DETECTING SPECULAR HIGHLIGHTS IN DERMATOLOGICAL IMAGES

Ali Madooei, Mark S. Drew

School of Computing Science
Simon Fraser University, Vancouver, Canada
{amadooei, mark}@cs.sfu.ca
http://www.cs.sfu.ca/∼amadooei

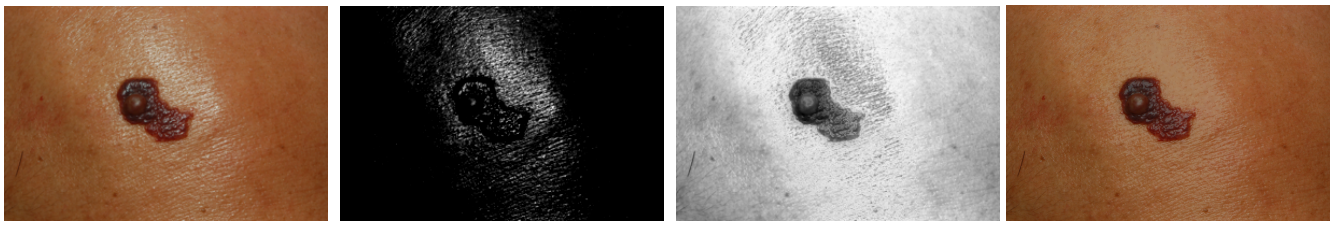

a) Input image b) Specular map c) Diffuse map d) Inpainted image

Fig. 1: a: Image of skin lesion with specular highlights. Under DRM assumptions, using NMF with sparseness constraint on specular component (see §3), the input image is decomposed to b: Specular and c: Diffuse components. d: The effect of highlights is attenuated by inpainting specular spots.

ABSTRACT

There is an increasing interest in computer-aided diagnosis of skin cancers through analysis of dermatological images. This is processing that attempts to correlate diagnosis with skin lesion appearance (i.e., the image), by extracting visual features such as colour, pigmentation, size, etc. Presence of glare (specular highlights) can confuse such systems; highlights may obscure skin surface details and appear as additional features that are not intrinsic to the lesion. In this paper, we put forward a simple method to detect specular highlights specific to dermatological images. Knowledge of the location of specularities is advantageous since it allows us to then deal with them, either by excluding them from further processing or by attempting to recover the image data in specular regions. The proposed method is built on the dichromatic reflection model and, in a novel step, uses non-negative matrix factorization with sparseness constraints to separate the specular component.

***Index Terms*—** Specular Highlight, Glare Removal, Skin Cancer, Medical Image Analysis, DRM, NMF

> There are two kinds of light–the glow that illuminates, and the glare that obscures
>
> *James Thurber*

1. INTRODUCTION

Highlights (specularities) are bright spots in an image that typically appear in regions with specular (mirror-like) reflectance such as on shiny objects. Unlike matte objects, specular objects reflect most of the light into a single direction.[1] Highlights occur on specularly reflecting objects when the geometry of the scene is such that the specular reflection is directly entering the camera. Hence, this law of behavior makes specular highlights a useful visual cue in recovering scene geometry, shape and surface material of an object.

Despite these useful properties, specular highlights also form a confounding factor for image processing and computer vision applications. Highlights generate discontinuities in images; they produce visual features disparate from those of their surroundings; and in extreme cases, they can cause loss of texture and colour information.

Moreover, many vision algorithms assume the presence of only diffuse surfaces in the scene (i.e. they assume objects are matte). An important implication of this assumption is that a simple Lambertian model [1] can often be used for image processing. Presence of specularity could potentially impair performance by violating this underlying assumption.

In dealing with specularity, various approaches exist that are motivated by different applications. Commonly, highlights are taken as outliers and disregarded in image data.

[1] For an ideal mirror, the surface normal bisects the angle between incidence and reflectance.

Alternatively, highlights are included in the image formation model in an attempt to recover the intrinsic features at specular spots. The latter approach often makes use of the foundational dichromatic reflectance model (DRM) [2] where the radiance is assumed to be a linear combination of specular and diffuse components.

This study is particularly concerned with dermatological image analysis. The reflectance of human skin can also be described by the DRM — a mixture of specular and diffuse reflectance. Skin's specular reflectance is often described as *skin glossiness* and is usually due to the existence of a thin layer of oil or sweat over the skin surface.[2] The resulting glare in skin images can, as it has been argued, disturb proper colour and texture feature extraction and thereby disrupt the performance of computer-aided diagnosis of skin lesions.

We present a simple yet effective method to detect specular highlights in skin lesion images. Building on the DRM, we use a blind source separation approach to decompose an RGB input skin image into diffuse/specular components (Fig. 1). The specular component provides a mask identifying highlight locations. Using this mask, we employ inpainting [1] to recover a highlight–attenuated image.

2. RELATED WORK

The separation of specular and diffuse reflectance is a long-standing research problem in computer graphics, computer vision and image processing. Previous work is extensive: space does not allow a review and interested reader is referred to [3]. The proposed techniques in literature differ in their assumptions and the information they use to achieve the task.

Most of these approaches have constraints that limit their usability in real applications. For dermatological images, those methods that can be applied do not produce good results (see Fig. 2), mainly because the model assumptions does not hold true.

Highlight removal in skin lesion imaging, although acknowledged as a distorting factor [4], has been mostly ignored. There are few studies that attempted to address this issue. For example, Gutenev et al. [5] considered specularities as image noise and applied a median filter to remove them. Similarly, Abbas et al. [6] applied a high-pass filter to reduce this defect. Ballerini et al. [7] characterized specular highlights as intense brightness patches with low colour saturation. They detected these areas by thresholding I and S channels of HSI colour space. The detected specular spots were excluded from further processing.

All the above-mentioned methods dealt with specularity as part of pre-processing. *To our knowledge, there is no study that reports an experimental procedure with specific results pertaining to the problem being examined here. This study is the first one to do so.*

[2]Dryness and glossiness of the skin can be affected by factors such as temperature and humidity, cleansing, sweating, etc.

3. PROPOSED METHOD

Let us begin by briefly describing the DRM and its implications to the problem under study here. The reflectance of a variety of materials, including human skin, can be described via the dichromatic reflectance model proposed by Shafer [2]. In this model, the reflected light $L(\theta, \lambda)$ from an object is described as an additive mixture of the light L_S reflected immediately at the object's surface and, the light L_D reflected from interaction with the material's body:

$$\begin{aligned} L(\theta,\lambda) &= L_D(\theta,\lambda) + L_S(\theta,\lambda) \\ &= m_D(\theta)C_D(\lambda) + m_S(\theta)C_S(\lambda) \end{aligned} \quad (1)$$

where L_D is diffuse ("body", or matte) reflection; L_S is specular ("surface", or interface) reflection; λ represents wavelength and θ encapsulates geometrical dependencies. Note that the diffuse and specular components are further decomposed into a geometry-dependent term m which indicates the magnitude of reflection and a term C depending on wavelength only. For simplicity, it is common to assume DRM components are perfect diffuse and neutral-interface specular reflectances. I.e., the body is assumed to be Lambertian and the specular reflection is assumed to have the same colour as the illumination.

Suppose the illuminant spectral power distribution is $E(x, y, \lambda)$, the spectral reflectance function is $S(x, y, \lambda)$ at each pixel (x, y), and the radiance can broadly be represented as $L(x, y, \lambda) = E(x, y, \lambda)S(x, y, \lambda_k)$.

For skin lesion images, we make the following assumptions: First, the illumination is spatially and spectrally uniform and any variation in intensity is caused by variation in skin surface curvature (i.e. the effect of shading ω). Thus $E(x, y, \lambda) = \omega(x, y)E(\lambda)$. Second, the image contains mostly skin surface, so that non-skin materials such as hair (if any) are neglectable. Third, we assume that the surface of healthy skin and the lesion can be represented by the same function $S(x, y, \lambda)$. The colour of human skin is largely due to the underlying melanin and haemoglobin pigmentations. Accordingly, skin colour variation is primarily caused by changes in the density and distribution of these two pigments. For example, blue-greyish colour in melanoma lesions indicates dense amounts of melanin in the lower layer of the skin (the dermis). We assume that changes in density and distribution of skin pigments simply scale the skin reflectance function, i.e. $S(x, y, \lambda) = \gamma(x, y)S(\lambda)$. The validity of this assumption, of course, will depend on the nature of the problem: such a simple reflectance function stands in contradistinction to a considerably more complex model used in e.g. [8]. Yet, as shown by our experiments, it is proven to be useful here.

Following these simplifying assumptions, we can rewrite the DRM for skin lesion images as follows:

$$L(x, y, \lambda) = m_D(x, y)E(\lambda)S(\lambda) + m_S(x, y)E(\lambda) \quad (2)$$

where we have lumped all geometric-dependent terms in scalar factors m_D and m_S. The implication of DRM is that RGB values of a skin lesion image live in a parallelogram spanned by the colour of the surface (diffuse colour), and the colour of the light (specular colour):

$$\begin{bmatrix} R(x,y) \\ G(x,y) \\ B(x,y) \end{bmatrix} = h_D(x,y) \begin{bmatrix} R_D \\ G_D \\ B_D \end{bmatrix} + h_S(x,y) \begin{bmatrix} R_S \\ G_S \\ B_S \end{bmatrix} \quad (3)$$

where h_D and h_S are coefficients, and colour is defined (at each point) as a weighted combination of diffuse and specular colour vectors.

The beauty of this simple linear model is in that it grants us the possibility of separating the diffuse and specular components from colour information. The decomposition task can be seen as a problem of blind source separation (BSS) where the source signals are coefficient vectors $[h_D, h_S]^T$ and, diffuse and specular colour vectors form the columns of the mixing matrix:

$$\begin{bmatrix} R(x,y) \\ G(x,y) \\ B(x,y) \end{bmatrix} = \begin{bmatrix} R_D & R_S \\ G_D & G_S \\ B_D & B_S \end{bmatrix} \times \begin{bmatrix} h_D(x,y) \\ h_S(x,y) \end{bmatrix} \quad (4)$$

Our strategy to recover the source and mixing matrices is based on non-negative matrix factorization (NMF) [9]. Given a non-negative data matrix $\boldsymbol{V}$ of size $N \times T$ (where N is the dimension of the feature space — 3 for RGB as here — and T is the number of samples, i.e. pixels here), NMF finds the non-negative matrices $\boldsymbol{W}$ and $\boldsymbol{H}$ of sizes $N \times M$ and $M \times T$ respectively, where M is the number of sources (2 for diffuse/specular here) such that $\mathbb{E}(\boldsymbol{W}, \boldsymbol{H}) = \|\boldsymbol{V} - \boldsymbol{W}\boldsymbol{H}\|^2$ is minimized. The non-negativity constraint fits well with the nature of image data in that quantities in DRM are non-negative.

As a novel, innovative step, we go on to impose another constraint on the NMF formulation. In the original DRM, the magnitude of specular reflection m_S is considered to be constant or zero for non-highlighted areas over the surface of an object. By our assumptions, we set the surface reflectance to $\gamma(x,y)S(\lambda)$ over the image plane, as if the entire image contains one object. Therefore, we can consider $h_S(x,y)$ to be constant and equal to zero for every (x,y) except on specular spots.[3] Since typically a small portion of the image contains specularity, the vector $\boldsymbol{h}_S$ of all h_S is sparse. And sparseness constraints can be adapted into NMF optimization problem [10].

Thus, we arrive at a very simple solution to detect specular highlights in skin lesion images. Given an image, we transform the image data into a 3-vector representation of R,G,B channels. The data is separated into diffuse/specular components by applying NMF with sparseness constraints on the specular component.

[3]The non-zero constant can be seen as an offset, and removed by a simple offset subtraction.

Highlight Attenuation – For many image processing algorithms, it suffices to treat specular pixels as outliers, omitting them from further processing. For some applications however, such as photo-editing and visualization, it is of interest to remove the effect of specular highlights.

To that aim, one possible strategy is to use inpainting technique (see e.g. [1]). This is a process often employed to reconstruct missing parts of images. Artists have long used manual inpainting to restore damaged paintings. Today, we use image processing to automate this laborious task. Image inpainting techniques are based on image interpolation: applying partial differential equations to propagate information from the neighbourhood around a lacuna (missing part of the image, i.e. the specular spot in our context) into the lacuna to fill it in.

4. EXPERIMENT AND RESULTS

Quantitative evaluation of specularity removal methods can only be performed on synthetic images. For real images, results are often compared to images (of the same scene) captured with polarizing filters over the camera and the light source, with consequently no highlights, in controlled laboratory settings.

Benchmark datasets exist for both synthetic and real images (with ground truth). However, our proposed method is restrictively tailored for dermatological images and therefore we cannot test it on those (more general) benchmark datasets.

Instead, to evaluate our method we conducted experiments on a set of 124 clinical images of skin lesions. These were selected amongst 550 images captured digitally during clinical routine at the British Columbia Cancer Agency [4] for patient screening and follow-up assessment purposes. The selected images were all those that exhibit (various degrees of) specular reflection.

For objective evaluation, we treat specular reflection as image noise and employ signal-to-noise ratio (SNR) to quantify it. In the context of image processing, SNR is sometimes defined as the reciprocal of the *coefficient of variation*: SNR $= \mu/\sigma$ where μ is the mean and σ is the standard deviation of pixel values in a local neighbourhood or over an entire image.

Ideally, our proposed method will reduce the level of noise (specularity) in the images, resulting in increased SNR values. This effect is indeed observed in that 120 of the 124 images (i.e. 96.77%) had SNR increased after application of our specularity detection and highlight attenuation method, with improvement range from 0.04 to 7.47dB.

Some samples are shown in Fig.2. Comparison with the highly celebrated highlight removal method of Mallick et al. [11] is also provided.

Discussion – Do we need an image processing technique to remove specular highlights from dermatological images?

[4]Located in Vancouver, Canada. Visit `http://bccancer.bc.ca/`

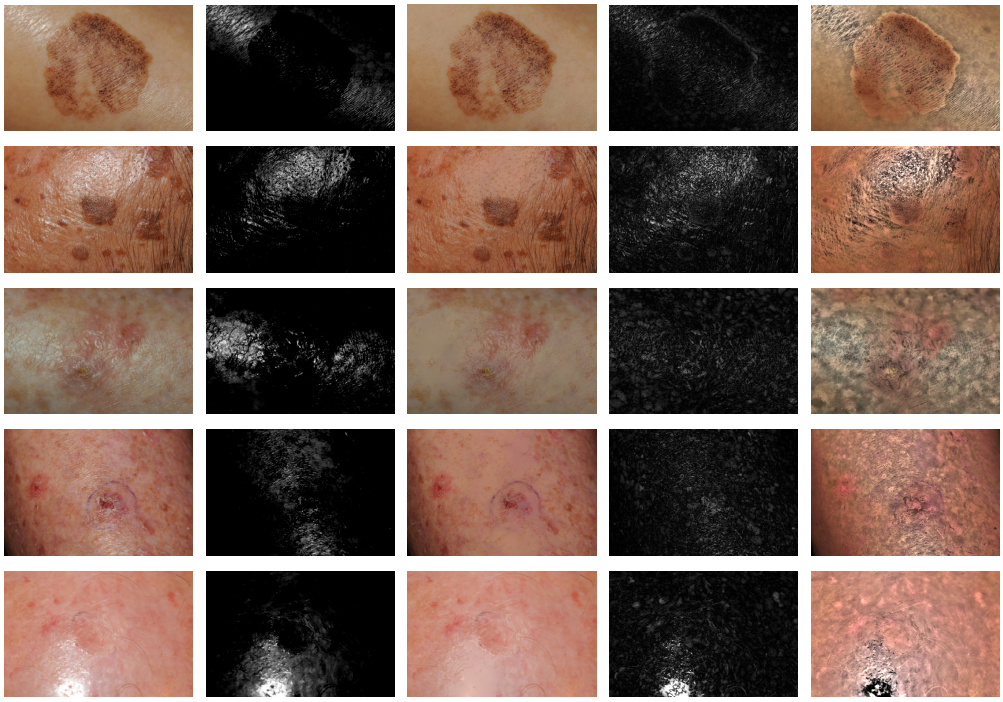

Fig. 2: The input image (1st column), specularity mask and highlight-attenuated image using our proposed method (2nd – 3rd columns), and using the method in [11] (4th - 5th columns.)

Yes, indeed. The best practice is of course to avoid imaging artifacts by implementing best imaging setup and using proper equipment such as polarization filters, etc. However, most clinicians continue to capture images under uncontrolled conditions. Moreover, most of existing data available to researchers is not taken under ideal conditions. Therefore, an image processing technique such as the proposed one may be found useful in dealing with such shortfalls.

5. CONCLUSION

The performance and diagnostic utility of computer-aided skin lesion analysis is hindered by the existence of specularities. We have presented a novel, simple and effective method to detect specular pixels, and to recover a highlight-attenuated image. The proposed method relies on reasonable assumptions and constraints derived as an underlying model, and from heuristics justified by good empirical performance.

Future work will include a quantitative evaluation of the merit and potential utility of employing the proposed technique as a pre-processing step for automated skin lesion classification.

6. REFERENCES

[1] Richard Szeliski, *Computer Vision: Algorithms and Applications*, Springer, 2011 edition, 2010.

[2] S. Shafer, "Using color to separate reflection components," *Color Research & App.*, vol. 10, no. 4, pp. 210–218, 1985.

[3] A. Artusi, F. Banterle, and D. Chetverikov, "A survey of specularity removal methods," *Computer Graphics Forum*, vol. 30, no. 8, pp. 2208–2230, 2011.

[4] E.J. Bae, S.H. Seo, Y.C. Kye, and H.H. Ahn, "A quantitative assessment of the human skin surface using polarized light digital photography and its dermatologic significance," *Skin Res. & Tech.*, vol. 16, no. 3, pp. 270–274, 2010.

[5] A. Gutenev, V.N. Skladnev, and D. Varvel, "Acquisition-time image quality control in digital dermatoscopy of skin lesions," *Computerized Medical Imaging and Graphics*, vol. 25, no. 6, pp. 495–499, 2001.

[6] Q. Abbas, M.E. Celebi, Fondn G.I., and M. Rashid, "Lesion border detection in dermoscopy images using dynamic programming," *Skin Res. & Tech.*, vol. 17, no. 1, pp. 91–100, 2011.

[7] L. Ballerini, R. Fisher, B. Aldridge, and J. Rees, "A color and texture based hierarchical k-nn approach to the classification of non-melanoma skin lesions," in *Color Medical Image Analysis*, pp. 63–86. Springer, 2013.

[8] A. Madooei and M. Drew, "A probabilistic approach to quantification of melanin and hemoglobin content in dermoscopy images," in *MICCAI*, pp. 49–56. Springer, 2014.

[9] P. Paatero and U. Tapper, "Positive matrix factorization," *Environmetrics*, vol. 5, no. 2, pp. 111–126, 1994.

[10] P. Hoyer, "Non-negative matrix factorization with sparseness constraints," *J. Mach. Learning Res.*, vol. 5, pp. 1457–1469, 2004.

[11] S. Mallick, T. Zickler, P. Belhumeur, and D. Kriegman, "Specularity removal in images and videos: A PDE approach," in *ECCV*, 2006, pp. 550–563.