

Hair Enhancement in Dermoscopic Images using Dual-Channel Quaternion Tubularness Filters and MRF-based Multi-Label Optimization

Hengameh Mirzaalian, Tim K Lee, *Member, IEEE*, and Ghassan Hamarneh, *Senior Member, IEEE*

Abstract—Hair occlusion is one of the main challenges facing automatic lesion segmentation and feature extraction for skin cancer applications. We propose a novel method for simultaneously enhancing both light and dark hairs with variable widths, from dermoscopic images, without the prior knowledge of the hair color. We measure hair tubularness using a quaternion color curvature filter. We extract optimal hair features (tubularness, scale, and orientation) using Markov random field theory and multi-label optimization. We also develop a novel dual channel matched filter to enhance hair pixels in the dermoscopic images while suppressing irrelevant skin pixels. We evaluate the hair enhancement capabilities of our method on hair-occluded images generated via our new hair simulation algorithm. Since hair enhancement is an intermediate step in a computer aided diagnosis system for analyzing dermoscopic images, we validate our method and compare it to other methods by studying its effect on: (i) hair segmentation accuracy, (ii) image inpainting quality, and (iii) image classification accuracy. The validation results on 40 real clinical dermoscopic images and 94 synthetic data demonstrate that our approach outperforms competing hair enhancement methods.

Index Terms—Melanoma; hair enhancement; hair segmentation; quaternion tubularness filters; light and dark objects.

I. INTRODUCTION

Malignant melanoma is one of the fastest growing cancers among the world’s white population with an average 3% increase in incidence for the last four decades. In the USA and Canada alone, there were about 73,720 cases of melanoma in 2011[1]. Early diagnosis of melanoma can lead to life-saving treatment. Dermoscopy is a noninvasive imaging technique that provides rich features of skin lesions, such as shape, color and texture, with excellent diagnostic value. To improve the efficacy and efficiency of early detection of melanoma, a considerable amount of research has focused on automating the segmentation, feature extraction and classification of dermoscopic images [2]. A review of the existing computerized methods to analyze single skin lesions in dermoscopic images has been recently reported in [3]. However, the existence



Fig. 1: Example dermoscopic images of skin lesions with no hair, dark hairs, and mixed colored (dark and light) hairs.

of hairs that occlude the lesion in the image is a common problem that complicates the segmentation process and renders subsequent feature detection and classification unreliable (Figure 1) [4], [5], [6], [7], [8]. Hair occlusion thus remains a key challenge facing computer aided diagnosis of malignant melanoma.

There have been limited published papers on hair enhancement and segmentation (i.e. masking out the occluding hair). Lee et al. [5], Xie et al. [7], Fiorese et al. [9], and Schmid-Saugeon et al. [10] used morphological top-hat (MTH) operators of different sizes to enhance hair structures. MTH was also used in [11] to prepare a set of candidate pixels. Then, clusters of connected neighbouring pixels with similar intensities, which were approximately linear, were classified as hair. Nguyen et al. [6] proposed a matched filtering method, with a fixed hair width assumption, to enhance both the dark and light hairs. Kiani and Sherfat [12] applied Prewitt filters. Abbas et al. [13], [14] applied derivative of Gaussian filters to detect hair structures. Zhou et al. [8] and Fleming et al. [4] modelled hairs as a long relatively straight curvilinear structure, with constant width and curvature, and applied the Steger’s line detection algorithm [15]. All the mentioned methods operated on scalarized (e.g. luminance channel) dermoscopic images.

We note that the appearance of vasculature in biomedical images resembles to some degree the appearance of hair in dermoscopic images. In particular, both structures are curvilinear with varying width. Despite notable differences (e.g. vessels bifurcate whereas hairs overlap; vessel images are typically single channel whereas dermoscopic images are color), the hair enhancement and segmentation problems stand to benefit from methods for the detection and analysis of tubular structures, as has been witnessed in state-of-the-art research on vascular image analysis [16], [17].

In this paper, as our first contribution, we present a novel methodological extension to the vesselness filter proposed by Frangi et al. [18]. To the best of our knowledge, Frangi et al.’s

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

H. Mirzaalian was a Scholar supported by the CIHR Skin Research Training Centre.

T. K. Lee is with the Photomedicine Institute, Department of Dermatology and Skin Science, University of British Columbia, and Vancouver Coastal Health Research Institute, and the Cancer Control Research, BC Cancer Agency, BC Canada.

H. Mirzaalian and G. Hamarneh are with the Medical Image Analysis Lab at the School of Computing Science, Simon Fraser University, BC, Canada.

filter and variants thereof have been applied to enhance images containing *only* dark or *only* bright objects. However, there is a problem with applying this filter in the presence of mixed color tubular structures, e.g. presence of mixed color hairs, *both* light and dark with respect to the background, which is common in skin images. The naïve approach of applying the Frangi filter twice in succession; first to enhance light structures and then apply it again to enhance dark structures (or vice versa), and then combine the two enhanced images would lead to erroneous enhancement (explained technically and qualitatively in Section II-B). Our solution to this problem is to take advantage of a novel dual-channel matched filter (Section III-B). We then apply it to a problem related to dermatological image analysis.

As our second contribution, we extend our previous work on Markov random field (MRF) based scale selection [19] to estimate the optimal features of the hairs such as: orientation, tubularness, and scale (where optimality is defined with respect to MRF energy with unary and binary terms). In [19], we showed that MRF-based vessel scale selection gives more accurate detection results for vascular structures in retinal images compared to the traditional scale selection (at each pixel, select the scale with maximum response)¹ (Section III-A).

As our third contribution, we propose a new (realistic) hair-on-skin simulation system (Section III-C). Since hairs are very thin with spatially-varying width, manually preparing ground truth masks with accurate hair width for a large number of hair pixels would be exorbitantly tedious. Therefore, generating simulated hair occluded images with known ground truth would be beneficial in this area for our research as well as for other research groups. We also provide an accompanying software, which is publicly and freely available for download².

As our fourth contribution, we make use of color information for hair enhancement (Section II-A). The idea of quaternion color curvature was proposed by Shi et al. [20]. All previous hair disocclusion methods have been applied to scalarized (i.e. single channel) dermoscopic images only. This is the first work that uses all color channels for hair image analysis (in dermatology applications). This is done by adopting the quaternion based Hessian of color images. So the novelty is (only) in the application domain.

We evaluate the goodness of our hair enhancement algorithm and compare its performance quantitatively (in terms of segmentation, inpainting, and classification accuracy measures) to other competing methods on both simulated data and 40 real dermoscopic images (Section IV).

II. BACKGROUND

We begin by giving a brief review of quaternion tubularness filtering (Section II-A), followed by highlighting its weakness in the presence of simultaneous light and dark tubular objects (Section II-B).




Fig. 2: Tubularness filter responses of a dermoscopic image. (a) A dermoscopic image. The small region pointed to by the white arrow is zoomed up at the lower-right corner. (b) Tubularness response ν (1). (c-d) The tubularness responses after considering the sign tests in (3).

A. Quaternion Tubularness Filter

Frangi et al. [18] proposed to measure the tubularness $\nu(\mathbf{x}, s)$ at pixel $\mathbf{x} = (x, y)$ for K different scales $s \in \{s^1, \dots, s^K\}$ using:

$$\nu(\mathbf{x}, s) = \exp\left(-\frac{R^2(\mathbf{x}, s)}{2\beta^2}\right) \left(1 - \exp\left(-\frac{s^l S^2(\mathbf{x}, s)}{2c^2}\right)\right) \quad (1)$$

$$R(\mathbf{x}, s) = \frac{\lambda_1(\mathbf{x}, s)}{\lambda_2(\mathbf{x}, s)} \quad S(\mathbf{x}, s) = \sqrt{\sum_{i \leq 2} \lambda_i^2(\mathbf{x}, s)} \quad (2)$$

where $\lambda_i(\mathbf{x}, s)$, $i = 1, 2$ ($|\lambda_1| \leq |\lambda_2|$), are the eigenvalues of the Hessian matrix of image I computed at scale s . R and S are measures of blobness and second order structureness, respectively. β and c are the parameters that control the sensitivity of the filter to the measures R and S . The parameter l is introduced by Lindeberg [21] to make the filter responses of the different scales comparable.

Figure 2(b) shows the computed tubularness for the dermoscopic image in Figure 2(a) based on (1). Since the sign of the largest eigenvalue is an indicator of the brightness or darkness of the pixels (i.e. dark-on-bright vs. bright-on-dark), and since dermoscopic images may contain both dark and light hairs, the following sign tests are used to determine the tubularness of the light, ν^- , and dark, ν^+ , hair structures [18]³:

$$\nu^-(\mathbf{x}, s) = \begin{cases} 0 & \text{if } \lambda_2(\mathbf{x}, s) > 0 \\ \nu(\mathbf{x}, s) & \text{if } \lambda_2(\mathbf{x}, s) < 0 \end{cases} \quad (3)$$

$$\nu^+(\mathbf{x}, s) = \begin{cases} 0 & \text{if } \lambda_2(\mathbf{x}, s) < 0 \\ \nu(\mathbf{x}, s) & \text{if } \lambda_2(\mathbf{x}, s) > 0 \end{cases}$$

Note that $\nu(\mathbf{x}, s) = \nu^-(\mathbf{x}, s) + \nu^+(\mathbf{x}, s)$. Figures 2(c-d) show examples of the computed ν^- and ν^+ for a dermoscopic image occluded by dark hairs. It can be seen that the dark hairs are enhanced correctly in the filter response ν^+ (Figure 2(d)).

To make use of color information in the computed tubularness (3), the tubularness in (1) can be measured using the eigenvalues of the quaternion Hessian matrix of the image [20], $H_Q = H_r \cdot i + H_g \cdot j + H_b \cdot k$, where H_r , H_g , and H_b are Hessian matrices of the RGB channels and i , j , and k are three imaginary units. The eigenvalues of H_Q can be obtained by quaternion singular value decomposition (QSVD). QSVD is a generalization of singular value decomposition (SVD) of real or complex numbers to quaternion numbers. Examples of the measured tubularness (3) based on SVD and QSVD are shown in Figures 3. It can be clearly seen that the hairs are

¹<http://www.cs.sfu.ca/~hamarneh/software/mrfscalesel/>

²www.cs.sfu.ca/~hamarneh/software/hairsim

³Superscripts $-$ and $+$ indicate the sign of the largest eigenvalue of the pixels being inside the light and dark tubular objects, respectively.




Fig. 3: Hair tubularness ν^+ (3) based on SVD and QSVD. (a) A dermoscopic image with simulated hairs (Section III-C explains our simulator). The small region pointed to by the white arrow is scaled up for clarity. (b) Luminance channel of image (a) used to compute the SVD-based tubularness. (c-d) The tubularness computed using SVD and QSVD. (e-h) Similar to (a-d), but for a real hair in a dermoscopic image.

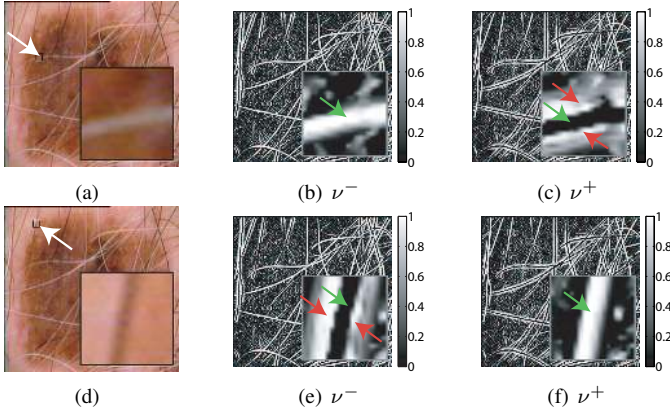


Fig. 4: Problem of enhancing mixed-colored hairs. (a) Close up of a white hair (white arrow). (b-c) Tubularness ν^- , and ν^+ computed according to the sign tests in (3). (d-f) Similar to (a-c), but show a close up of an image region occluded by a dark hair. The green arrows in (b, c, e, and f) indicate regions of correct enhancements or correct suppression of the image in ν^- and ν^+ , whereas the red arrows indicate regions of erroneous enhancement.

better delineated when considering color information in the RGB channels.

B. Tubularness of Mixed-Colored Objects

Applying the tubularness filter in Section II-A to skin images exhibiting both light and dark (mixed) hairs causes incorrect enhancement. As an example, a light and a dark hair along with their tubularness ν^- and ν^+ are shown in Figure 4. It can be seen that although the light hairs are enhanced as desired in ν^- (indicated by the green arrows in Figure 4(b)), the skin pixels along two narrow bands sandwiching the dark hair are erroneously enhanced as well (red arrows in Figure 4(e)). Similar erroneous enhancement can be noticed for ν^+ (Figure 4(c)). In short, neither sign tests in (3) work for both light and dark hairs simultaneously. Thus, enhancing or detecting light and dark colored hairs sequentially and separately in skin images with mixed colored hairs will generate incorrect results. A mathematical exposition of this problem is detailed next.

Consider the following simplified example. Given an image containing a bright line and a dark line parallel to the y -axis (resembling mixed color hair), as shown in Figure 5(a), we model the intensity profiles of the lines by Gaussian functions: $f^i = G(x - x_i, s_o)$, where $i \in \{1, 2\}$, x_1 and x_2 represent the x -coordinates of the centerlines of the light and dark lines, respectively, and s_o is the standard deviation of the Gaussian profiles.

The eigenvalues of the Hessian matrix of f^i functions at scale s are $\lambda_1^i(x, s) = 0$ and $\lambda_2^i(x, s) = \partial^2 f^i / \partial x^2$ (second derivative of f^i), which lead to high tubularness (1) for the pixels within the following intervals (Figure 5):

$$\Gamma_{\text{in}}^i(x) : x \in |x - x_i| < F_1 \quad (4)$$

$$\Gamma_{\text{out}}^i(x) : x \in 0.5\sqrt{2(s + s_o)} < |x - x_i| < F_2 \quad (5)$$

where $x = x_i \pm F_1 = x_i \pm 0.5\sqrt{2(s + s_o)}$ and $x = x_i \pm F_2 = x_i \pm 0.5\sqrt{6(s + s_o)}$ are the inflection points of f^i and $\partial f^i / \partial x$ (first derivative of f^i), respectively.

Applying the negative sign test (ν^- in (3)), the tubularness is enhanced not only for the pixels within Γ_{in}^1 , but also within Γ_{out}^2 (the blue curves in Figure 5(b)). Similarly, considering the positive sign test causes false enhancements within Γ_{out}^1 (the pink curves in Figure 5(b)).

Therefore, in the presence of mixed-color lines, an appropriate tubularness filter must be designed to enhance the pixels only within Γ_{in}^i , not in Γ_{out}^i .

A naïve approach would be to attempt to detect erroneous enhancements appearing as parallel bands. However, this is not a robust strategy, as hairs themselves are often parallel to each other as shown in Figure 1.

To this end, we propose the dual channel filters in Section III-B, which lead to an image with the enhanced tubularness only within Γ_{in}^i .

III. METHOD

Our hair enhancement framework consists of two main stages: MRF-based extraction of hair features using the multi-scale quaternion tubularness filter responses (Section III-A) and the suppression of incorrect enhancements using our novel dual matched filters (Section III-B). In this section, we first provide details about these two steps. Next, we describe the details of our hair simulator, which allows for a comprehensive validation of the goodness of detecting hairs with varying thickness (Section III-C).

A. MRF-based Hair Features

We extend multi-scale quaternion tubularness to compute hair features: tubularness, scale, and orientation. Note that we estimate the hair direction $\theta(\mathbf{x}, s)$ as the angle between the x -axis of the image and the eigenvector corresponding to $\lambda_1(\mathbf{x}, s)$, which points along the minimum principal curvature direction of the tubular object (minimum intensity variation).

Given the multi-scale light tubularness $\nu^-(\mathbf{x}, s)$ and tubularness direction $\theta(\mathbf{x}, s)$, tubularness, scale, and orientation at each pixel \mathbf{x} of image I can be computed using the maximum tubularness response (MTR) as:




Fig. 5: The issue of simultaneously detecting dark and light lines by applying Frangi et al. filter: An image containing a light and a dark ridge (simulating mixed color hair shafts) is shown in (a). The intensity profiles of the lines are modeled by Gaussian functions (orange curve in (b)). Filter responses ν^- and ν^+ computed according to the sign tests in (3) enhance tubularness of the pixels within Γ_{in}^i and Γ_{out}^i , which results in erroneous enhancement inside Γ_{out}^1 and Γ_{out}^2 . Applying the dual channel filters in Section III-B, we discard the erroneous enhancement within Γ_{out}^i as shown in (c).

$$\begin{aligned}
 \nu_{MTR}^-(\mathbf{x}) &= \max_{s \in \{s^1, \dots, s^K\}} \nu^-(\mathbf{x}, s) \\
 \mathcal{S}_{MTR}^-(\mathbf{x}) &= \arg \max_{s \in \{s^1, \dots, s^K\}} \nu^-(\mathbf{x}, s) \\
 \theta_{MTR}^-(\mathbf{x}) &= \theta(\mathbf{x}, \mathcal{S}_{MTR}^-(\mathbf{x}))
 \end{aligned} \quad (6)$$

where $\nu_{MTR}^-(\mathbf{x})$, $\mathcal{S}_{MTR}^-(\mathbf{x})$, and $\theta_{MTR}^-(\mathbf{x})$ denote the MTR-based tubularness, scale, and orientation of the light structures, respectively. By replacing ν^- by ν^+ in (6), similar features can be defined for dark structures denoted by $\nu_{MTR}^+(\mathbf{x})$, $\mathcal{S}_{MTR}^+(\mathbf{x})$, and $\theta_{MTR}^+(\mathbf{x})$, respectively.

In this paper, rather than using MTR-based hair features, we utilize MRF-based vessel scale-selection [19] and extend it to extract MRF-based hair orientation, scale, and tubularness features.

The MRF-based approach [19] formulates the scale selection as a graph labelling, in which image I is modelled by a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where vertices \mathcal{V} correspond to the pixels $\mathbf{x}_p \in I$ and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ connect two neighbouring pixels $(\mathbf{x}_p, \mathbf{x}_q)$, and the labels are chosen to be correspond to the different scales $\ell \in \{s^1, s^2, \dots, s^K\}$. The MRF-optimization seeks the labelling ℓ for each pixel \mathbf{x}_p by minimizing an energy functional of the form:

$$E(\ell) = \eta \sum_{\mathbf{x}_p \in \mathcal{V}} \varphi_p(\ell_p) + (1 - \eta) \sum_{(\mathbf{x}_p, \mathbf{x}_q) \in \mathcal{E}} \varphi_{pq}(\ell_p, \ell_q) \quad (7)$$

where $\ell_q \equiv \ell(\mathbf{x}_q)$ represents the label of the q^{th} pixel of the image, φ_p is a unary term which measures the likelihood of labelling \mathbf{x}_p with a specific label disregarding the labels of any of the neighbours, φ_{pq} is the spatial regularization term penalizing different label configurations of neighbouring vertices, and η is the weight of the spatial regularization term.

Therefore, the MRF-based scales of the light hairs \mathcal{S}_{MRF}^- are obtained by minimizing (7) with respect to ℓ , where $\nu^-(\mathbf{x}, s)$

of the different scales are set as the unary terms:

$$\begin{aligned}
 \mathcal{S}_{MRF}^- &\equiv \ell = \arg \min_{\ell \in \{s^1, \dots, s^K\}} E(\ell) \\
 \text{s.t. } \varphi_p(\ell_p = s^i) &= \nu^-(\mathbf{x}_p, s^i)
 \end{aligned} \quad (8)$$

Given \mathcal{S}_{MRF}^- , we introduce MRF-based tubularness ν_{MRF}^- and orientation θ_{MRF}^- at each pixel as:

$$\begin{aligned}
 \nu_{MRF}^-(\mathbf{x}) &= \nu^-(\mathbf{x}, \mathcal{S}_{MRF}^-(\mathbf{x})) \\
 \theta_{MRF}^-(\mathbf{x}) &= \theta^-(\mathbf{x}, \mathcal{S}_{MRF}^-(\mathbf{x})).
 \end{aligned} \quad (9)$$

By replacing ν^- by ν^+ in (9), we can define MRF-based features of the dark structures: tubularness $\nu_{MRF}^+(\mathbf{x})$, scale $\mathcal{S}_{MRF}^+(\mathbf{x})$, and direction θ_{MRF}^+ . We use α -expansion graph cuts [22] to solve the multi-label MRF optimization. The regularization penalty is set as the Pott's energy $\varphi_{pq}(\ell_p, \ell_q) = \min\{1, |\ell_p - \ell_q|\}$ due to its discontinuity preserving characteristic [23]. We compute the optimal value of η in (8) empirically; using the evaluation methods mentioned in Section IV, we validate the goodness of the final enhanced images for different values of $\eta \in [0, 1]$. Our results indicate that $\eta = 0.8$ gives the highest segmentation accuracy.

Figure 6 compares the MTR-based features (tubularness, scales, and orientations) (Figures 6(e-g)) with the MRF-based features (Figures 6(h-j)). It can be noticed that the undesirable blurring of the borders of the hairs is reduced substantially in ν_{MRF}^+ compared with ν_{MTR}^+ . In addition, the MRF-based scales and orientations (e.g. \mathcal{S}_{MRF}^+ and θ_{MRF}^+) are more regularized compared with MTR-based features (e.g. \mathcal{S}_{MTR}^+ and θ_{MTR}^+).

B. Dual Matched Filters for Detecting Light and Dark Hairs

Figure 7 summarizes the magnitude response of the tubularness computed as ν^- and ν^+ for a pixel **b** inside the hair and two neighboring skin pixels **a** and **c**. Note that high responses are obtained in:

- $\nu^+(\mathbf{a})$, $\nu^-(\mathbf{b})$, and $\nu^+(\mathbf{c})$ for light tubular structure.
- $\nu^-(\mathbf{a})$, $\nu^+(\mathbf{b})$, and $\nu^-(\mathbf{c})$ for dark tubular structure.

We now wish to model and take advantage of the above tubularness appearance characteristics of light and dark hairs. To this end, we model a light hair at scale s and orientation θ , in each of the ν^- and ν^+ as M_1 and M_2 , respectively:

$$\begin{aligned}
 M_1(\mathbf{x}, s, \theta) &= \exp\left(-\frac{x'^4 + \gamma y'^4}{2s^2}\right) \\
 M_2(\mathbf{x}, s, \theta) &= \exp\left(-\frac{x'^4 + \gamma y'^4}{2(3s)^2}\right) \left(1 - \Pi\left(\frac{y'}{0.5s}\right)\right)
 \end{aligned} \quad (10)$$

where $\mathbf{x} = (x, y)$, $x' = x \cos(\theta) + y \sin(\theta)$, $y' = -x \sin(\theta) + y \cos(\theta)$, Π is the rectangular (or boxcar) function, and γ specifies the ellipticity of the hyper-Gaussian function. Note how the functions M_1 and M_2 match the responses summarized in Figure 5, modulo scale and rotation. The choice of the parameters $3s$ and $0.5s$ in M_2 are related to the width of the hair at scale s ; the standard deviation $3s$ is used to cover the bandwidth of the pixels within Γ_{out} (5) and to block the filter responses within Γ_{in} (5), we set the width of the boxcar to s ($=2 \times 0.5s$).

Examples of these matched filters for different scales s and orientations θ are shown in Figure 8. We now combine the




Fig. 6: Comparing MTR-based features of hairs (tubularness, scales, and orientations) with the MRF-based features. (a-d) Calculated tubularness of the image in Figure 2(a) according to (3) at four scales. (e) Maximum tubularness response ν_{MTR}^+ using (6) for the tubularness of the different scales in (a-d). (f-g) Estimated MTR-based scales (S_{MTR}^+) and MTR-based tubularness direction (θ_{MTR}^+). Different scales and orientations are represented with different colors. (h-j) Estimated MRF-based tubularness (ν_{MRF}^+), scale (S_{MRF}^+), and tubularness direction (θ_{MRF}^+), respectively (Section III-A). Note how the undesirable blurring around the hairs in (e) (ν_{MTR}^+) is reduced in (h) (ν_{MRF}^+). Furthermore, in (i-j), the MRF-based scales and orientations (S_{MRF}^+ and θ_{MRF}^+) are more regularized compared with the corresponding MTR-based features (S_{MTR}^+ and θ_{MTR}^+) in (f-g).

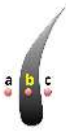
	ν^-			ν^+			
	a	b	c	a	b	c	
Light Hair	L	H	L	H	L	H	
Dark Hair	H	L	H	L	H	L	

Fig. 7: The magnitude response of the tubularness computed in ν^- and ν^+ for a pixel inside the hair, b, and two neighbouring skin pixels, a and c, respectively. H: High and L: Low. The green cells represent correct enhancements or suppressions in ν^- and ν^+ , whereas the red cells represent erroneous enhancements.




Fig. 8: Our proposed dual channel filterbank to detect light and dark hairs. (a) First channel (M_1) and (b) second channel (M_2) with three different scales (rows) and orientations (columns).

correlation values between ν^- and M_1 and between ν^+ and M_2 to compute the likelihood P^- of a pixel \mathbf{x}_p belonging to a light (or P^+ for dark) hair as:

$$P^-(\mathbf{x}_p|s, \theta) = \bar{K}_{\mathbf{x}_q}(\nu^-(\mathbf{x}_q, s), M_1(\mathbf{x}_q, s, \theta)) \times \bar{K}_{\mathbf{x}_q}(\nu^+(\mathbf{x}_q, s), M_2(\mathbf{x}_q, s, \theta)) \quad (11)$$

$$\text{s.t. } \mathbf{x}_q \in |\mathbf{x}_p - \mathbf{x}_q| \leq 3s$$

$$\bar{K}_D(f, g) = \frac{K_D(f, g)}{\sqrt{K_D(f, f)}\sqrt{K_D(g, g)}}$$

$$K_D(f, g) = \sum_D (f - \bar{f})(g - \bar{g})$$

where $\mathbf{x}_q \in |\mathbf{x}_p - \mathbf{x}_q| \leq 3s$ represents neighbouring pixels of \mathbf{x}_p in radius of $3s$, $\bar{K}_D(f, g)$ is a kernel computed as the normalized cross correlation between f and g over the domain D , and \bar{f} is the average of f over D . Since we compute the mutual correlation over the two channels in (11), we call this operator as a dual-channel matched filter.

For each pixel \mathbf{x}_p , we compute (11) using the MRF-based estimation of the scales and orientations (Section III-A):

$$P^-(\mathbf{x}_p) = P^-(\mathbf{x}_p|s = S_{\text{MRF}}^-, \theta = \theta_{\text{MRF}}^-). \quad (12)$$

In a similar manner, we define the probability of a pixel being part of a dark hair, P^+ . The final enhancement of the light and dark hairs in the image is thus enabled by our proposed dual-channel, quaternion tubularness, MRF-based formulation.

C. Hair Occlusion Simulator

Since hairs are very thin with spatially-varying width, manually preparing a ground truth mask with accurate hair-width for a large number of hair pixels would be exorbitantly tedious (let alone for a large number of images). Therefore, we have developed, for the first time, a hair simulation algorithm to accurately and comprehensively validate the goodness of detecting hairs with varying thickness.

At a high level, our simulator is based on, first, creating a hair mask (e.g. Figure 9(c)) and then, blending it with a hair free image (e.g. Figure 9(f)), i.e. polluting the hair free image with artificial hair. This way, we are able to control both the hair free image, and the added artificial hair. We are able to choose different type of skin images (e.g. different color or tan). Also, we are able to control the added hair, e.g. the number, color, thickness. There are two stages involved in preparing the masks: first, create curves (degree 2 splines) that are fitted to the medial axis of each hair; and second; thicken these medial curves with a spatially varying hair width.

The first step, preparing medial curves of the hairs, is carried out using medial curves that are either detected manually by a human operator on real dermoscopic images, or generated automatically via a random curve synthesizer (Figure 9(a)). The second step, hair-thickening, is as follows. Each hair medial curve is morphologically dilated with a disk structuring element of varying radius. Let A and B denote the end points of a hair in the mask (Figure 9(b)). The radius r of the disk at pixel p along the hair is obtained by:

$$r(p) = \min\left\{T, \frac{\Gamma(p, A)}{200}, \frac{\Gamma(p, B)}{200}\right\} \quad (13)$$




Fig. 9: Simulating hair-occluded images. (a) Hair medial axes (degree 2 splines). (b) An example of thickening a hair by dilation using (13). (c) Binary mask of dilated hairs. (d) Gaussian filtering of (c). (e) Hair coloring of (d) using (14). (f) Hair-free dermoscopic image. (g) Pixel-wise multiplication of (f) by an inverted version of (d). (h) Sum of (g) and (e). (i-l) Examples of simulated images.

where $\Gamma(p, q)$ measures the geodesic distance between two points p and q . The basic idea of using (13) is to gradually increase the size of the disk, and hence the hair width, from the end points toward the midpoint of the hair, with the maximum radius T reached at the mid-point. We set $T = 4$ pixels, yielding a maximum width of about 0.1 mm, the average width of a human hair.

The generated ground truth binary mask containing thick hairs is denoted by I_{gt} (Figure 9(c)). I_{gt} is used to pollute hair-free dermoscopic images, I^0 (Figure 9(f)), resulting in a simulated hair-occluded image with a realistic appearance denoted by I^H (Figure 9(i-l)). Corrupting the hair-free images is completed as follows. To control the color of the hair, each of the RGB channels of I^H is created individually, then the three channels are composited together. For example, the red channel I_R^H is obtained as follows:

$$I_R^H = I_R^0(\mathbb{1} - I_{\text{gt}} * G_\sigma) + C_R(I_{\text{gt}} * G_\sigma) \quad (14)$$

where $\mathbb{1}$ is a matrix of ones equal to the size of the mask I_{gt} , G_σ is a Gaussian kernel of width σ , and $C = [C_R, C_G, C_B]$ is used to control the color of the generated hairs. Equation (14) represents a linear combination between the colors of the hair free image I^0 and the overlaid hair C . Since we do not want the skin color in I^0 to appear at the centre line of the hair, applying the Gaussian function G_σ , we control the weight of the colors of I^0 and C , e.g. gradually increase the weight of I^0 from the hair skeleton toward the outer side, and vice versa, gradually decrease the weight of C .

Example simulated images are shown in Figure 9. Our hair simulation software is provided publicly for download⁴. Note

⁴www.cs.sfu.ca/~hamarneh/software/hairsim (or directly from: www.cs.sfu.ca/~hamarneh/software/hairsim/restricted/HairSim.zip using username: researcher and password: skin0hair).

that, we sampled colors of the hairs from our dataset [24], which includes yellow, brown, white, black, gray ones, and used them to simulate the hair colors.

IV. RESULTS

We evaluate our hair enhancement method on 40 real dermoscopic images as well as 94 synthetic data generated by our simulator (Section III-C), with 0.033 mm/pixel isotropic resolution; 10% of real and 50% of synthetic images are occluded by mixed color hairs. In our implementation, the free parameters in (1) are set as follows: $\beta=0.7$, $c=0.05$, and $l=1.1$, $K=4$, and $[s^1, s^K] \in [0.033, 0.2]$ mm to cover the maximum line width of the hairs. We refer to the work by Bouattour et al. [26] as a theoretical framework for optimal choice of the free parameters. The same scale values $[s^1, s^K]$ are used to construct the dual channel filters (Section III-B) and the size of the filters are set as $3s$. Note that in our implementation we did not apply any preprocessing.

Since image enhancement is generally an intermediate step, and not the final objective, in an automatic dermoscopic image analysis pipeline, we evaluate the effect of our hair enhancement method on a number of subsequent objectives. First, we start by comparing the effect of the enhancement on hair segmentation accuracies for different methods (Section IV-A). As feature extraction is crucial for computer aided diagnosis (CAD) of lesions in dermoscopic images, and since hair occlusion corrupts the extracted features, in the second set of experiments (Section IV-B), we evaluate the effect of our hair enhancement algorithm and several competing methods on an inpainting method designed to rectify the effect of hair disocclusion by estimating the underlying missing pixel data. Third, we evaluate the performance of a CAD system by calculating lesion feature classification accuracies, after the images are processed using our proposed method vs. competing methods (Section IV-C).

A. Segmentation Evaluation

We begin by showing qualitative results of hair enhancement of real dermoscopic images in Figure 10. The figure compares the use of morphological top-hat operator based method (MTH) [5], [7], matched filters (MFL) [6], parallel line detectors (PL) [25], the original Frangi et al.'s filters (SVD) [18], the quaternion version of SVD (QSVD) [20], and our proposed method using the dual channel filters without and with considering quaternion information denoted by DCT (dual channel tubularness) and DCQT (dual channel quaternion tubularness), respectively. For fair comparison, we extend all the methods to perform multi-scale enhancement of light, dark, and mixed-colored hairs, which improved the methods' performance. For example, focusing on MTH, we sequentially apply close and top-hat operators to enhance dark and light hairs, respectively, using structuring elements of different sizes. Note in Figure 10 how MTH and MFL fail to detect several hairs. Note also how SVD and QSVD are overwhelmed with false positives.

To quantitatively evaluate the quality of the results from applying different hair enhancement methods, we first produce




Fig. 10: Hair enhancement using different methods and the resulting hair segmentation. (a) Two real dermoscopic images (first and third rows) and their corresponding hair mask ground truth (second and fourth rows). (b-g) The first and third rows represent the enhanced images in (a) using (b) MTH [7], [5], (c) MFL [6], (d) PL [25], (e) SVD [18], (f) QSVD [20], (g) DCT, and (h) DCQT, and the second and fourth rows represent the binary images resulting Otsu’s thresholding of the enhanced images.

TABLE I: Numerical comparison between different hair enhancement methods in terms of accuracy, precision, and f-score (15) of the segmented hairs, and running time of the methods (Section IV-A).

Method	Results for real data					Results for synthetic data				
	ACC _{seg}	PSN _{seg}	REC _{seg}	FSC _{seg}	Time	ACC _{seg}	PSN _{seg}	REC _{seg}	FSC _{seg}	Time
MTH	0.77±0.01	0.41±0.03	0.34±0.10	0.28±0.06	0.03±0.01	0.82±0.01	0.45±0.05	0.42±0.02	0.41±0.02	0.03±0.02
MFL	0.70±0.01	0.36±0.01	0.47±0.08	0.30±0.02	0.09±0.02	0.69±0.02	0.31±0.04	0.51±0.02	0.43±0.02	0.09±0.10
PL	0.77±0.01	0.44±0.01	0.37±0.06	0.28±0.03	0.33±0.31	0.72±0.02	0.31±0.05	0.47±0.02	0.41±0.05	0.20±0.06
SVD	0.65±0.01	0.34±0.01	0.67±0.06	0.38±0.02	0.01±0.00	0.44±0.01	0.25±0.03	0.87±0.04	0.34±0.06	0.01±0.01
QSVD	0.63±0.01	0.33±0.01	0.68±0.06	0.38±0.02	0.03±0.01	0.40±0.01	0.30±0.01	0.88±0.01	0.38±0.02	0.03±0.01
DCT _{MTR}	0.81±0.01	0.56±0.02	0.52±0.07	0.50±0.07	0.10±0.03	0.84±0.02	0.54±0.05	0.60±0.02	0.57±0.03	0.05±0.07
DCT _{MRF}	0.84±0.01	0.59±0.02	0.52±0.07	0.52±0.07	0.12±0.03	0.84±0.01	0.55±0.05	0.60±0.02	0.57±0.03	0.07±0.15
DCQT _{MTR}	0.85±0.01	0.58±0.02	0.54±0.08	0.54±0.07	0.12±0.07	0.84±0.01	0.65±0.06	0.62±0.03	0.58±0.03	0.07±0.28
DCQT _{MRF}	0.86±0.01	0.59±0.02	0.56±0.08	0.55±0.07	0.14±0.05	0.85±0.01	0.66±0.06	0.62±0.03	0.60±0.03	0.12±0.07

a binary hair mask (i.e. segmentation) by applying Otsu’s thresholding on the enhanced image [11]. Then, similar to the post processing step by Lee et al. [5], each pixel in the hair region of the mask is checked to ensure that it is located within a thin long structure, i.e. the hair structure; otherwise, it is rejected as noise. Let I_{bw} and I_{gt} denote the filtered and the ground truth hair masks, respectively. We calculate the accuracy (ACC_{seg}), precision (PSN_{seg}), recall (REC_{seg}), and f-score (FSC_{seg}) of the segmented hairs according to:

$$ACC_{seg} = (TP + TN) / (TP + FN + TN + FP) \quad (15)$$

$$PSN_{seg} = TP / (TP + FP)$$

$$REC_{seg} = TP / (TP + FN)$$

$$FSC_{seg} = 2TP / (2TP + FP + FN)$$

where:

$$TP = \sum_{\mathbf{x}} I_{bw}(\mathbf{x})I_{gt}(\mathbf{x}), TN = \sum_{\mathbf{x}} (1 - I_{bw}(\mathbf{x}))(1 - I_{gt}(\mathbf{x}))$$

$$FP = \sum_{\mathbf{x}} (1 - I_{bw}(\mathbf{x}))I_{gt}(\mathbf{x}), FN = \sum_{\mathbf{x}} I_{bw}(\mathbf{x})(1 - I_{gt}(\mathbf{x})).$$

The numerical results in Table I indicate that the best accuracy and precision are achieved by our method. Note that

DCQT_{MTR} and DCQT_{MRF} represent a variant of our proposed method considering MTR-based and MRF-based hair features (Section III-A), respectively.

A comparison between SVD and QSVD in Table I declares that QSVD provides inferior results. Note that QSVD’s performance is better than SVD *if* either light or dark (i.e. not mixed color) hair exist (as shown in Figure 3). However, in the presence of *erroneous* colored structures, which is happened in the presence of mixed color tubular objects, then QSVD will enhance them too. As mentioned in Section II-B, erroneous structures appear as bands parallel to the hairs as shown in Figure 4. Due to the erroneously enhances in these bands by QSVD, which are undesirable, we see inferior results by QSVD in Table I compared with SVD. When these erroneous structures are suppressed using the dual channel filters, it can be seen that a better performance is achieved by DCQT compared with DCT.

The running time of the different hair enhancement methods are reported in the last column of Table I. Although our method is computationally expensive, it achieves more accurate results. In a CAD system for automatic analysis of the dermoscopic images there are no strict requirements for speed, whereas




Fig. 11: The inpainted images resulting from applying different hair enhancement approaches. (a) Sample of ground truth hair masks (top) and hair-occluded images (bottom). (b-f) Binary hair masks (top) and their corresponding inpainted images (bottom) resulting from (b) MTH [7], [5], (c) MFL [6], (d) PL [25], (e) SVD [18], (f) QSVD [20], (g) DCT, and (h) DCQT. Note that the last four rows represent the results for simulated hair-occluded images.

accuracy is paramount.

B. Inpainting Evaluation

Since computer-aided diagnosis (CAD) systems that analyze dermoscopic images rely on performing feature extraction for lesion classification [3], [9], a method that properly discloses hair while accurately retaining the original skin-features is desirable. Applying the inpainting approach in [5], in Figure 11, we show examples of the inpainted images resulting from applying competing hair enhancement methods (MTH [7], [5], MFL [6], PL [25], SVD [18], QSVD [20]) and our proposed method (DCT and DCQT). Note how our method is more faithful in retaining skin lesion features in the inpainted image while removing the hairs.

To provide quantitative results of the goodness of the inpainted images, using the same notation in Section

III-C, we introduce a root mean square error (RMSE) computed between the inpainted image I^H and the hair free image I^0 for the pixels occluded by the simulated hairs $\mathcal{D} = \{\mathbf{x} | I_{\text{gt}}(\mathbf{x}) = 1\}$:

$$\text{RMSE} = \sqrt{\sum_{\mathbf{x} \in \mathcal{D}} (I^H(\mathbf{x}) - I^0(\mathbf{x}))^2 / |\mathcal{D}|}. \quad (16)$$

Our underlying assumption is that better hair enhancement approach will identify the hair-occluded pixel, without over- or under-segmentation, and therefore the subsequent inpainting method will estimate the intensities for these pixels only, without unnecessarily estimating the intensity of over-segmented pixels nor ignoring under-segmented pixels. To validate that our enhancement method leads to better inpainting, we present numerical RMSE results in Table II. The results indicate that our method has the least RMSE compared to the other

methods. Note that RMSE cannot be measured for real hair occluded images since the ground truth of the image to be restored is not known.

C. Classification Evaluation

Here, we validate the goodness of the inpainted (hair segmented and the pixels underneath reconstructed as in [5]) images by comparing the skin patch-classification accuracies of a trained random forest classifier. In particular, we test the presence or absence of pigmentation, such as pigment network or streaks in a 60×60 skin patch. Small inpainted patches are tested in this validation step instead of the entire skin images because small patches will limit the variations of the lesion features and, hence, allows the classifier to build better models for feature classification. Our underlying assumption is that features extracted from a better inpainted image is more accurate and result in a better estimation of the class of the skin patch. Note that the patches are selected randomly from the images.

Our feature vector to perform classification is a set of 32 means and variances of pixel intensities in different color channels of the skin patches in RGB, HSI, and Luv and the means and variances of the wavelet transform coefficients of the grayscale skin patches using Haar wavelets and two decomposition levels [27]. We first, train a classifier for skin patch classification (absence vs. presence of pigment network or streaks) using hair-free image patches. Then, to classify a patch of hair-occluded image, we feed its corresponding inpainted image to the random forest classifier (the number of trees is set to 80).

The classification evaluation is performed with real and synthetic images separately. Using 40 images from the atlas of dermoscopy by Argenziano et al. [24], we randomly select 320 hair-free 60×60 -pixel patches to train our classifier. Then a different set of 320 unseen 60×60 -pixel hair-occluded patches are used for testing. When evaluating 94 synthetic images, we use 752 hair-free 60×60 -pixel patches and their corresponding hair-occluded patches (generated using our hair simulator mentioned in Section III-C) for training and testing, respectively. The number of absence vs. presence cases for real and synthetic data are: 128 vs. 192 and 292 vs. 460, respectively. In Table III, we report the precision, recall, and f-score of the random forest classifier, which are computed as follows:

$$\begin{aligned} \text{PSN}_{\text{RF}} &= \text{TP}/(\text{TP}+\text{FP}) \\ \text{REC}_{\text{RF}} &= \text{TP}/(\text{TP}+\text{FN}) \\ \text{FSC}_{\text{RF}} &= 2\text{TP}/(\text{TP}+\text{FN}+\text{FN}) \end{aligned} \quad (17)$$

where TP, FP, and FN are the numbers of true positives, false positives, and false negatives of the lesion classifier. The numerical results in Table III indicate that our method outperforms the other methods.

V. CONCLUSIONS

Handling occluding hairs is an important pre-processing step for all computer-based applications related with skin

TABLE II: Numerical comparison between different hair enhancement methods in terms of RMSE (16) (Section IV-B).

Results for synthetic data	
Method	RMSE (16)
MTH	0.1373±0.0148
MFL	0.2532±0.0502
PL	0.2416±0.0560
SVD	0.2221±0.0211
QSVD	0.2517±0.0246
DCT _{MTR}	0.1476±0.0066
DCT _{MRF}	0.1475±0.0066
DCQT _{MTR}	0.0922±0.0094
DCQT _{MRF}	0.0905±0.0094

TABLE III: Numerical comparison between different hair enhancement methods in terms of the lesion classification precision (17) and recall(17) (Section IV-C). Note that the first row represents the results for the original image, the hair **occluded one**.

Method	Results for real data			Results for synthetic data		
	PSN _{RF}	REC _{RF}	FSC _{RF}	PSN _{RF}	REC _{RF}	FSC _{RF}
-	0.79	0.74	0.76	0.76	0.62	0.68
MTH	0.84	0.86	0.85	0.89	0.90	0.89
MFL	0.74	0.82	0.78	0.83	0.79	0.80
PL	0.82	0.84	0.82	0.87	0.89	0.88
SVD	0.77	0.85	0.80	0.93	0.90	0.91
QSVD	0.76	0.85	0.80	0.91	0.90	0.90
DCT _{MTR}	0.80	0.85	0.82	0.95	0.90	0.92
DCT _{MRF}	0.81	0.85	0.83	0.95	0.91	0.93
DCQT _{MTR}	0.89	0.91	0.90	0.95	0.90	0.92
DCQT _{MRF}	0.90	0.93	0.90	0.96	0.93	0.94

images. Unfortunately, existing programs that enhance and remove dark hairs cannot be applied to both light and dark color hairs without generating artifacts. In this paper, we extended the classic Frangi filter, which was designed to handle single color objects, to dual color objects. The optimal Frangi features (scale, tubularness and orientation) were determined by MRF-modelling and multi-label optimization. In addition, the problem was formulated using a true color approach based on quaternion Hessian matrix. As the result, we constructed a dual-channel quaternion tubularness filter, which achieved the best accuracy when it was compared qualitatively and quantitatively to other state of art programs using both synthetic and real skin images.

Although we applied the proposed filter to one skin application, which is related to feature extraction and lesion classification in dermoscopic images, however, the technique can be applied to many other skin applications. One of the examples is to perform hair-counting for hair-loss therapy [28], as many of the patients have gray hairs. It might have non-medical application, e.g. in images containing both light and dark tubular structures, e.g. aerial images. Furthermore, we have developed a hair simulation algorithm, which allowed us to properly validate the goodness of detecting the varying thickness of the hairs. We make our simulation software publicly available online at www.cs.sfu.ca/~hamarneh/software/hairsim. Our simulator is useful for validating and benchmarking future works on hair enhancement and segmentation.

In the current work, to compute the tubularness at each pixel, we consider information of its immediate neighbouring pixels. However, it might be helpful to take more global approaches into account, e.g. applying minimal path based

methods to leverage likelihoods in the direction of the hair structures, which is considered as a part of our future work.

ACKNOWLEDGEMENTS

This work was supported in part by a scholarship from CIHR Skin Research Training Centre and by grants from the Natural Sciences and Engineering Research Council of Canada, and Canadian Dermatology Foundation. Also, we would like to thank Dr. Giuseppe Argenziano at the University of Naples for sharing the dermoscopy data [24].

REFERENCES

- [1] R. Siegel, D. Naishadham, and A. Jemal. Cancer statistics. *A Cancer Journal for Clinicians*, 62:10–29, 2012. 1
- [2] E. Celebi, H. Iyatomi, G. Schaefer, and V. Stoecker. Lesion border detection in dermoscopy images. *CMIG*, 33(2):148 – 153, 2009. 1
- [3] K. Korotkov and R. Garcia. Computerized analysis of pigmented skin lesions: A review. *Artificial Intelligence in Medicine*, 56:69–90, 2012. 1, 8
- [4] M. Fleming, C. Steger, J. Zhang, J. Gao, A.B. Cognetta, I. Pollak, and C.R. Dye. Techniques for a structural analysis of dermatoscopic imagery. *CMIG*, 22(5):375–389, 1998. 1
- [5] T. Lee, V. Nguyen, R. Gallagher, A. Coldman, and D. McLean. DullRazor: a software approach to hair removal from images. *Comput Biol Med*, 27(3):533–543, 1997. 1, 6, 7, 8, 9
- [6] N. Nguyen, T. Lee, and S. Atkins. Segmentation of light and dark hair in dermoscopic images: a hybrid approach using a universal kernel. *SPIE*, 7623:148 – 153, 2010. 1, 6, 7, 8
- [7] F. Xie, S. Qin, Z. Jiang, and R. Meng. PDE-based unsupervised repair of hair-occluded information in dermoscopy images of melanoma. *CMIG*, 33(4):275 – 282, 2009. 1, 6, 7, 8
- [8] H. Zhou, M. Chen, R. Gass, M. Rehg, L. Ferris, J. Ho, and L. Drogowski. Feature-preserving artifact removal from dermoscopy images. *SPIE*, 6914(1):69141B, 2008. 1
- [9] M. Fiorese, E. Peserico, and A. Silletti. Virtualshave: Automated hair removal from digital dermatoscopic images. *IEEE EMBC*, pages 4378–4381, 2011. 1, 8
- [10] P. Schmid-Saugeon, J. Guilloid, and J. Thiran. Towards a computer-aided diagnosis system for pigmented skin lesions. *CMIG*, 27(1):65 – 78, 2003. 1
- [11] A. Afonso and M. Silveira. Hair detection in dermoscopic images using percolation. *IEEE EMBC*, pages 4378–4381, 2012. 1, 7
- [12] K. Kiani and A. Sharafat. E-shaver: An improved DullRazor; for digitally removing dark and light-colored hairs in dermoscopic images. *Computers in Biology and Medicine*, 41(3):139–145, 2011. 1
- [13] Q. Abbas, M. Celebi, and I. Garcia. Hair removal methods: A comparative study for dermoscopy images. *Biomedical Signal Processing and Control*, 6:395–404, 2011. 1
- [14] Q. Abbas, I. Garcia, M. Celebi, and W. Ahmad. A feature-preserving hair removal algorithm for dermoscopy images. *Skin Research and Technology*, 19:27–36, 2013. 1
- [15] C. Steger. An unbiased detector of curvilinear structures. *IEEE TPAMI*, 20:113–125, 1998. 1
- [16] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM Computing Surveys*, 36:81–121, 2004. 1
- [17] D. Lesage, E. Angelini, I. Bloch, and G. Funka-Lea. A review of 3d vessel lumen segmentation techniques: Models, features and extraction schemes. *MIA*, 13(6):819 – 845, 2009. 1
- [18] F. Frangi, J. Niessen, K. Vincken, and M. Viergever. Multiscale vessel enhancement filtering. *MICCAI*, pages 130–137, 1998. 1, 2, 6, 7, 8
- [19] H. Mirzaalian and G. Hamarneh. Vessel scale selection using MRF optimization. In *IEEE CVPR*, pages 3273–3278, 2010. 2, 4
- [20] L. Shi, B. Funt, and G. Hamarneh. Quaternion color curvature. pages 338–341, 2008. 2, 6, 7, 8
- [21] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *IEEE CVPR*, pages 465–470, 1996. 2
- [22] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE TPAMI*, 23:2001, 1999. 4
- [23] O. Veksler. Graph cut based optimization for mrfs with truncated convex priors. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1 –8, june 2007. 4
- [24] G. Argenziano, H. Soyer, V. Giorgio, D. Piccolo, P. Carli, M. Delfinoand A. Ferrari, R. Hofmann, D. Massi, G. Mazzocchetti, M. Scalvenzi, and H. Wolf. Interactive atlas of dermoscopy. *Edra Medical Publishing and New Media*, 2000. 6, 9, 10
- [25] T. Koller, G. Gerig, G. Szekely, and D. Dettwiler. Multiscale detection of curvilinear structures in 2D and 3D image data. *IEEE ICCV*, page 864, 1995. 6, 7, 8
- [26] S. Bouattour and D. Paulus. Vessel enhancement in 2D angiographic images. *Functional Imaging and Modeling of the Heart*, 4466:41–49, 2007. 6
- [27] H. Mirzaalian and T.K. Lee and G. Hamarneh. Learning features for streak detection in dermoscopic color images using localized radial flux of principal intensity curvature. *IEEE MMBIA*, pages 97–101, 2012. 9
- [28] P. Vallotton and N. Thomas. Automated body hair counting and length measurement. *Skin Research and Technology*, 14:493–497, 2008. 9



Hengameh Mirzaalian Hengameh got her BSc and MSc degree in electrical engineering from Isfahan University of Technology, Isfahan, Iran. Currently, she is pursuing her PhD program in the Medical Image Analysis Lab at Simon Fraser University, BC, Canada. Her research interests include medical image analysis, image segmentation, and pattern recognition. In particular, she is working on computational techniques for skin lesion tracking and classification. Homepage: www.sfu.ca/~hma36



Dr. Tim K. Lee received his PhD in computer science in 2001. He is currently a senior scientist at the Cancer Control Research Program, BC Cancer Agency, and an associate professor in the Department of Dermatology and Skin Science, University of British Columbia. His research interests include processing and analyzing dermoscopic images and wide-area skin images, developing non-invasive laser speckle techniques for skin cancer detection and determining risk factors of melanoma.



Dr. Ghassan Hamarneh is a Professor of Computing Science at Simon Fraser University (SFU). He was a visiting professor at INRIA and cole Central Paris (2010–2011). Before joining SFU in 2003, he was a postdoctoral fellow at the University of Toronto and the Hospital for Sick Children (2001–2003). He completed his doctoral studies (2001) at Chalmers University of Technology and, as a predoctoral research fellow (2000–2001), at the Department of Computer Science, University of Toronto. He received his Master's degree with distinction in

Signals and Systems (digital communications) from Chalmers University in Sweden (1997). Dr. Hamarneh is a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE) since 2010 (and a member since 2003), a Senior Member of the Association for Computing Machinery, and a member of the Medical Image Computing and Computer Assisted Intervention (MICCAI) society. He is an Associate Editor of IEEE Transactions on Medical Imaging (TMI). He is the co-founder and director of the Medical Image Analysis (MIA) Lab at SFU (2003) and a founding member of the IEEE Engineering in Medicine and Biology Chapter in Vancouver. He is an author on over 200 publications in the area of medical image analysis. He publishes in and reviews for the main journals (e.g. IEEE TMI, IEEE TPAMI, IEEE TIP, and J. MedIA) and conferences in his area (e.g. IEEE ISBI, IEEE ICCV, IEEE CVPR, and MICCAI). He recently served as program committee member for SPIE Medical Imaging 2015, MICCAI 2014, and MICCAI 2014 workshops Machine Learning in Medical Imaging (MLMI) and Bayesian Models (BAMBI). His group won several honours, most recently best papers at MICCAI 2014, and MICCAI MLMI 2013. Homepage: www.cs.sfu.ca/~hamarneh