

# Illumination-based Transformations Improve Skin Lesion Segmentation in Dermoscopic Images

Kumar Abhishek\*, Ghassan Hamarneh, and Mark S. Drew

School of Computing Science, Simon Fraser University, Canada  
{kabhishe, hamarneh}@sfu.ca, mark@cs.sfu.ca

## Abstract

*The semantic segmentation of skin lesions is an important and common initial task in the computer aided diagnosis of dermoscopic images. Although deep learning-based approaches have considerably improved the segmentation accuracy, there is still room for improvement by addressing the major challenges, such as variations in lesion shape, size, color and varying levels of contrast. In this work, we propose the first deep semantic segmentation framework for dermoscopic images which incorporates, along with the original RGB images, information extracted using the physics of skin illumination and imaging. In particular, we incorporate information from specific color bands, illumination invariant grayscale images, and shading-attenuated images. We evaluate our method on three datasets: the ISBI ISIC 2017 Skin Lesion Segmentation Challenge dataset, the DermoFit Image Library, and the PH2 dataset and observe improvements of 12.02%, 4.30%, and 8.86% respectively in the mean Jaccard index over a baseline model trained only with RGB images.*

## 1. Introduction

Skin conditions are the most common reason for visits to general practitioners in studied populations [29], and the prevalence of skin cancer in the United States has been higher than all other cancers combined over the last three decades [27]. Melanoma, a type of skin cancer which represents only a small fraction of all skin cancers in the USA, is responsible for over 75% of skin cancer related fatalities and a projected 6,850 deaths in 2020 in the USA alone [31]. However, studies have shown that early diagnosis can drastically improve patient survival rates. While skin cancers can be diagnosed by visual examination, it is often difficult to distinguish malignant lesions from healthy skin. As a

result, computer aided diagnoses of skin lesions have been widely used to automate the assessment of dermoscopic images. The segmentation of skin lesion images is therefore a crucial step in the diagnosis and subsequent treatment. Segmentation refers to the process of delineating the lesion boundary by assigning pixel-wise labels to the dermoscopic images, so as to separate the lesion from the surrounding healthy skin. However, this is a complicated task, primarily because of the large variety in the shape, color, presentation, and contrast of skin lesions, originating from intra- and inter-class variations as well as image acquisition.

Recent years have witnessed the successful applications of machine learning, particularly deep learning-based approaches, to the semantic segmentation of skin lesions. Numerous contributions have been made in terms of new architectures (such as fully convolutional network models [41], deep residual networks [39], deep auto-context architectures [24], etc.), shape [23] and texture [42] priors, synthesis-based augmentations [25, 3], and loss functions [38, 4, 36].

Although deep learning-based approaches have made significant improvements to the segmentation performance, they are reliant on a large amount of training data in order to yield acceptable performances. Deep learning-based approaches also tend to ignore knowledge about illumination in skin lesion images and other such physics-based properties, an area that has been explored in the past. Madooei et al. [20] proposed a new 2D log-chromaticity color space and showed that color intensity triplets in skin images lie on a plane, and used Otsu’s algorithm to segment skin lesions, demonstrating superior performance even on low-contrast lesions. In another work [19], they also presented pre-processing techniques for improved segmentation of skin lesions. They calculated an illumination invariant grayscale ‘intrinsic’ image and used it to attenuate shading and lighting intensity changes in dermoscopic images. Moreover, they also presented a novel RGB-to-grayscale conversion

\*Corresponding author

algorithm for dermoscopic images using principal component analysis in the optical density space. In a more recent work, Guarracino et al. [13] proposed an unsupervised approach for skin lesion segmentation from dermoscopic images by choosing certain color bands. Ng et al. [16] demonstrated an improvement in segmentation performance with the use of color constancy algorithms in a fully convolutional network-based segmentation model. For more detailed literature surveys, we point the interested readers to review papers on skin lesion segmentation [8] and incorporating color information in computer-aided skin lesion diagnosis [18]. However, very little research has been done on the applicability of color image theory and illumination information to a deep learning-based semantic segmentation framework.

We propose a novel deep semantic segmentation algorithm for dermoscopic images, which leverages prior illumination and color theory knowledge. In particular, we build upon previous works and leverage specific color bands, intrinsic images, and skin image information to yield improved segmentation results. To the best of our knowledge, this is the first work that incorporates such information in a deep learning-based framework.

The rest of the paper is structured as follows: Section 2 describes the proposed approach, Section 3 describes the dataset, the experiments, and the evaluation metrics, Section 4 contains quantitative and qualitative analyses of the proposed approach, and Section 5 concludes the paper.

## 2. Method

In this work, we extract color information from the RGB dermoscopic images and use them along with the original image to train a deep semantic segmentation model. In particular, we use (a) variations of certain color bands (Section 2.1), (b) a color-theory-based grayscale estimate (Section 2.3), (c) an illumination-invariant intrinsic grayscale image (Section 2.2), and (d) a shading-attenuated image obtained from the dermoscopic image (Section 2.4). Figure 1 shows an overview of the proposed approach. In the subsequent sections, we describe the methods for obtaining these images.

### 2.1. Choosing color bands

We choose the two color bands which have been shown to be efficient for skin lesion segmentation [13]: the red color channel from the normalized RGB image, and the complement of the value channel from the HSV color space representation of the image (denoted by  $R'$  and  $V^*$  respectively) and concatenate them to the original RGB dermoscopic image. They are defined as:

$$R' = \frac{R}{R + G + B}, \text{ and} \quad (1)$$

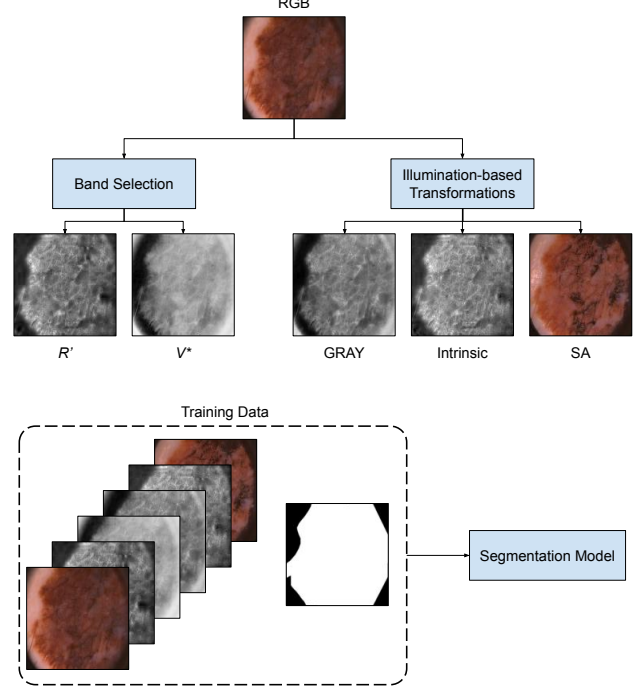


Figure 1: An overview of the proposed approach. Various color bands and transformations are computed as explained in Section 2 and concatenated channel-wise to the original RGB dermoscopic image in order to train the segmentation model.

$$V^* = 1 - V, \quad (2)$$

where  $R, G, B$  denote the channels from the original image and  $V$  denotes the Value channel from the HSV representation of the image. For computational efficiency, instead of converting the image from the RGB to the HSV color space, the  $V$  channel can directly be calculated as:

$$V = \max\left(\frac{R}{M}, \frac{G}{M}, \frac{B}{M}\right), \quad (3)$$

where  $M = 2^n - 1$  denotes the number of gray-levels in an  $n$ -bit image ( $M = 255$  for our 8-bit color images).

### 2.2. Intrinsic images

We follow the approach proposed by Finlayson et al. [11] to derive an illumination-invariant grayscale ‘intrinsic’ image using entropy minimization. Given an RGB image, let  $R_k, k \in \{1, 2, 3\}$  denote the channel-wise intensities, and the 3-vector chromaticities can be obtained by dividing each color channel by the geometric mean of the channels.

$$c_k = \frac{R_k}{R_p} \quad \forall k \in \{1, 2, 3\}, \quad (4)$$

where  $R_p = \sqrt[3]{\prod_{k=1}^3 R_k}$ . Finlayson et al. note that while it is possible to obtain 2-vector chromaticities by dividing by one of the color channels, the choice of dividing by the geometric mean ensures that there is no bias towards any particular channel.

From [11], assuming the light as a Planckian radiator (and using Wein's approximation [37]) and the camera sensors to be fairly narrow-band, the channel  $R_k$  can be written as:

$$R_k = \sigma L k_1 \lambda_k^{-5} \exp\left(-\frac{k_2}{T \lambda_k}\right) S(\lambda_k) q_k, \quad (5)$$

where  $\sigma$  is the Lambertian shading,  $L$  is the overall light intensity,  $T$  is the temperature of the lighting color,  $S(\lambda_k)$  is the spectral reflectance of the surface (which is the skin in our case) as a function of the wavelength  $\lambda_k$ ,  $q_k$ s are the camera sensor sensitivity functions, and  $k_1, k_2$  are constants. Therefore, the log-chromaticities (obtained by taking the logarithm of Eqn. 4) can be written as:

$$\rho_k \equiv \log(c_k) = \log\left(\frac{s_k}{s_p}\right) + \left(\frac{e_k - e_p}{T}\right), \quad (6)$$

where  $s_k = k_1 \lambda_k^{-5} S(\lambda_k) q_k$  and  $e_k = -k_2 / \lambda_k$ . Note that this expression does not have the shading and the intensity information. Eqn. 6 is the parametric equation of a straight line with  $T$  as the parameter, and although the surface information is present in the intercept of the line, the direction is given by  $e \equiv (e_k - e_p)$ , which is independent of the surface.

With 2D log-chromaticities, it is possible to obtain the intrinsic image by projecting  $\rho_k$  in a direction orthogonal to  $e$  (denoted by  $e^\perp$ ), followed by taking its exponential. However, dividing by the geometric mean (Eqn. 4) yields 3D log-chromaticities, and therefore, the task is to find a projector  $P$  which can project  $\rho$  onto the 2D chromaticity space, which is a plane. Note that the log-chromaticities  $\rho$  are orthogonal to  $u \equiv (1, 1, 1)^T / \sqrt{3}$ , and so the projector  $P_u^\perp$  can be used to characterize the plane. Since  $P_u^\perp$  has two non-zero eigenvalues, it can be decomposed as:

$$P_u^\perp = I - uu^T = U^T U, \quad (7)$$

where  $I$  is the identity matrix and  $U$  is a  $2 \times 3$  orthogonal matrix, which projects the three  $\rho$  vectors onto a coordinate system in the plane as two vectors denoted by  $\chi$ . It should be noted that straight lines in  $\rho$  remain straight in  $\chi$ .

$$\chi \equiv U \rho \quad (8)$$

The next step is to find the optimal angle  $\theta_{\text{proj}}$  to project along in the  $\{\chi_1, \chi_2\}$  plane, for which the entropy for the marginal distribution along a 1D line orthogonal to the lighting direction is minimized. The resulting projected log grayscale image  $\mathcal{I}_{\text{projected}}$  is given by

$$\mathcal{I}_{\text{projected}} = \chi_1 \cos\theta_{\text{proj}} + \chi_2 \sin\theta_{\text{proj}} \quad (9)$$

To compute the best projection angle, only the middle 90% of the data is used. This is done to exclude the outliers by using data between the 5<sup>th</sup> and the 95<sup>th</sup> percentiles. Then, Scott's rule [30] is used to estimate the bin width for constructing the histogram as:

$$\text{bin width} = 3.5 * \text{STD (projected data)} * \sqrt[3]{N}, \quad (10)$$

where  $\text{STD}(\cdot)$  denotes the standard deviation and  $N$  is the size of the grayscale image data for a given angle  $\omega$ . Next, for each angle, probabilities  $p_i$  for each bin  $i$  are computed by dividing the bin by the sum of the bin counts, and the entropy is calculated as:

$$\eta = - \sum_i p_i(I) \log(p_i(I)) \quad (11)$$

The angle which yields the lowest entropy is chosen as the projection angle, and finally the projected log-image  $\mathcal{I}_{\text{projected}}$  is exponentiated to yield the intrinsic image. The entire approach is shown in Algorithm 1.

---

**Algorithm 1:** Intrinsic image by entropy minimization

---

**Input:** RGB image  $\mathcal{I}$

**Output:** Grayscale intrinsic image  $I_{\text{Intrinsic}}$

---

construct 2D log-chromaticity representation of

$\mathcal{I}_{\text{projected}}$ ;

**for**  $\omega \leftarrow 1$  **to** 180 **do**

    calculate histogram bin width;

    compute histogram with middle 90% data;

    compute  $\eta_\omega$ , the entropy for the angle  $\omega$ ;

**end**

$\theta_{\text{proj}} \leftarrow \arg \min_\omega \eta_\omega$ ;

$\mathcal{I}_{\text{projected}} \leftarrow \chi_1 \cos\theta_{\text{proj}} + \chi_2 \sin\theta_{\text{proj}}$ ;

$I_{\text{Intrinsic}} \leftarrow \exp(\mathcal{I}_{\text{projected}})$

**return**  $I_{\text{Intrinsic}}$ ;

---

### 2.3. Grayscale images of skin lesions

Madooei et al. [19] proposed a RGB-to-grayscale conversion algorithm for dermoscopic images based on the optics and the reflectance properties of skin surfaces. Unlike a traditional grayscale representation calculated as the weighted sum of the red, the green, and the blue channels [26], this grayscale image preserves the lesion while suppressing the healthy skin, thereby increasing the contrast between the healthy and the affected regions. Based on the skin models proposed by Hiraoka et al. [15] and Tsumura et al. [34], the spectral reflection of skin at a pixel  $(x, y)$  under

polarized light can be written as:

$$S(x, y, \lambda) = \exp \left[ -\rho_m(x, y)\alpha_m(\lambda)l_m(\lambda) - \rho_h(x, y)\alpha_h(\lambda)l_h(\lambda) \right], \quad (12)$$

where  $\rho_{\{m,h\}}$ ,  $\alpha_{\{m,h\}}$ , and  $l_{\{m,h\}}$  denote the densities ( $\text{cm}^{-3}$ ), cross-sectional areas ( $\text{cm}^2$ ) for scattering absorption, and mean path lengths for photons in the epidermis and dermis layers of the human skin for melanin and hemoglobin (denoted by subscripts  $m$  and  $h$  respectively). Substituting this expression in Eqn. 5 followed by taking the logarithms on both sides yields:

$$\begin{aligned} \log(R_k(x, y)) &= -\rho_m(x, y)\sigma_m - \rho_h(x, y)\sigma_h \\ &+ \log(k_1 I(x, y)\sigma(x, y)) \\ &+ \log(\lambda_k^{-5}) - \frac{k_2}{\lambda_k T}, \end{aligned} \quad (13)$$

where  $\sigma_{\{m,h\}} = \alpha_{\{m,h\}}l_{\{m,h\}}$ . Eqn. 13 suggests that the pixels from a skin image lie on a plane in the optical density space described by  $[-\log R_1, -\log R_2, -\log R_3]$ . As such, Madooei et al. observe that in almost all the skin lesion images analyzed, the first eigenvector in the principal component analysis (PCA) explains a very high fraction of the total variance, and thus contains most of the information in the image. As such, the first principal component can be used to obtain a grayscale skin lesion image. The approach has been described in Algorithm 2.

---

**Algorithm 2:** Grayscale skin lesion image

---

**Input:** RGB image  $\mathcal{I}_{\text{RGB}}$

**Output:** Grayscale image  $\mathcal{I}_{\text{GRAY}}$

$\mathcal{I}'_{\text{RGB}} \leftarrow \text{channel-wise-flatten}(\mathcal{I}_{\text{RGB}});$   
 $\mathcal{I}^*_{\text{RGB}} \leftarrow \mathcal{I}'_{\text{RGB}}$  image in optical density space;  
 $P \leftarrow$  1st principal component from PCA( $\mathcal{I}^*_{\text{RGB}}$ );  
 $\mathcal{I}_{\text{GRAY}} \leftarrow \text{reshape}(P)$   
**return**  $\mathcal{I}_{\text{GRAY}};$

---

## 2.4. Shading-attenuated skin lesion images

The non-flat nature of skin surfaces, especially lesions, and the effect of light intensity falloffs towards the edges of the skin lesions can induce shading in dermoscopic images, which can degrade the segmentation (and classification) performance. Madooei et al. [19] proposed to use the intrinsic images generated by Finlayson et al. [11] to perform illumination normalization in dermoscopic images, thereby performing shading-attenuation. They proposed to use the intrinsic image to normalize the intensity values. Given a dermoscopic image, its intrinsic image is first calculated. The RGB image is then converted to the HSV color

space. In order to normalize the intensities, the Value (V) channel of the HSV image is used. First, both the intrinsic image and the Value channel are normalized. The image intensity histogram of the Value channel is then mapped to that of the intrinsic image. Finally, this new normalized and histogram-mapped Value channel is used to replace the original Value channel in the HSV image, and the resultant image is then mapped back to the RGB color space. The authors demonstrated a significant attenuation in the shading and the intensity falloff using this approach. The entire approach is summarized in Algorithm 3.

---

**Algorithm 3:** Shading-attenuated skin lesion image

---

**Input:** RGB image  $\mathcal{I}_{\text{RGB}}$

**Output:** Shading-attenuated RGB image  $\mathcal{I}_{\text{SA}}$

compute  $I_{\text{Intrinsic}}$  from  $\mathcal{I}_{\text{RGB}}$  using Algorithm 1;  
 $\mathcal{I}_{\text{HSV}} \leftarrow \text{rgb2hsv}(\mathcal{I}_{\text{RGB}});$   
normalize  $I_{\text{Intrinsic}}, V;$   
 $V^* \leftarrow V$  histogram-matched to  $I_{\text{Intrinsic}};$   
 $\mathcal{I}_{\text{SA}} \leftarrow \text{hsv2rgb}(\mathcal{I}_{\text{HSV}}^*)$   
**return**  $\mathcal{I}_{\text{SA}};$

---

## 3. Datasets and Experimental Details

### 3.1. Datasets

We evaluate our proposed approach on three skin lesion image datasets, namely the ISIC ISBI 2017 dataset, the DermoFit dataset, and the PH2 dataset.

#### 3.1.1 ISBI ISIC 2017

The ISIC ISBI 2017: Skin Lesion Analysis Towards Melanoma Detection Challenge [10] Part 1 dataset contains skin lesion images and the corresponding manually annotated lesion delineations belonging to three diagnoses: melanoma, seborrheic keratosis, and benign nevi. The dataset is split into training, validation, and testing subsets containing 2000, 150, and 600 images respectively.

#### 3.1.2 DermoFit

The DermoFit Image Library [1, 5] contains 1300 skin lesion images belonging to ten diagnosis classes, along with the corresponding binary segmentation masks. We divide the original dataset into training, validation, and testing splits in the ratio of 60 : 10 : 30.

#### 3.1.3 PH2

The PH2 Image Database [21] contains a total of 200 dermoscopic images of common nevi, atypical nevi, and






Figure 2: Transformation results for 5 images from the ISIC 2017 training set.

melanomas, along with their lesion segmentations annotated by an expert dermatologist.

### 3.2. Experiments and Evaluation

Since the goal of this work is to demonstrate the effectiveness of the various color theory and illumination-based transformations for enhancing segmentation performance, we use U-Net [28] as the baseline segmentation network. The U-Net consists of a symmetric encoder-decoder architecture, with skip connections between symmetrically corresponding layers in the encoder and the decoder, which help in recovering the full spatial resolution [17] and address the problem of gradient vanishing [32]. For evaluating upon the ISIC 2017 dataset, we train seven segmentation models where the inputs to the corresponding networks are the following:

- **RGB Only:** The original 3-channel RGB dermoscopic image.
- **All Channels:** The original RGB dermoscopic image channel-wise concatenated with  $R'$ ,  $V^*$  (Section 2.1), intrinsic image (denoted by Intrinsic; Section 2.2), grayscale image (denoted by GRAY; Section 2.3), and shading-attenuated image (denoted by GRAY; Section 2.4). The result is a 10-channel image.
- Next, to determine the contribution of each of the transformations described in Section 2, we drop one component at a time from the 10-channel image above, and denote it by **No  $\mathbf{x}$** , where  $\mathbf{x}$  denotes the dropped channel. The models are:

- **No  $R'$ :** 9-channel image.
- **No  $V^*$ :** 9-channel image.
- **No GRAY:** 9-channel image.
- **No Intrinsic:** 9-channel image.
- **No SA:** 7-channel image.

For each of these aforementioned models, the input layer of the segmentation network is modified to handle the corresponding number of channels, and the rest of the architecture remains the same. The models are trained to predict the pixelwise labels for the semantic segmentation task. All images and their corresponding ground truth segmentation masks are resized to  $128 \times 128$  using nearest neighbor interpolation from Python’s SciPy [35] library. All networks are trained with Dice loss [22] using mini-batch stochastic gradient descent with a batch size of 40 (since a larger batch size exceeded our GPU memory) and a learning rate of  $1e - 3$ . We apply real time data augmentation strategies (random horizontal and vertical flips and rotations) during training. All the code was written in Python and the PyTorch framework was used to implement the deep segmentation models.

For the evaluation of the three methods, we report the metrics used by the official challenge - pixel-wise accuracy, sensitivity, specificity, Dice similarity coefficient, and Jaccard index (also known as the intersection over union). They are given by:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (14)$$

Table 1: Quantitative results for the seven methods on 600 images from the ISIC 2017 test set. (mean  $\pm$  standard error).

Method	Accuracy	Dice Coefficient	Jaccard Index	Sensitivity	Specificity
RGB Only	0.9029 $\pm$ 0.0053	0.7781 $\pm$ 0.0086	0.6758 $\pm$ 0.0095	0.7471 $\pm$ 0.0091	0.9683 $\pm$ 0.0031
All Channels	<b>0.9220 <math>\pm</math> 0.0045</b>	<b>0.8386 <math>\pm</math> 0.0078</b>	<b>0.7570 <math>\pm</math> 0.0089</b>	<b>0.8706 <math>\pm</math> 0.0077</b>	0.9516 $\pm$ 0.0037
No $R'$	0.9185 $\pm$ 0.0046	0.8243 $\pm$ 0.0078	0.7363 $\pm$ 0.0090	0.7949 $\pm$ 0.0085	0.9735 $\pm$ 0.0030
No $V^*$	0.9189 $\pm$ 0.0049	0.8263 $\pm$ 0.0077	0.7381 $\pm$ 0.0089	0.7892 $\pm$ 0.0087	0.9786 $\pm$ 0.0025
No Intrinsic	0.9092 $\pm$ 0.0056	0.7997 $\pm$ 0.0094	0.7139 $\pm$ 0.0103	0.7662 $\pm$ 0.0104	<b>0.9803 <math>\pm</math> 0.0024</b>
No GRAY	0.9116 $\pm$ 0.0052	0.8163 $\pm$ 0.0080	0.7260 $\pm$ 0.0091	0.8041 $\pm$ 0.0090	0.9643 $\pm$ 0.0033
No SA	0.9198 $\pm$ 0.0050	0.8274 $\pm$ 0.0083	0.7445 $\pm$ 0.0093	0.8137 $\pm$ 0.0088	0.9603 $\pm$ 0.0044

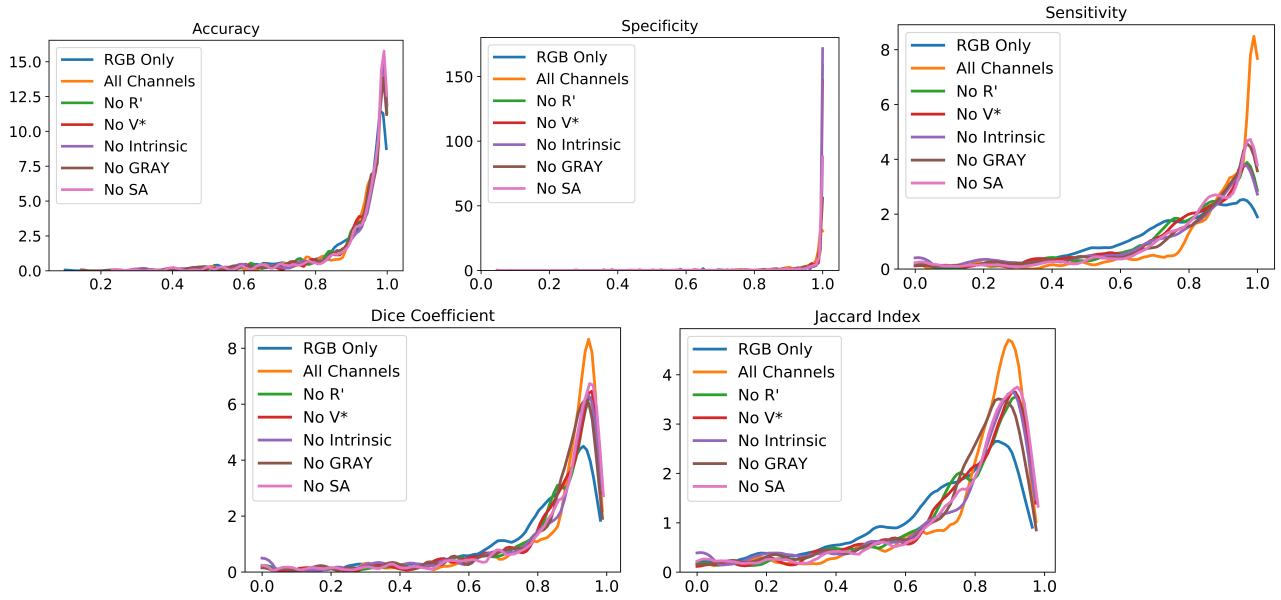


Figure 3: Kernel density estimates for the five metrics for all the segmentation methods evaluated on the ISIC 2017 test set.

$$\text{Sensitivity} = \frac{TP}{TP + FN}, \quad (15)$$

$$\text{Specificity} = \frac{TN}{TN + FP}, \quad (16)$$

$$\text{Dice coefficient, } Dice(A, B) = 2 \frac{|A \cap B|}{|A| + |B|}, \quad (17)$$

$$\text{Jaccard index, } J(A, B) = \frac{|A \cap B|}{|A \cup B|}, \quad (18)$$

where  $TP$ ,  $TN$ ,  $FP$ ,  $FN$  denote true positive, true negative, false positive, and false negative respectively, and  $A$ ,  $B$  denote two binary masks. As with the challenge, all metrics are reported at 128 confidence threshold.

## 4. Results and Discussion

Figure 2 shows the normalized red channel ( $R'$ ), the complement of the Value channel ( $V^*$ ), the intrinsic image (Intrinsic) using the approach by Finlayson et al. [11], the grayscale converted image (GRAY) and the shading-attenuated image (SA) using the approach by Madooei et al. [19] for 5 dermoscopic images (with their ISIC image IDs) and their corresponding ground truth segmentation masks from the ISIC 2017 training set. We notice that the presence of artifacts such as markers (second row) and rulers (third and fourth rows) lead to poor results, particularly for shading-attenuated images. While the shading-attenuation results are acceptable for some images, a large number of images yield poor results, such as the last row in Figure 2.




Figure 4: Qualitative results for the segmentation performance of all the methods on 6 images from the ISIC 2017 test set. Incorporating the two color bands and illumination-based transformations improves the segmentation consistently, and the performance drop is the most significant when Intrinsic is not used.

Table 1 shows the quantitative results for the 600 test images evaluated using the seven trained segmentation networks. We observe that ‘All Channels’ outperforms ‘RGB Only’ in all metrics except specificity, where the difference is quite small. Using all the transformations yields an improvement of 12.02% and 7.76% over the baseline (‘RGB Only’) in the mean Jaccard index and the mean Dice similarity coefficient metrics respectively. We also note that we are within 1% of the Jaccard index of the top 3 entries on the challenge leaderboard [2], without using any additional external data [7], post-processing, or an ensemble of models [40, 6] and without optimizing the network architecture or any other hyperparameters.

To further capture the improvement in the segmentation performance, we plot the kernel density estimates of the metrics for all the methods (Figure 3). We use the Epanechnikov kernel to estimate their probability density functions, and the plots have been clipped to the range of the values of the respective metrics. The plots show higher peaks (indicating higher densities) at larger values of all the metrics for the proposed method(s).

Figure 4 shows six samples from the test dataset and their corresponding ground truth segmentation masks, along with the prediction outputs from the seven models. The samples have been chosen to cover almost all possible variations in the images, such as the size of the lesion, the contrast between the lesion and the surrounding skin, and the presence of artifacts (ruler, gel bubble, etc.). We note that apart from

the improved segmentation performance, incorporating the proposed transformations into the input to the model also considerably improves the false positive and the false negative labels.

Next, we analyze the contribution of each of the color theory and illumination-based transformations towards improving the segmentation performance. From Table 1, we can see that dropping the normalized red channel ( $R'$ ), the complement of the Value channel ( $V^*$ ), and the shading-attenuated image (SA) have the least impact on the Dice coefficients. Of these, the first two can possibly be explained by the fact that these are relatively simpler transformations as compared to the other three, and are therefore easier for the network to learn. As for the SA component, as already noted previously and shown in the SA column in Figure 2, a large number of images yield very poor results. Since we use JPEG compressed images, most of the high frequencies (in the Fourier domain representation) are discarded during JPEG compression, which leads to the entropy minimization step producing sub-optimal projection angles. We confirm this by plotting the projection angles calculated for the 2000 and 780 images in the ISIC 2017 and the DermoFit training sets (Figure 5). We observe that the projection angles are spread across the entire range, which is in contrast to Finlayson et al. [11] where the minimum entropy angles are between  $147^\circ$  and  $161^\circ$  for their HP912 camera. As such, we do not expect the SA images to provide a considerable improvement when used in a segmentation model,

Table 2: Quantitative results for the two methods on the 390 test images from the DermoFit dataset and 200 images from the PH2 dataset (mean  $\pm$  standard error).

Dataset	Method	Accuracy	Dice Coefficient	Jaccard Index	Sensitivity	Specificity
DermoFit	RGB Only	$0.9024 \pm 0.0038$	$0.8437 \pm 0.0053$	$0.7418 \pm 0.0069$	$0.8080 \pm 0.0078$	<b><math>0.9534 \pm 0.0030</math></b>
	No SA	<b><math>0.9124 \pm 0.0030</math></b>	<b><math>0.8674 \pm 0.0042</math></b>	<b><math>0.7737 \pm 0.0055</math></b>	<b><math>0.8721 \pm 0.0053</math></b>	$0.9347 \pm 0.0032$
PH2	RGB Only	$0.8546 \pm 0.0133$	$0.7989 \pm 0.0128$	$0.6944 \pm 0.0140$	$0.8032 \pm 0.0166$	$0.9543 \pm 0.0031$
	No SA	<b><math>0.8926 \pm 0.0091</math></b>	<b><math>0.8537 \pm 0.0071</math></b>	<b><math>0.7559 \pm 0.0091</math></b>	<b><math>0.8442 \pm 0.0110</math></b>	<b><math>0.9607 \pm 0.0022</math></b>

which is consistent with the quantitative results on the ISIC 2017 test set.

On the other hand, we observe that the intrinsic image (Intrinsic) and the grayscale converted image (GRAY) are crucial to the segmentation performance improvement. Since these transformations rely on the log-chromaticity and the optical density space representations respectively, and therefore are not so easily learned by a deep semantic segmentation model. The dip in performance is the most when the Intrinsic image is dropped, indicating that it is the most important illumination-based transformation for improving the segmentation. Figure 4 shows that ‘No Intrinsic’ also results in higher false positives and false negatives (most clearly visible in the second and the third rows).

Finally, for the DermoFit dataset, we train two models: ‘RGB Only’ and ‘No SA’. As discussed, SA images do not contribute much to improving the segmentation performance (as shown for the ISIC 2017 dataset, Table 1), while also being computationally intensive (Algorithm 3). As such, we use ‘RGB Only’ as the baseline to evaluate the performance of ‘No SA’. As for the PH2 dataset, given the small number of images, we use the entire PH2 dataset as a test set for the two models trained on the DermoFit dataset to evaluate the generalizability of the trained models.

Table 2 shows the quantitative results for evaluating these two trained models on the DermoFit test set and the entire PH2 dataset. We observe that ‘No SA’ improves the mean Jaccard index for the DermoFit and the PH2 datasets by 4.30% and 8.86% respectively over the ‘RGB Only’ baseline.

## 5. Conclusion

Motivated by the potential value of leveraging information about the physics of skin illumination and imaging in a data hungry deep learning setting, in this work, we proposed a novel semantic segmentation framework for skin lesion images by augmenting the RGB dermoscopic images with additional color bands and intrinsic, grayscale, and shading-attenuated images. We demonstrated the efficacy of the proposed approach by evaluating on three datasets: the ISIC ISBI 2017 Challenge dataset, the DermoFit Image




Figure 5: Histogram of projection angles for the training images from the ISIC 2017 and the DermoFit datasets. The projection angles for these images are spread across the entire range, whereas it is restricted to a small range for Finlayson et al. [11].

Library, and the PH2 database and observed a considerable performance improvement over the baseline method. We also performed ablation studies to ascertain the contribution of each of the transformations on the segmentation performance improvement. We hypothesize that, despite being useful for improving prediction accuracy, deep learning does not happen to stumble upon these illumination-based channels given the large search space, the fixed architecture, and the local gradient-descent optimizer. Future work could explore augmenting with other channels [14, 12, 9, 33] and, perhaps more interestingly, investigate architectures, losses, or training strategies to ensure such channels are automatically learnt.

## Acknowledgements

Partial funding for this project is provided by the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors are grateful to the NVIDIA Corporation for donating Titan X GPUs used in this research.



## References

- [1] DermoFit Image Library. <https://licensing.edinburgh-innovations.ed.ac.uk/i/software/dermofit-image-library.html> [Accessed: March 4, 2020].
- [2] ISIC 2017: Skin lesion analysis towards melanoma detection part 1: Lesion segmentation phase 3: Final test submission leaderboard. <https://challenge.kitware.com/#phase/584b0afacad3a51cc66c8e24> [Accessed: November 24, 2019].
- [3] Kumar Abhishek and Ghassan Hamarneh. Mask2Lesion: Mask-constrained adversarial skin lesion image synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention Workshop on Simulation and Synthesis in Medical Imaging (MICCAI SASHIMI)*, pages 71–80. Springer International Publishing, 2019.
- [4] Nabila Abraham and Naimul Mefraz Khan. A novel focal Tversky loss function with improved attention U-Net for lesion segmentation. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 683–687. IEEE, 2019.
- [5] Lucia Ballerini, Robert B Fisher, Ben Aldridge, and Jonathan Rees. A color and texture based hierarchical K-NN approach to the classification of non-melanoma skin lesions. In *Color Medical Image Analysis*, pages 63–86. Springer, 2013.
- [6] Matt Berseth. ISIC 2017-skin lesion analysis towards melanoma detection. *arXiv preprint arXiv:1703.00523*, 2017.
- [7] Lei Bi, Jinman Kim, Euijoon Ahn, and Dagan Feng. Automatic skin lesion analysis using large-scale dermoscopy images and deep residual networks. *arXiv preprint arXiv:1703.04197*, 2017.
- [8] M Celebi, Quan Wen, Hitoshi Iyatomi, Kouhei Shimizu, Huiyu Zhou, and Gerald Schaefer. A state-of-the-art survey on lesion border detection in dermoscopy images. In *Dermoscopy Image Analysis*, pages 97–129. CRC Press, Sept. 2015.
- [9] M. Emre Celebi, Hitoshi Iyatomi, and Gerald Schaefer. Contrast enhancement in dermoscopy images by maximizing a histogram bimodality measure. In *2009 16th IEEE International Conference on Image Processing (ICIP)*. IEEE, Nov. 2009.
- [10] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. Skin lesion analysis toward melanoma detection: A challenge at the 2017 International Symposium on Biomedical Imaging (ISBI), hosted by the International Skin Imaging Collaboration (ISIC). In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 168–172. IEEE, 2018.
- [11] Graham D Finlayson, Mark S Drew, and Cheng Lu. Intrinsic images by entropy minimization. In *European Conference on Computer Vision*, pages 582–595. Springer, 2004.
- [12] David Delgado Gomez, Constantine Butakoff, Bjarne Kjær Ersboll, and William Stoecker. Independent histogram pursuit for segmentation of skin lesions. *IEEE Transactions on Biomedical Engineering*, 55(1):157–161, Jan. 2008.
- [13] Mario Rosario Guarracino and Lucia Maddalena. SDI+: A novel algorithm for segmenting dermoscopic images. *IEEE Journal of Biomedical and Health Informatics*, 23(2):481–488, 2018.
- [14] Ghassan Hamarneh, Artur Chodorowski, and Tomas Gustavsson. Active contour models: application to oral lesion detection in color images. In *SMC 2000 Conference Proceedings. 2000 IEEE International Conference on Systems, Man and Cybernetics. Cybernetics Evolving to Systems, Humans, Organizations, and their Complex Interactions (Cat. No. 00CH37166)*, volume 4, pages 2458–2463. IEEE, 2000.
- [15] M Hiraoka, M Firbank, M Essenpreis, M Cope, SR Arridge, P Van Der Zee, and DT Delpy. A Monte Carlo investigation of optical pathlength in inhomogeneous tissue and its application to near-infrared spectroscopy. *Physics in Medicine & Biology*, 38(12):1859, 1993.
- [16] Jia hua Ng, Manu Goyal, Brett Hewitt, and Moi Hoon Yap. The effect of color constancy algorithms on semantic segmentation of skin lesions. In *Medical Imaging 2019: Biomedical Applications in Molecular, Structural, and Functional Imaging*, volume 10953, page 109530R. International Society for Optics and Photonics, 2019.
- [17] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, pages 6389–6399, 2018.
- [18] Ali Madooei and Mark S. Drew. Incorporating colour information for computer-aided diagnosis of melanoma from dermoscopy images: A retrospective survey and critical analysis. *International Journal of Biomedical Imaging*, 2016:1–18, 2016.
- [19] Ali Madooei, Mark S Drew, Maryam Sadeghi, and M Stella Atkins. Automated pre-processing method for dermoscopic images and its application to pigmented skin lesion segmentation. In *Color and Imaging Conference*, volume 2012, pages 158–163. Society for Imaging Science and Technology, 2012.
- [20] Ali Madooei, Mark S Drew, Maryam Sadeghi, and M Stella Atkins. Intrinsic melanin and hemoglobin colour components for skin lesion malignancy detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 315–322. Springer, 2012.
- [21] Teresa Mendonça, Pedro M Ferreira, Jorge S Marques, André RS Marcal, and Jorge Rozeira. PH2 - a dermoscopic image database for research and benchmarking. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 5437–5440. IEEE, 2013.
- [22] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571. IEEE, 2016.
- [23] Zahra Mirikharaji and Ghassan Hamarneh. Star shape prior in fully convolutional networks for skin lesion segmenta-

- tion. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 737–745. Springer, 2018.
- [24] Zahra Mirikharaji, Saeed Izadi, Jeremy Kawahara, and Ghassan Hamarneh. Deep auto-context fully convolutional neural network for skin lesion segmentation. In *Biomedical Imaging (ISBI 2018), 2018 IEEE 15th International Symposium on*, pages 877–880. IEEE, 2018.
- [25] Federico Pollastri, Federico Bolelli, Roberto Paredes, and Costantino Grana. Augmenting data with GANs to segment melanoma skin lesions. *Multimedia Tools and Applications*, May 2019.
- [26] Charles A. Poynton. Frequently asked questions about color, Mar 1997. <https://www.poynton.ca/PDFs/ColorFAQ.pdf> [Accessed: February 25, 2020].
- [27] Howard W Rogers, Martin A Weinstock, Steven R Feldman, and Brett M Coldiron. Incidence estimate of nonmelanoma skin cancer (keratinocyte carcinomas) in the US population, 2012. *JAMA Dermatology*, 151(10):1081–1086, 2015.
- [28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [29] JK Schofield, D Fleming, D Grindlay, and H Williams. Skin conditions are the commonest new reason people present to general practitioners in England and Wales. *British Journal of Dermatology*, 165(5):1044–1050, 2011.
- [30] David W Scott. Multivariate density estimation and visualization. In *Handbook of Computational Statistics*, pages 549–569. Springer, 2012.
- [31] Rebecca L. Siegel, Kimberly D. Miller, and Ahmedin Jemal. Cancer statistics, 2020. *CA: A Cancer Journal for Clinicians*, 70(1):7–30, Jan. 2020.
- [32] Saeid Asgari Taghanaki, Kumar Abhishek, Joseph Paul Cohen, Julien Cohen-Adad, and Ghassan Hamarneh. Deep semantic segmentation of natural and medical images: A review. *arXiv preprint arXiv:1910.07655*, 2019.
- [33] Saeid Asgari Taghanaki, Kumar Abhishek, and Ghassan Hamarneh. Improved inference via deep input transfer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 819–827. Springer, 2019.
- [34] Norimichi Tsumura, Hideaki Haneishi, and Yoichi Miyake. Independent-component analysis of skin color image. *JOSA A*, 16(9):2169–2176, 1999.
- [35] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [36] Xiaohong Wang, Henghui Ding, and Xudong Jiang. Dermoscopic image segmentation through the enhanced high-level parsing and class weighted loss. In *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, Sept. 2019.
- [37] Gunther Wyszecki and WS Stiles. Color science: Concepts and methods, quantitative data and formulae. *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition*, by Gunther Wyszecki, WS Stiles, pp. 968. ISBN 0-471-39918-3. Wiley-VCH, July 2000., page 968, 2000.
- [38] Yuan Xue, Tao Xu, and Xiaolei Huang. Adversarial learning with multi-scale loss for skin lesion segmentation. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, Apr. 2018.
- [39] Lequan Yu, Hao Chen, Qi Dou, Jing Qin, and Pheng-Ann Heng. Automated melanoma recognition in dermoscopy images via very deep residual networks. *IEEE Transactions on Medical Imaging*, 36(4):994–1004, 2017.
- [40] Yading Yuan. Automatic skin lesion segmentation with fully convolutional-deconvolutional networks. *arXiv preprint arXiv:1703.05165*, 2017.
- [41] Y. Yuan, M. Chao, and Y. Lo. Automatic skin lesion segmentation using deep fully convolutional networks with Jac-card distance. *IEEE Transactions on Medical Imaging*, 36(9):1876–1886, Sept 2017.
- [42] Lei Zhang, Guang Yang, and Xujiong Ye. Automatic skin lesion segmentation by coupling deep fully convolutional networks and shallow network with textons. *Journal of Medical Imaging*, 6(2):024001, 2019.