptually intuitive and semantically meaningful approach for colour and colour-related feature detection. As a case study, the task of automatic detection and segmentation of blue-white veil feature in dermoscopy images is examined. Our proposed method, as shown in our experiments, outperforms the prior art for this task, while it attempts to mimic the human perception of skin lesion colours.

1. Introduction

Colour assessment is important in the clinical diagnosis of many conditions, especially in diagnosis of skin diseases. For instance, presence of multiple colours with an irregular distribution in a mole strongly suggests malignancy. Therefore, the use of colour has been substantial in skin lesion classification.

In the computer analysis of medical images, in particular skin images, colour plays an important role too. For instance, colour can be used for lesion segmentation [1] or as one of the image features [2] for lesion classification.

Sometimes a particular colour is associated with a particular medical condition. Hence, automatic identification (detection) of the colour itself is of interest. As an example, blue-greyish or blue-whitish areas in dermoscopy (skin-surface microscopy) images are known to be significant indicators of invasive melanoma. Celebi et al. [3] developed a colour based method for automatic detection and segmentation of blue-white areas.

In most computerised colour analysis programs (as well as in [3]), colour models are based on statistical parameters and methods. These methods are used for computational convenience; they do not model human perception and interpolation of colour.

Recently, an interesting study by Seidenari et al. [4] made an effort to develop a computer program for colour assessment in dermoscopic images by mimicking the human perception of lesion colours. Their method identifies a colour by matching its RGB values (in nearest neighbour fashion) to a small set of “most representative colours”, interactively selected by user from a set of skin lesion images. Although the idea of [4] is intriguing, the study suffers from a number of flaws. The most serious of these are lack of connection to colour science and failing to utilise state-of-the-art techniques of computer vision. For example, using Euclidean distance metric in RGB space to compute colour differences and conduct colour matching is, in essence, contradictory




Figure 1. Melanoma images with Blue-white Veil

to the goal of mimicking human colour perception because RGB is not a perceptually uniform colour space.

This paper puts forward a framework for colour detection and assessment of a somewhat similar nature. Our method is inspired by the work of [4] yet different in that i) we make careful consideration in utilising colour models that are more congruent with human colour perception; ii) we systematically produce a discrete set of (Munsell) colours best describing a colour family to be detected (i.e. here, the colour of “blue-white veil” under dermoscopy, a strong indicator of melanoma); iii) we incorporate current colour analysis and computer vision techniques.

For the purpose of this publication, the proposed framework is implemented to automatically detect and segment blue-white veil areas in dermoscopy images. This is described next.

Blue-white veil detection for diagnosis of Melanoma

Melanoma, the most unforgiving skin cancer, is among the cancers with rising incidence and mortality rate [5]. Early detection of melanoma is paramount to patients’ prognosis towards greater survival. The clinical diagnosis of early melanoma is acknowledged as challenging [6] and has provoked increased interest in computer-aided diagnosis systems through automatic analysis of dermoscopy images.

Dermoscopy is a non-invasive imaging technique, popularly used for screening of pigmented skin lesions. It uses optical magnification and cross-polarized lighting (or fluid immersion, to remove scatter of light) which allows better visualization of skin morphological characteristics.

The clinical assessment of skin lesions under dermoscopy consists of assessment of shape, size, colour, border and elevation [7, 8]. Common colours under dermoscopy are light brown, dark brown, black, blue, blue-grey, red, yellow, and white [9]. There is more to these colours than meets the eye; for instance, the colour blue (under dermoscopy) indicates melanin localized within deeper parts of the skin [9].

Among dermoscopic features, blue-white veil is the single most important finding in making a diagnosis of invasive melanoma (with specificity of 97%) [10]. Fig. 1 shows dermoscopy images of melanoma with a blue-white veil feature; the

blue-white veil regions are structureless areas of confluent blue pigment with a ground-glass haze (as if the image were out of focus there) [11]. Blue-white veil is associated with “superficial fibrosis with melanophages and/or malignant cells in the papillary dermis” [12].

To our knowledge, the only study that reports a method, experimental procedure, and results specifically pertaining to detection of blue-white veil is the one by Celebi et al. Their approach involves pixel classification using explicit thresholding, where the threshold values are induced by a trained decision tree.

This paper puts forward an alternative method, by incorporating colour analysis and computer vision techniques to address the task under study. The proposed method is an attempt to mimic the human perception of lesion colours and outperforms the state-of-the-art [3] as shown in our experiments. The proposed method is described next.

2. Method

Let us look at the problem of blue-white veil detection from a further level of abstraction. There are two questions of interest to us: how a dermatologist identifies the presence of certain colours under dermoscopy? And, how can we develop computer programs to mimic this dazzling human ability?

Although it seems natural for human to associate names (labels) with colours, the task of colour naming conceals complex and unsolved problems in the field of computational colour science. If we ignore the great deal of uncertainty in naming colours, as well as the cultural and language dependencies, we can safely say that, when it comes to colour naming, people are influenced by the colours they saw previously.

Indeed, a dermatologist needs training to be able to identify the blue-white veil feature. Thus, to mimic the colour assessment performed by human observer, we propose to identify the colour feature of interest, i.e. blue-white veil, by colour matching to a *discrete* set of colours best describing that colour label.

To this aim, we first describe a method to identify and extract the set of colours best describing blue-white veil data, followed by a method to perform colour matching to automatically detect and segment this feature.

Discretization of blue-white veil colour data

We used a set of 105 images selected from [11], consisting of 43 images containing sizeable blue-white veil areas with the remaining 62 free of this feature. This dataset has been labeled and used by Celebi et al. For each image a number of small circular regions that contain either veil or non-veil pixels were manually determined and extracted as veil and non-veil data.

We analysed the veil data by mapping their colour values to the Munsell colour system [13]. In colorimetry, the Munsell colour system is one of the most fundamental and influential colour-modellings. Albert Munsell, its creator, was an American scientist and artist; although his work originally was devised more by intuition than exact science, yet it offers both perceptual and quantitative colour definitions. The perceptual definition (given in the form of a book with printed colour patches) is appropriate for the use of artists such as painters and designers, whereas the quantitative definition provides measurement standards that are appropriate for technical and scientific use.

Quantitatively, the Munsell system specifies a local colour

by giving its hue (H), value (V), and chroma (C), in the form HV/C. The value is a number between 0 and 10. The chroma is a positive number, whose bound depends on hue and value, as given by the MacAdam limits. The hue specification consists of a letter designator (B, BG, G, GY, Y, YR, R, RP, P, PB), and a number designator which is greater than 0 and less than or equal to 10 [14].

Since there is no direct conversion from standard RGB to Munsell colour quantities, we consider computing an approximate transform from CIELAB specification to Munsell specification.

The CIELAB (a.k.a. $L^*a^*b^*$) model is another well known colour system. The three coordinates are: L^* which represents the lightness of colour (luminosity layer), a^* which indicates colour differences along the red-green axis, and b^* that indicates colour differences along the blue-yellow axis. The space spanned by a^* and b^* represents the colour chromaticity. The sRGB values are directly mapped to CIELAB coordinate system through a non-linear transformation¹.

An alternative notation to $L^*a^*b^*$ is the $L^*C_{ab}^*h_{ab}$ expression where $C_{ab}^* = \sqrt{a^{*2} + b^{*2}}$ and $h_{ab} = \tan^{-1}(b^*/a^*)$ [15]. This alternative notation offers an advantage in that it is easier to relate it to the earlier systems based on physical samples, like the Munsell colour notation.

For colour conversion from CIELAB to Munsell, we follow [16]. There, L^* directly corresponds to Munsell value. The hue angle h_{ab} corresponds to Munsell hue; ideally, the ten hue designators should be evenly spaced in terms of their hue angle, with no dependency on chroma or value. Experiments [15] show that yellow occurs at $h_{ab} = 90^\circ$, so 5Y Munsell hue is set to correspond to 90 degrees. The other nine hues are assumed to be evenly spaced around the circle, in a counter clockwise direction. Finally, C_{ab}^* is taken to be 5 times Munsell chroma. “This is a crude approximation but the best available so far” [17].

The above-mentioned conversion is approximate: if one takes a Munsell sample, calculates its coordinates in CIELAB (when illuminated by Illuminant C), and then converts it back to the Munsell system by above-stated conversion routine, the resulting Munsell specification would not agree exactly with the Munsell specification of the original sample. “Nevertheless, the agreement is good enough for most practical purposes” [16].

Note that we are interested to achieve the best generalization in building the discrete set of blue-white veil data. For that task, as a pre-processing step, each image is converted to its ‘superpixel’ representation. Superpixels capture image redundancy by grouping pixels into perceptually atomic regions. They also preserve information over scales and sampling resolutions. There are a handful of methods to create superpixel representation. In our implementation, we used SLIC [18] which is a simple and efficient method based on a spatially localized version of k-means clustering².

Algorithm 1 summarizes the approach taken here to convert

¹A precise transformation requires knowledge of observer and light source; since the images in our dataset are taken under uncontrolled condition, we assumed standard observer and illumination (D65) for mapping from sRGB to CIELAB.

²“SLIC (Simple Linear Iterative Clustering) performs local clustering of pixels in the 5-D space defined by the L, a, b values of the CIELAB color space and the x, y pixel coordinates to efficiently generate compact, nearly uniform superpixels” [18].




Figure 2. Colour Palette for blue-white veil detection.

blue-white veil data to a discrete set of colour Munsell patches. Interestingly, the 146,353 pixels under analysis mapped to only 179 Munsell colour patches. Among these, 97% of the veil data is described by only 116 colour patches. Figure 2 shows these 116 colours organized on a palette, and grouped according to their Munsell hue.

Algorithm 1 – Discretization of blue-white veil colour data

- 1: **for** each image in database **do**
 - 2: Convert from sRGB to CIELAB
 - 3: Replace each pixel value with superpixel representation
 - 4: **for** each pixel marked as veil **do**
 - 5: Compute the approximate Munsell specification
 - 6: **end for**
 - 7: **end for**
 - 8: Create frequency table of computed Munsell colour patches
 - 9: Keep the most representative colours (in terms of highest frequency) and organize them in a palette
-

We also analysed non-veil data by the same principle. The 254,739 pixels from non-veil areas mapped to 220 Munsell colour patches, among which only 6 patches were overlapping with the 179 veil patches. These 6 contribute (altogether) to less than 1% of veil data and are not considered among the 116 patches in the blue-white veil colour palette.

The colour palette (Figure 2) serves as the human observer’s (i.e. dermatologist’s) prior knowledge for colour assessment. It can be used to extract blue-white veil regions from an input image through colour matching. This is described next.

Blue-white veil detection via colour matching

We identify blue-white veil feature in each image through a thresholded nearest neighbour matching. That is, each pixel of a test image is matched to one colour patch of the colour palette in a nearest neighbour fashion. If the (colour of) pixel is *similar (close) enough* to the one of the colour patches, it is identified as blue-white veil. Here, ‘close enough’ is evaluated by comparing the distance (between any given colour vector and its nearest match) against a threshold value³.

To reduce the computational cost of above-mentioned pixel-based colour matching, we can segment colours in any given image and instead match the colour vector of the centroid of each cluster (segment) to colours in the colour palette. By segmentation, we decompose an image into visually homogeneous regions and effectively replace the pixel-based approach by a region-based one, while preserving salient features of the overall image. Note that in dermoscopy images there are a handful of distinct colour regions (such as blue-white veil region) where the presence or absence of a feature significantly affects the diagnosis while the information within such a region is often less important and can be neglected.

Colour image segmentation can be done in an automated and unsupervised fashion (where the number and the shape of the image clusters are unknown). In our implementation, we use EDISON [19] software, a mean-shift based segmentation tool⁴.

Colour matching in the fashion described above is sensitive to colour representation employed and colour distance metric used. We consider working in the CIELAB colour space. This choice is motivated by two reasons: first, this colour space is an approximately perceptually uniform⁵ colour model. Thus the difference between two colours can be measured using e.g. the Euclidean distance metric. Second, CIE colour spaces are device independent, which make them suitable for colour matching and colour comparison.

It is to be noted that the difference (or distance) between two colours is a metric of interest in colour science. In this study we used the CIEDE2000 colour difference formula as indicated in [20]⁶. “The CIEDE2000 formula provides an improved procedure for the computation of colour differences from experimental data” [20].

Algorithm 2 summarizes the proposed method for detection and segmentation of blue-white veil areas in dermoscopy images.

3. Experiment

We tested our proposed method on a set of 223 images selected from [11] and used by Celebi et al. The image set consists of 173 images containing blue-white veil areas and a remaining 50 free of this feature.

³The threshold can be set experimentally using the training data. From the training data we can also define the posterior probability of a pixel being veil given its colour in terms of the likelihood of observing its associated colour patch given the class label and the prior probability of classes.

⁴Mean shift (MS) is a statistically robust mode-seeking algorithm that is based on clustering in both space and colour.

⁵That is, colour differences (in CIELAB spaces) agree more consistently with human visual perception. Note that the common sRGB colour space does not yield this property.

⁶A Matlab implementation is available on the author’s [20] website at <http://www.ece.rochester.edu/~gsharma/ciede2000/dataNprograms/deltaE2000.m>

Algorithm 2 – Blue-white veil detection

```
1: Load a skin lesion image
2: Convert from sRGB to CIELAB
3: Segment using EDISON
4: for each segmented region do
5:   Find the best match from colour palette
6:   if The best match is within the threshold distance then
7:     Classify region as veil
8:   else
9:     Classify region as non-veil
10:  end if
11: end for
```

The results of [3] is taken as the baseline. A summary of this baseline is given in algorithm 3.

Experimental results are presented in Table 1. We report sensitivity and specificity. Sensitivity measures how good a test is at detecting positives. It is defined as true positives/(true positives + false negatives). Specificity shows how good a test is at avoiding false alarms. It is defined as true negatives/(true negatives + false positives). These measures of performance are of utmost importance for medical image analysis systems. Any computer program to aid diagnosis is to be sensitive (as it would save more lives) and specific (since it would reduce financial and emotional burden on patients).

For any test, there is usually a trade-off between these performance measures. For example, the baseline achieves high specificity at the cost of low sensitivity. We believe this can be explained by the limitation of decision trees used in [3]. Decision trees are simple to use and easy to understand, yet (as can be seen) their good detection rate usually arrives at the expense of high false negatives. Our proposed method on the other hand improves on sensitivity while preserves specificity. To ensure that this improvement did not arrive at the cost of introducing high false positives, we also include and compare the F-score of methods. The latter is the harmonic mean of precision (true positives/(true positives + false positives)) and recall (sensitivity).

Method	Sensitivity	Specificity	F-score
Celebi [3]	0.65	0.97	0.68
Proposed	0.71	0.97	0.70

Table 1. Proposed method vs. Celebi et al.

Accurate detection and segmentation of blue-white veil feature can be useful for computer analysis of skin lesion images. However, in clinical assessments, presence or absence of this feature is associated with diagnosis. Accordingly, in a different experiment we aimed to determine only the presence (or absence) of veil feature in a set of 300 images taken from various sources. The image set is divided into two subsets of 200 ‘easy’ and 100 ‘challenging’ images. An image is considered challenging if the blue-white veil area was too small, too pale, occluded, or had variegated colour.

Our experiment produced accuracy of 87% and 67% on easy and challenging sets respectively ⁷. Figure 3 illustrates the output of proposed method on some of the images in each set. Note that

⁷In order to exclude very small areas which are without clinical relevance, a minimum value (threshold) for areas of veil region can be considered. We set this area threshold to 0.5% of image area.

for this set of images, since the lesion border was not available to us, we did not run the baseline method and only report our detection results as presented.

Algorithm 3 –Blue-white veil detection by Celebi

```
Load a dermoscopy image of skin lesion.
Extract lesion border.
Dilate the border by 10% of its area.
Extract region outside the dilated border of size 20% of lesion area.
for each pixel in extracted region do
  if  $R > 90$  and  $R > B$  and  $R > G$  then
    Mark the pixel as healthy skin.
  else
    Ignore the pixel and continue.
  end if
end for
Set  $\bar{R}_s$  as the mean of red channel values for pixels marked healthy skin.
for each pixel in the image do
   $nB = B/R + G + B$ 
   $rR = R/\bar{R}_s$ 
  if  $nB \geq 0.3$  and  $-194 \leq rR < -51$  then
    Classify pixel as veil
  else
    Classify pixel as non-veil
  end if
end for { Note that above-mentioned requires extraction of lesion border. Thus, the complexity and computation cost of [3] is influenced by the border extraction algorithm employed. In their experiments [3], the lesion borders were obtained manually “to separate the problem of feature extraction from the problem of automated border detection”. Also, the above requires a search for normal (healthy) skin colour in the background, outside the lesion border. Thus its performance is also constrained subject to the accuracy of skin colour filter used. A revised version of this method is given in [21] which requires less computation and performs equally well. }
```

4. Conclusion

The importance of colour assessment in medical images has been discussed. In the quest to computerise this process, most studies focus on statistical methods and parameters for colour description. These descriptors are mainly chosen for computational convenience; they do not model human perception and interpretation of colour.

Our proposed method on the other hand mimics the human process of identifying lesion colours while providing statistical outputs that can be ported to any computer-aided diagnosis system. In addition to its strong performance, the proposed method is intuitive and easy to understand, which makes it suitable for non-engineers (*viz.* clinicians) to employ and apply.

Furthermore, the proposed method can easily be extended to account for detecting other colour features in dermoscopy images. In fact, it can be seen as a scaffolding for colour based detection and assessment problems of similar nature.

The blue-white veil colour palette can be used for training

purposes, as well as a reference tool (such as designers' colour chart) for dermoscopy trainees and neophytes. For that matter, a colour palette of other common colours can be generated as well. Note that the proposed method can automatically generate colour palettes given labelled training data. The advantage of this is twofold since statistical data can be extracted from training data and associated with colour patches on colour palettes. In fact, one can use this method to extract a colour map of lesions with common colours (under dermoscopy) and link it to statistical data to e.g. associate a feature with a diagnosis.

Finally, this method can aid clinical and laboratory investigations to e.g. confirm the high diagnostic relevance of presence or absence of colour and colour related features.

It is important to note that our method is bound by the representativeness of the training data. The method can benefit from consensus colour labelling (naming) of training data by different dermatologists.

References

- [1] A Madooei, MS Drew, M Sadeghi, and MS Atkins. Automated pre-processing method for dermoscopic images and its application to pigmented skin lesion segmentation. In *20th Color and Imaging Conf.*, pages 158–163, 2012.
- [2] A Madooei, MS Drew, M Sadeghi, and MS Atkins. Intrinsic melanin and hemoglobin colour components for skin lesion malignancy detection. *Medical image computing and computer-assisted intervention: MICCAI*, 15:315–322, 2012.
- [3] ME Celebi, H Iyatomi, WV Stoecker, RH Moss, HS Rabinovitz, G Argenziano, and HP Soyer. Automatic detection of blue-white veil and related structures in dermoscopy images. *Computerized Medical Imaging and Graphics*, 32(8):670–677, 2008.
- [4] S Seidenari, G Pellacani, and C Grana. Computer description of colours in dermoscopic melanocytic lesion images reproducing clinical assessment. *British J. of Dermatology*, 149(3):523–529, 2003.
- [5] World Health Organization. Skin cancers. Available at: <http://www.who.int/uv/faq/skincancer/en/index1.html>, Accessed August 22, 2012.
- [6] Spectrophotometric intracutaneous analysis: a new technique for imaging pigmented skin lesions. 146.
- [7] R Mackie. *An Illustrated Guide to the Recognition of Early Malignant Melanoma*. Schering, 1986.
- [8] W Stolz, A Riemann, AB Cagnetta, and et al. ABCD rule of dermoscopy: a new practical method for early recognition of melanoma. *European J. of Dermatology*, 1994.
- [9] H Soyer, G Argenziano, S Chimenti, and V Ruocco. Dermoscopy of pigmented skin lesions. *EJD*, 11(3):270–276, 2001.
- [10] S Menzies. *Atlas of Surface Microscopy of Pigmented Skin Lesions*. McGraw Hill, 2003.
- [11] G Argenziano, HP Soyer, and et al. *Interactive atlas of dermoscopy*. Edra Medical Publishing & New Media, Milan, Italy, 2000.
- [12] FA Bahmer and et al. Terminology in surface microscopy. *Journal of the American Academy of Dermatology*, 23:1159–1162, 1990.
- [13] ER Landa and MD Fairchild. Charting color from the eye of the beholder. *American Scientist*, 93(5):436–443, 2005.
- [14] P Centore. An open-source inversion algorithm for the munsell renotation. *Color Research & Application*, 37(6):455464, 2012.
- [15] M Fairchild. *Color Appearance Models*. John Wiley & Sons, 2005.
- [16] P Centore. The Munsell toolbox. Code available at <http://www.99main.com/~centore/MunsellToolbox/MunsellToolbox.html>, 2010.
- [17] Munsell Color Science Laboratory website, accessed April 9, 2013. Complete color FAQ.. Available at <http://www.cis.rit.edu/mcsl/outreach/faq.php?catnum=5/#696/>.
- [18] R Achanta and et al. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(11):2274–2282, 2012.
- [19] B Georgescu and CM Christoudias. The edge detection and image segmentation (EDISON) system. Code available at <http://coewww.rutgers.edu/riul/research/code/EDISON/>, 2003.
- [20] G Sharma, W Wu, and EN Dalal. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1):2130, 2005.
- [21] A Madooei, MS Drew, M Sadeghi, and MS Atkins. Automatic detection of blue-white veil by discrete colour matching on dermoscopy images. To appear in *Medical image computing and computer-assisted intervention: MICCAI*, 2013.




Figure 3. (a-m): Easy images; (n-r): challenging images – The green border indicates veil areas detected by proposed method.