

# Segmentation of light and dark hair in dermoscopic images: a hybrid approach using a universal kernel

Nhi H. Nguyen<sup>a,b</sup>, Tim K. Lee<sup>\*a,b,c</sup>, M. Stella Atkins<sup>a</sup>

<sup>a</sup>School of Computing Science, Simon Fraser University, Burnaby BC, Canada;

<sup>b</sup>Cancer Control Research Program, BC Cancer Agency, Vancouver BC, Canada;

<sup>c</sup>Photomedicine Institute, Department of Dermatology and Skin Science, University of British Columbia and Vancouver Coastal Health Research Institute, Vancouver BC, Canada

## ABSTRACT

The main challenge in an automated diagnostic system for the early diagnosis of melanoma is the correct segmentation and classification of moles, often occluded by hair in images obtained with a dermoscope. Hair occlusion causes segmentation algorithms to fail to identify the correct nevus border, and can cause errors in estimating texture measures. We present a new method to identify hair in dermoscopic images using a universal approach, which can segment both dark and light hair without prior knowledge of the hair type. First, the hair is amplified using a universal matched filtering kernel, which generates strong responses for both dark and light hair without prejudice. Then we apply local entropy thresholding on the response to get a raw binary hair mask. This hair mask is then refined and verified by a model checker. The model checker includes a combination of image processing (morphological thinning and label propagation) and mathematical (Gaussian curve fitting) techniques. The result is a clean hair mask which can be used to segment and disocclude the hair in the image, preparing it for further segmentation and analysis. Application on real dermoscopic images yields good results for thick hair of varying colours, from light to dark. The algorithm also performs well on skin images with a mixture of both dark and light hair, which was not previously possible with previous hair segmentation algorithms.

**Keywords:** hair segmentation, image segmentation, matched filtering, dermoscopic images

## 1. INTRODUCTION

The presence of hair in dermoscopic images is a challenge in an automated diagnostic system for the early diagnosis of melanoma. Hair occlusion can cause segmentation algorithms to perform badly due to confusion between hair and mole borders. Hence, the removal of hair is an important pre-processing step in such systems.

Existing solutions to this problem include the popular DullRazor<sup>1</sup>, which uses morphological operations to remove thick, dark hair from dermoscopic images. It works by applying morphological closing operation separately on all three channels of the RGB input image, with three structure elements in the horizontal, diagonal and vertical directions. The maximum result out of the three structure elements is taken as the response at each pixel. A hair mask is generated for each channel by hard-thresholding the difference between the original channel value and the response at each pixel. Then hair pixels are repainted using linear interpolation with two nearby non-hair pixels. While being generally effective, this software has a limitation on its applicability: it only works on thick dark hair, as the authors clearly state in their paper.

Schmid-Saugeon *et al.*<sup>2</sup> use a similar approach but the morphological closing operation is applied to the three components of the  $L^*u^*v$  uniform colour space. However, hair removal is mentioned very briefly as a preprocessing step in the diagnostic system for pigmented skin lesions in this paper, without any study on the accuracy of the hair disocclusion scheme, as the authors favour physical hair removal over image processing techniques. This method has the same applicability limitation as DullRazor.

Fleming *et al.*<sup>3</sup> take a different approach to segmenting hair in dermoscopic images. Based on the observation that hairs are long, straight curvilinear structures with relatively constant width and curvature, they develop a tracing algorithm to follow and connect the curvilinear hair segments. First, they apply Steger's line detection algorithm<sup>4</sup> on the input to get a list of line segments present in the image. Then linear discriminant analysis is applied on the feature vector for each line

segment in order to classify it as hair or not hair. The output of this step is a list of hair segments, with possible gaps. A graph search is done to fill the gaps and connect the hair segments. Then length filtering is implemented as the final step to reject the remaining false positives. This algorithm, despite being complicated, does not address the problem of intersecting hair. It also does not mention the case of light hair, and the results they present are on a small number of images with dark hair only.

Zhou *et al.*<sup>5</sup> implements a similar algorithm. Their approach uses the same method to identify the line segments (Steger's line detection algorithm<sup>4</sup>), but it adds an extra step to detect intersections and analyze them to reconnect the line segments broken at intersections. Then it fits a curve to each hair segment and evaluates the fit to reject non-hair segments. The result in this paper clearly shows all the light hairs being left out in the segmentation; in other words, this algorithm is applicable to dark hair only.

Most existing methods for dermoscopic hair segmentation in the literature overlook or ignore the case of hair lighter than the background, the skin. This is partly due to the difficulty of the problem since light hair has similar characteristics to skin lines (ie. lines that form the skin patterns or dermatoglyphics). Furthermore, light hair is usually thinner than dark hair, which makes segmentation difficult; in fact, most existing hair segmentation algorithms focus on thick dark hair only.

Our method aims at segmenting light and dark hair in dermoscopic images using a universal kernel. It differs from existing methods in its ability to detect hair without prior knowledge of the hair colour. It has the advantage of detecting hair of different colours in one run, even ones in contrasting colours such as black and blonde, which has not been done by any existing method.

## 2. METHOD

The hair segmentation process consists of two phases. The first phase is to identify the centerline of possible hair. The output of this phase is a skeletonised hair mask. The second phase takes this skeletonised mask and reconstructs the hair, as well as eliminates the false centrelines detected in phase 1. Figure 1 is a visual overview of the algorithm.

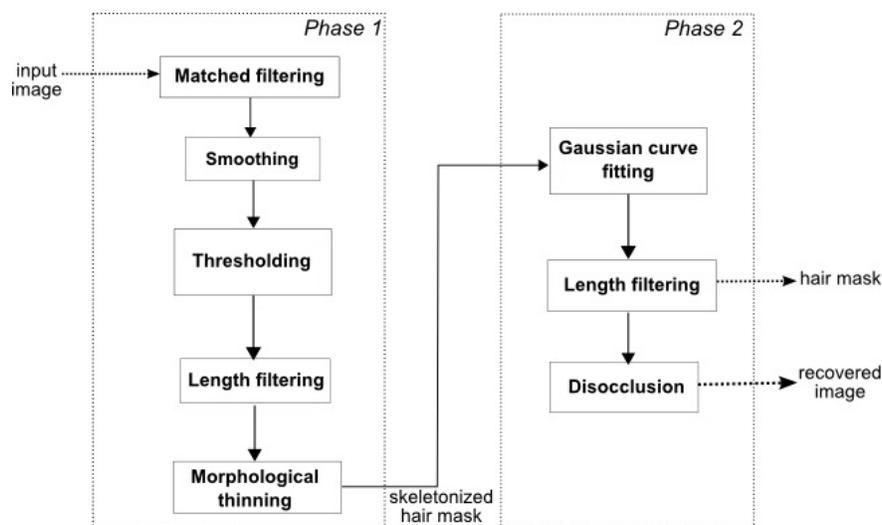


Figure 1: The hair disocclusion process

We make the following assumptions in this paper:

- (i) Hair has a Gaussian cross section profile with a constant width of  $2\sigma$  pixels along its length, where  $\sigma$  is the standard deviation of the Gaussian function.

- (ii) Hair is piece-wise linear, ie. it can be approximated with line segments.

## 2.2 Phase 1: Identify possible hairs

Phase 1 of the algorithm identifies possible hairs based on the observed properties of hair, most notably the Gaussian nature of the cross section intensity profile. First, the input image is converted to grayscale and all possible hairs are amplified using a modified matched filtering process. Then smoothing is done on the filter response image to bridge the “ghost” response to the true positive response, preparing it for the morphological thinning step. This is due to the output of our modified matched filter having three peaks at each matched position (section 2.2.1). After smoothing, the response is thresholded to produce a raw hair mask. The method used is a modified version of the local entropy thresholding algorithm<sup>6</sup>. The raw hair mask is then refined by a length filtering process to eliminate small blobs that are unlikely to be hair, based on size criteria. The final step in phase 1 is morphological thinning. It is applied on the refined hair mask to produce a skeletonised version containing only the centerlines of possible hairs.

### 2.2.1 Matched filtering

The first step in the segmentation pipeline is applying matched filtering on a grayscale version of the input image. Dermoscopic images are converted to grayscale before processing by taking the V channel of the HSV version of the input. Based on the observation that the intensity profile of hair in dermoscopic images can be approximated by a Gaussian curve, we use a Gaussian kernel similar to the one in Chaudhuri *et al.*'s paper<sup>7</sup> to amplify the hair. The kernel can be mathematically expressed as:

$$K(x, y) = -e^{-\frac{x^2}{2\sigma^2}} \quad (1)$$

for  $|y| < L/2$ , with  $L$  being the length of the piecewise linear hair segment. We use a series of 18 rotating kernels (10 degree angular resolution) to cover all possible directions of hair. Response at each pixel is taken as the maximum response out of the 18 kernels.

Amplification is done by a modified matched filtering process. General matched filtering can only amplify signals with shapes matching the kernel shape exactly. Our modified filtering is designed to match the shape of both dark hair and light intensity profiles by including an extra step: after normalised cross correlation, the absolute value of the response is taken as the result. Given that normalised cross correlation gives us the correlation coefficients, and that we are aiming at amplifying similar shapes as well as completely opposite shapes (Gaussian and inversed Gaussian), this process provides a response in which all hair is amplified, regardless of colour. Figure 2(a) and 2(b) show an example of an input image and its response.

The angle of the maximum response is also saved in a separate matrix. This matrix contains information about the direction of the strongest signal at each pixel. It can be used to calculate the direction as well as curvature of the hair.

### 2.2.2 Processing the filter response: smoothing, thresholding and length filtering

A raw hair mask is produced from the filter response using Gaussian smoothing, thresholding and length filtering.

The design of the kernel and taking the absolute value of the filter response makes it possible to amplify both dark and light hair using a universal kernel. However, it also introduces two lower peaks alongside the main peak where the signal matches. Applying thresholding right on this response would create many false positives, resulting in an incorrect hair mask where each correctly identified hair is accompanied by two false hairs running in parallel. The smoothing step aims at solving this problem. We use Gaussian smoothing to bridge the “ghost” responses to the peak response, creating a single peak instead of three peaks at each matched point. Gaussian smoothing is chosen for its low-pass property, especially in areas where the three-peak response occurs. In such cases, it attenuates the highest peak, bringing it closer to the lower peaks and increases the intensity levels of the “valleys” between the peaks, effectively bridging the three peaks into one. This prepares the response for the thresholding step.

We use the local entropy thresholding method in Chanwimaluang and Fan's paper<sup>6</sup>, based on the original idea of Pal and Pal<sup>8</sup>, to threshold the filter response to get the raw binary hair mask. Local entropy is a measure of information present in

a binary image. It is introduced by Pal and Pal, with the intention of preserving the information present in the spatial structure in an image, which is not possible with previous histogram-based measures of entropy. It uses a co-occurrence matrix, which is a well-known measure of correlation of a particular property between pixels in an image. The idea is to select a threshold that maximises the total second order local entropy of the object and the background, based on the co-occurrence matrix. The output of this step is a raw hair mask containing all the possible hairs.

Length filtering is applied on this raw hair mask to eliminate false positives with small size. These false positives are usually resulting from pigment network segments and skin lines, whose intensity profiles resemble that of hair. However, their length is usually less than that of normal hair. Observing that hair is thin, long structures, we use length as a criterion to filter out objects in the hair mask whose length is less than a certain threshold. Length filtering is done using label propagation with the eight-connected neighbourhood.

Figure 2(c) demonstrates a raw hair mask resulting from this step.

### 2.2.3 Skeletonising of the hair mask

The raw hair mask outlines the positions of all the possible hairs. However, it does not ensure correct size; and it also tolerates false positives. Hence this raw hair mask is subject to verification and reconstruction (phase 2). The skeletonising step reduces the raw hair mask to single pixel lines denoting the center lines of hairs. We use morphological thinning for this purpose. As the "ghost" responses are symmetric and the smoothing step connects them to the real peak, morphological thinning returns the centerline of the hair, given that the response from hair is strong enough.

Figure 2(d) shows a skeletonised hair mask, obtained from the raw hair mask in Figure 2(c).

## 2.3 Phase 2: Verification and reconstruction of hair

The output of phase 1 is a skeletonised hair mask which is an estimation of the hair center lines. We use this skeletonised hair mask as the input to our verification step, which utilizes a model-based checker. The model is built on the observation that hair has a Gaussian intensity profile (or its inverse) and has length exceeding a certain threshold.

### 2.3.1 Gaussian curve fitting

The skeleton obtained at the end of phase 1 provides a guideline as to where the center of the hair is. Using it as an input, we follow each pixel on the skeleton, take  $2\sigma$  pixels in the perpendicular direction of the hair centering at this pixel and check if the  $2\sigma+1$  pixels conform to a Gaussian curve. If they do, we expand the skeleton at that point to cover the whole hair width ( $2\sigma$ ); otherwise we reject the point (ie. label it as background). This process rejects most of the false positives and adds more true positives, which were not in the raw hair mask, to the refined hair mask.

The expanding process may leave single-pixel gaps within the reconstructed hair pixels, as only the pixels in the perpendicular direction to the hair at each centerline pixel are processed. To overcome this problem, we implement a post-processing step after reconstruction, checking for gaps and fill them. A background pixel is classified as a gap when it is surrounded by two hair pixels in the horizontal, vertical or diagonal direction. When we find a gap pixel, we relabel it as a hair pixel.

### 2.3.2 Length filtering

After reconstructing the hair, length filtering is applied once more on the hair mask to eliminate short segments. This step uses the same algorithm as the length filtering step in phase 1 (section 2.2.2).

### 2.3.3 Disocclusion

Disocclusion is the process of recovering the pixels originally occluded by hair. As this project focuses on segmentation rather than disocclusion, we use a simple algorithm to recover the pixels. We adopt the method used by DullRazor<sup>1</sup>, which is linear interpolation of the neighbouring non-hair pixels to repaint the hair pixels.

Figure 2 demonstrates the whole process on an image with phantom hair of both dark and light colour and varying width. The segmented hair is highlighted in black in Figure 2(e). Comparing to the ground truth in Figure 2(f), we can see that the segmentation result contains a small number of false positives (skin classified as hair) and false negatives (hair pixels classified as skin). The false positives result from strong noise in the background, whereas the false negatives are due to

hair intersections, too small width of hair and hairs running too close to each other in parallel. The intersection misses can be explained by the “halo” effect of the filter, with which response from the stronger hair overwrites the other one (note that the “halo” effect, ie. three-peak response, creates a response wider than the original signal, thus the gaps at intersections). The case of hairs running too close to each other in parallel is also a consequence of this effect. In the case of fine hairs, the small width causes the verification step to fail and hence, the hair is rejected.

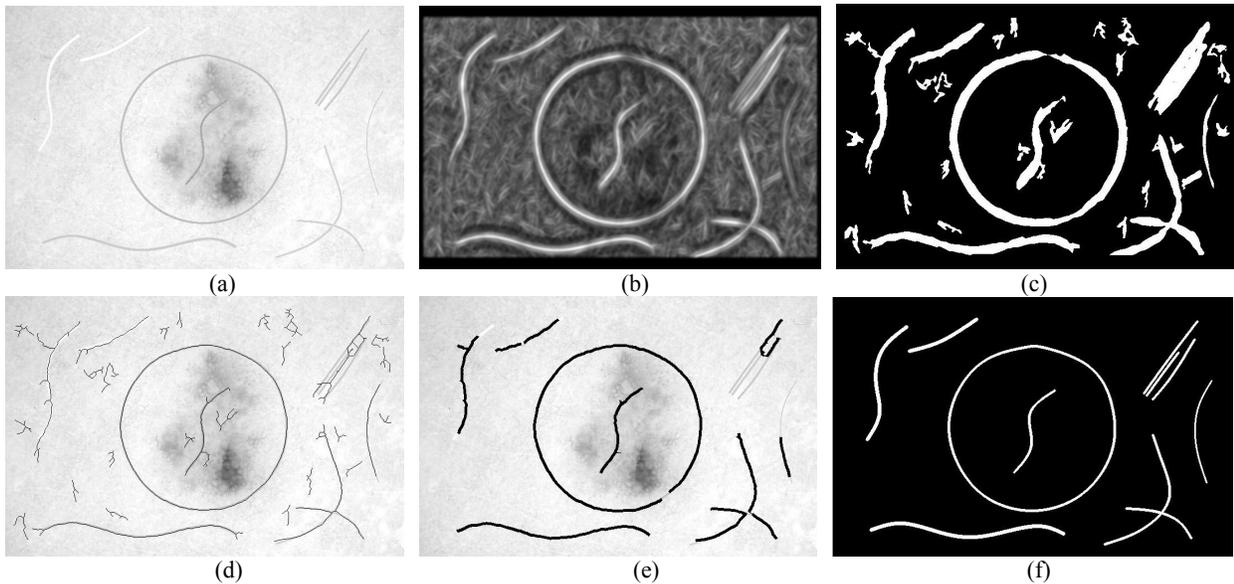


Figure 2: Demonstration of steps in the algorithm:

(a): The input image with dark and light phantom hair, (b) The filter response, (c) The raw hair mask, (d) The skeleton hair mask, (e) The final hair mask superimposed on the input, (f) The ground truth.

### 3. EXPERIMENTS

#### 3.1 Using synthetic images

We designed a pilot experiment with a synthetic input image with phantom hair to quantitatively evaluate the algorithm. The input image was built by superimposing phantom hairs on a real dermoscopic image with a melanoma and no hair (size: 728x480 pixels). We used two circular phantom hairs, one darker and one lighter than the background. The two circular phantom hairs were selected as they were the simplest case that covered all possible directions, varying curvature and both dark and light colours of hair. The hairs were 4 pixels wide and had perfect Gaussian profiles of different depths. The skin background was non-uniform with the presence of a nevus. We used the following set of parameters:

- Kernel size:  $\sigma = 2$  (the assumed half-width of hair),  $L = 20$  (the length of the piecewise linear segment)
- Angular resolution:  $\alpha = 10$  degrees
- Minimum length: The minimum length of hair used in length filtering is 50 pixels.

The synthetic input, its corresponding filter response, hair mask and output of the pilot experiment are shown in Figure 3. We can see in Figure 3(d) that most of the hairs were correctly detected by the algorithm, save a small number of false positives and false negatives. The errors were due to strong noise in the background; which became problematic in the thresholding and verification steps.

The segmentation result was compared with the ground truth to calculate statistics such as true positive, false positive, true negative and false negative. From these values we computed statistical measures of the performance of the algorithm: specificity, sensitivity, accuracy and diagnostic accuracy. Specificity, sensitivity and accuracy are well known statistical measures of the performance of a binary classification system. However; in problems where true negatives

outnumber other measures (true positives, false positives and false negatives) such as hair segmentation, any metric that uses true negative in its computation is biased toward this true negative rate. For this reason, we decided to use diagnostic accuracy as the main measure for the accuracy of the algorithm. Diagnostic accuracy is calculated as<sup>9</sup>:

$$\text{DiagnosticAccuracy} = \frac{TP}{TP + FP + FN} \times 100\% \quad (2)$$

where TP, FP, FN stand for the number of True Positives, False Positives and False Negatives respectively.

Quantitative analysis on the result yielded 73% diagnostic accuracy.

Another example for synthetic hairs is shown in Figure 2.

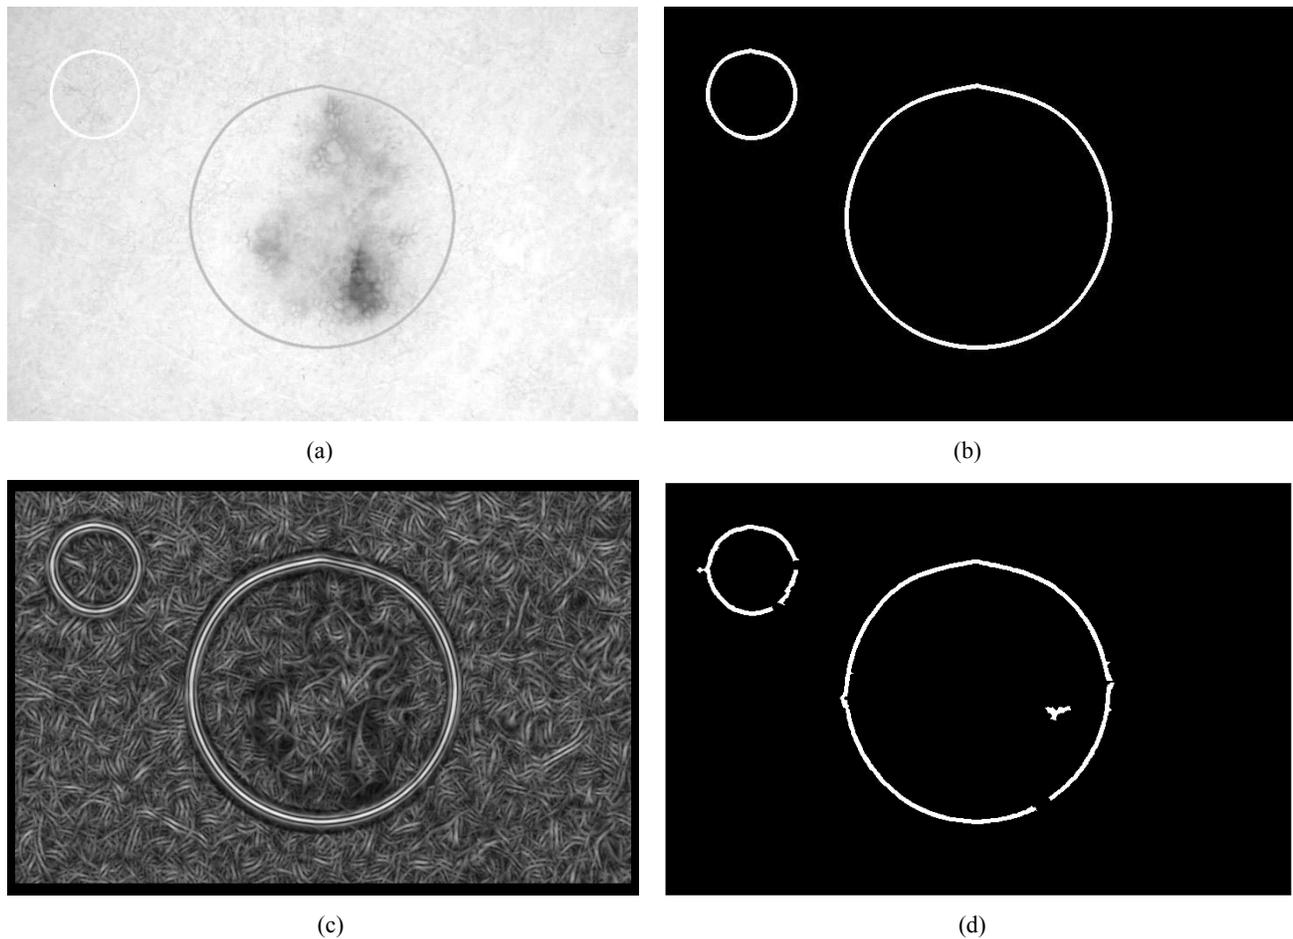


Figure 3: Pilot experiment: (a): The input image with dark and light phantom hair, (b) The ground truth, (c) The filter response, (d) The hair mask as produced by the algorithm

### 3.2 Using real dermoscopic images

Application on real dermoscopic images (taken from two dermatology atlases<sup>10,11</sup>) produced satisfying results for thick hair of both dark and light colours. The algorithm worked very well on images with sparse, mixed colour hair. No other existing method is capable of this. Figure 4 and Figure 5 demonstrate the algorithm on two dermoscopic images, with the segmented hairs highlighted in green and superimposed on the input image for easy visual assessment.

Limitation of this method includes applicability on fine and dense hair with many intersections. The presence of a large, prominent pigment network is also a challenge.

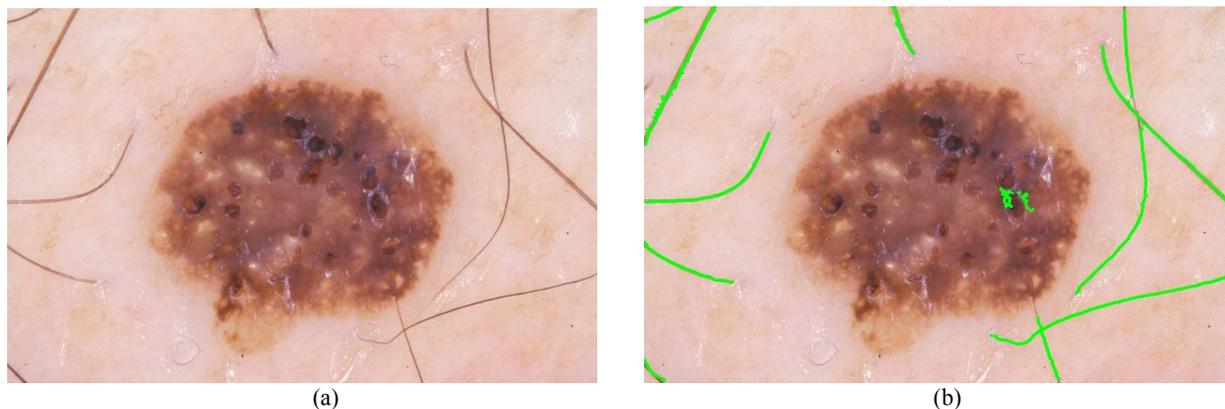


Figure 4: Example of results on a dermoscopic image with dark hair: (a) Input, (b) Output



Figure 5: Example of results on a dermoscopic image with mixed (light and dark) hair: (a) Input, (b) Output

#### 4. CONCLUSION

Our method of segmenting hair in dermoscopic images using a universal kernel is capable of segmenting both dark and light hair of constant width in dermoscopic images, without prior knowledge of the hair colour. It has its limitation in the cases of fine hairs and hairs with many intersections. However, these limitations can be solved by allowing for slight width variation in the verification step, and a region growing algorithm to bridge gaps at intersection.

#### REFERENCES

- [1] Lee, T., Ng, V., Gallagher, R., Coldman, A. and McLean, D., "Dullrazor®: A software approach to hair removal from images," *Computers in Biology and Medicine*, 27(6), 533-543 (1997).
- [2] Schmid-Saugeon, P., Guillod, J. and Thiran, J.P., "Towards a computer-aided diagnosis system for pigmented skin lesions," *Computerized Medical Imaging and Graphics*, 27(1), 65-78 (2003).
- [3] Fleming, M.G., Steger, C., Zhang, J., Gao, J., Cognetta, A.B., Pollak, I. and Dyer, C.R., "Techniques for a structural analysis of dermatoscopic imagery," *Computerized medical imaging and graphics*, 22(5), 375-389 (1998).

- [4] Steger, C., "An unbiased detector of curvilinear structures," IEEE Transactions on Pattern Analysis and Machine Intelligence, 20(2), 113-125 (1998).
- [5] Zhou, H., Chen, M., Gass, R., Rehg, J.M., Ferris, L., Ho, J. and Drogowski, L., "Feature-preserving artifact removal from dermoscopy images," Proc. SPIE 6914, 69141B-69141B (2008).
- [6] Chanwimaluang, T. and Fan, G., "An efficient blood vessel detection algorithm for retinal images using local entropy thresholding," Proceedings of the 2003 International Symposium on Circuits and Systems, 5, (2003).
- [7] Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M. and Goldbaum, M., "Detection of blood vessels in retinal images using two-dimensional matched filters," IEEE Transactions on Medical Imaging, 8(3), (1989).
- [8] Pal, N.R. and Pal, S.K., "Entropic thresholding," Signal Processing, 16, 97-108 (1989).
- [9] Grin, C.M., Kopf, A.W., Welkovich, B., Bart, R.S. and Levenstein, M.J., "Accuracy in the clinical diagnosis of malignant melanoma," Archives of Dermatology, 126(6), 763 (1990).
- [10] Argenziano, G., Soyer, HP, De Giorgi, V., Piccolo, D., Carli, P., Delfino, M. et al, [Interactive atlas of dermoscopy. Dermoscopy: a tutorial (Book) and CD-ROM], Edra-Medical Publishing & New Media, Milan (Italy) (2000).
- [11] Marghoob, A.A., Braun, R.P. and A.W. Kopf, [Atlas of dermoscopy], Taylor & Francis (2005).