
Estimating Skin Tone and Effects on Classification Performance in Dermatology Datasets

Newton M. Kinyanjui,^{1,4} Timothy Odonga,^{1,4} Celia Cintas,¹ Noel C. F. Codella,²
Rameswar Panda,³ Prasanna Sattigeri,² and Kush R. Varshney^{1,2}

IBM Research, ¹Nairobi, Kenya, ²Yorktown Heights, NY, USA, ³Cambridge, MA, USA

⁴Carnegie Mellon University Africa, Kigali, Rwanda

Abstract

Recent advances in computer vision and deep learning have led to breakthroughs in the development of automated skin image analysis. In particular, skin cancer classification models have achieved performance higher than trained expert dermatologists. However, no attempt has been made to evaluate the consistency in performance of machine learning models across populations with varying skin tones. In this paper, we present an approach to estimate skin tone in benchmark skin disease datasets, and investigate whether model performance is dependent on this measure. Specifically, we use individual typology angle (ITA) to approximate skin tone in dermatology datasets. We look at the distribution of ITA values to better understand skin color representation in two benchmark datasets: 1) the ISIC 2018 Challenge dataset, a collection of dermoscopic images of skin lesions for the detection of skin cancer, and 2) the SD-198 dataset, a collection of clinical images capturing a wide variety of skin diseases. To estimate ITA, we first develop segmentation models to isolate non-diseased areas of skin. We find that the majority of the data in the two datasets have ITA values between 34.5° and 48° , which are associated with lighter skin, and is consistent with under-representation of darker skinned populations in these datasets. We also find no measurable correlation between performance of machine learning model and ITA values, though more comprehensive data is needed for further validation.

1 Introduction

As machine learning is becoming more frequently applied to support decisions that affect the lives of many, understanding how to accurately measure latent dataset characteristics and demographic representation is gaining increasing interest to prevent the potential negative consequences of dataset imbalances [1], henceforth referred to as “dataset bias”. Dataset bias is a critical issue because it is one of the causes of machine learning-based systems placing certain groups of people at a systematic disadvantage [2]. Recognition and mitigation of unwanted bias throughout the machine learning pipeline is necessary in order to build machine learning systems that are trusted in their eventual domains of deployment [3].

Skin diseases continue to bear significant negative impacts on human health globally. A 2017 study found that skin diseases contribute 1.79% to the global burden of disease measured in disability-adjusted life years [4]. Skin diseases such as skin cancer account for about 7% of new cancer cases globally [5], with a cost to the US healthcare system that exceeded 8 billion in 2011 [6]. Within skin cancer cases, there is evidence of some outcome disparities with respect to ethnicity: although people of color are roughly 20 to 30 times less likely to develop melanoma than lighter skinned individuals, for certain melanoma sub-types they have been found to have lower [7, 8, 9] or higher [9] survival rates. Some studies have found that for people of color, the diagnosis of skin cancer

may occur at a more advanced stage, leading to lower rates of survival and poorer outcomes [10, 11]. However, increased screening using currently available diagnostic resources also carries risks and can lead to significant harms, such as unnecessary surgeries, disfigurement, disability, morbidity, and over-diagnosis [7].

Computer vision has been studied in the context of dermatology image analysis for decades, and several review articles have been written [12, 13, 14]. The success of deep learning models has led to studies applying the technology to dermatological use cases [15, 16]. Models using convolutional neural networks (CNNs) have been applied to problems such as skin cancer diagnosis and were found to outperform trained dermatologists in controlled settings and datasets [17, 18, 19]. However, as most of the publicly available datasets of skin images come from lighter skinned populations, due to the extreme disparities in disease prevalence, there are concerns about how to best collect data, train, and evaluate models for darker skinned populations [20, 1]. Also, because of the significant risks of harm from over-diagnosis with increased screening in low-risk dark skin populations, there is a need to better discriminate between life-threatening and stable presentations of disease [7, 1].

In this paper, we work towards answering questions around quantifying distributions according to estimated skin tone in datasets where this information is currently unavailable, and measuring downstream effects on classifier performance. Specifically, our contributions are as follows:

- We propose a pipeline to automatically estimate skin tone for images in two public benchmark skin disease datasets using the individual typology angle (ITA), which has been used previously as a measure of skin tone in absence of manually curated information [21].
- We have generated both a subset of manually delineated ground truth segmentation masks, as well as automatically generated masks, for non-diseased skin in both public benchmarks.
- We quantitatively confirm that the two benchmark skin disease datasets under-represent ITA values that would generally be correlated with darker skin populations.
- No correlation between performance of machine learning model and ITA values are measureable at this time, though more data is needed for conclusive results.

We hope that this work will encourage additional work to help focus resources toward further data collection to fill existing gaps, and improve classifier evaluations for comprehensive assessments across demographics.

2 Related work

Over recent years, there have been significant advances in automated skin lesion analysis, with hundreds of deep learning models implemented for skin cancer diagnosis. Much of this work has been enabled and supported by the International Skin Imaging Collaboration (ISIC) [22, 23, 24], which has organized a public repository of annotated dermoscopic images, and hosted 4 consecutive years of public challenge benchmarks. In 2016, the first work demonstrating classification accuracy higher than the average of expert dermatologists was described [17], employing an ensemble of methods that included hand-coded feature extraction, sparse coding methods, support vector machines, CNNs for skin lesion classification, and fully convolutional networks for skin lesion segmentation. Other models have also been implemented by researchers, such as a computationally efficient skin lesion classification model that uses the MobileNet architecture implemented by [25], and an Inception architecture trained on a large dataset of over 100,000 images [19].

Outside of dermatology, there has been work on evaluating fairness in computer vision with respect to skin type. Recent studies evaluated bias in automated facial analysis algorithms and datasets with respect to phenotypic groups [26, 27]. They found substantial disparities in the accuracy of classifying darker females, lighter females, darker males, and lighter males in gender classification systems with darker-skinned females as the most misclassified group [26, 27]. A related study on gender classification systems tested the stability to skin type of a gender classification system by varying the skin type of a face keeping all other features fixed, statistically showing that the effect of skin type on classification outcome was minimal [28]. The study revealed that well-performing gender classification systems are already invariant to skin type and thus the skin type by itself has a minimal effect on classification disparities [28]. Another study investigated equitable predictive performance in state-of-the-art object detection systems on pedestrians with different skin types,

finding higher precision on lower Fitzpatrick skin types (lighter skin) than higher skin types (darker skin) [29].

Most of the work in automated skin lesion analysis has focused on creating benchmark models. Our work focuses on the analysis of some of the benchmark models with respect to different skin tones that could represent different ethnicities. The work in this paper builds on the work done from previous ISIC challenges by investigating the accuracy of a skin lesion classification model over a range of skin tone categories.

3 Datasets

Public benchmark datasets, in addition to fostering direct comparisons among various algorithms to facilitate advancement in terms of classification performance, are also capable of supporting detailed analysis of that performance with respect to various characteristics of the dataset [1]. Therefore, we focus our analysis on two of the most widely used dermatology datasets in the computer vision literature: the ISIC 2018 Challenge dataset the SD-198 dataset.

ISIC2018: This dataset consists of a collection of dermoscopic images, separated into datasets for 3 tasks of image segmentation (Task 1), clinical feature detection (Task 2), and disease classification (Task 3). Dermoscopic images are acquired through a digital dermatoscope, with relatively low levels of noise and consistent background illumination. The training dataset for Task 3 is the largest among the tasks and used in this work. It consists of 10,015 dermoscopic images publicly available through the ISIC archive [23, 24], falling into one of 7 skin diseases; melanoma, melanocytic nevus, basal cell carcinoma, actinic keratosis, benign keratosis, dermatofibroma and vascular lesion.

SD-198: The SD-198 dataset, a large scale benchmark dataset for visual recognition of skin diseases, contains 6,548 clinical images from 198 skin disease classes, varying according to scale, color, shape and structure, and consists of images downloaded from DermQuest, an online medical resource for dermatologists [30]. Clinical images are collected via various devices, most of which are digital cameras and mobile phones [30]. Higher levels of noise and varying illumination in clinical images makes segmentation more challenging for this dataset in comparison to the dermoscopic images in ISIC2018. Each image in this dataset has been labeled by experts to denote the disease class.

In this work, we preprocess the SD-198 dataset to exclude classes containing images with no observable non-diseased skin. Eventually 136 disease classes are retained, and henceforth the preprocessed dataset is referred to as “SD-136”. Some of the classes excluded from the SD-136 dataset include classes of lesions inside the mouth, such as fibroma, geographic tongue, and stomatitis. Other diseases such as arsenical keratosis, pustular psoriasis, and mal perforans contain images of lesions on palms and soles of the feet from which it is difficult to determine the individual’s skin tone. Other disease classes such as stasis ulcer and eczema contained images of severely scarred skin from which it is visually impossible to differentiate non-diseased and diseased skin. This preprocessing step was done manually and eventually 4,467 images were retained from the original 6,548 images.

Since there are no existing ground truth segmentation masks for the SD-136 dataset, we manually segmented a subset of 343 images using the ImageJ image processing software. We were particularly interested in segmenting regions with non-diseased skin from other regions of the image containing diseased skin, shadows, and other artifacts. We used these ground-truth masks in training the segmentation model for the SD-136 dataset.

4 Methods

The work presented is a two-step approach. First, we quantify the skin tone categories represented in benchmark datasets. For this, we segment skin images to extract non-diseased background regions, and use a metric to characterize the skin tone of that region. Second, we evaluate the classification performance across different skin tone categories.

This is achieved through a series of steps, summarized in Fig. 1. First, we train a model to segment skin disease images to obtain the non-diseased skin in the image. Second, we select and compute the metric to stratify the non-diseased skin into a skin tone category. After that, a classification model is trained to classify skin images into one of the skin diseases in the dataset. Finally the performance of

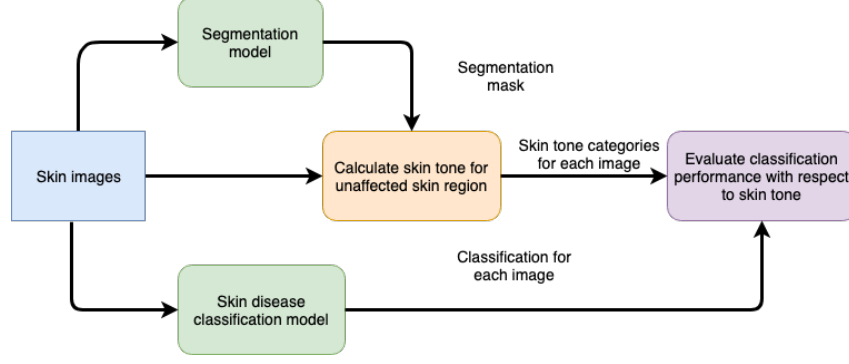


Figure 1: Block diagram of methodology.

the classifier on samples in each skin tone category is evaluated. In the following subsections we will explain each step in detail.

4.1 Quantification of representation of skin tone categories in benchmark datasets

4.1.1 Skin lesion segmentation

Segmentation of the skin lesion from the non-diseased skin is done using a Mask R-CNN model [31]. Mask R-CNN was selected because it was one of the top performing skin lesion boundary segmentation models in the 2018 ISIC challenge [22] and also because it has been shown to be highly effective and efficient in performing semantic segmentation [32].

To obtain a segmentation model for the ISIC2018 dataset, a Mask R-CNN pretrained on the COCO dataset is finetuned with the lesion boundary segmentation data from ISIC 2018 Challenge Task 1. The data is split into training and validation data using 90% to 10% train validation split. The images are resized to 600×450 pixels to correspond to the size used for classification. Random horizontal flips are done on the images for data augmentation during training. The segmentation model for the ISIC2018 dataset is trained for 25 epochs. This model is used to predict segmentation masks for the all the classification data. An additional step of thresholding based on contour extraction on the predicted masks is implemented using the OpenCV library to improve the segmentation masks.

The segmentation model from the ISIC2018 dataset is finetuned with the 343 images from SD-136 having ground truth segmentation masks. The data is split into training and validation data using 90% to 10% train validation split. The images are resized to 450×450 pixels. Random horizontal flips are done on the images for data augmentation during training. The segmentation model for the SD-136 dataset is trained for 50 epochs. This model is used to predict segmentation masks for the all the classification data in SD-136. To get binary masks, a thresholding step is further done on the grayscale masks predicted from the model.

The quality of segmentation is evaluated using the accuracy and the false negative rate of the predicted segmentation masks with respect to the ground-truth segmentation masks. The false negative rate is considered because it is more dangerous to wrongly classify a diseased region as non-diseased in our analysis. To further evaluate the quality of segmentation, mean absolute error is computed between the ITA estimates from ground truth masks and ITA estimates from predicted masks.

4.1.2 Skin tone categorization

With the segmentation masks obtained from the previous step, we are able to obtain pixels in the non-diseased region for each image. We use these pixels to categorize the skin tone of each sample.

There is no universal method for characterizing skin type or skin tone among dermatologists. The Fitzpatrick skin type, used in [26, 27, 28], is a dermatologist’s determination of a person’s risk of sunburn. It is by definition, however, a subjective human determination [33]. In contrast, the melanin index is measured objectively via reflectance spectrophotometry, and has a strong correlation with the Fitzpatrick type and is useful in assigning it [34]. The metric we use to quantify skin tone in this work is the individual typology angle (in degrees) because it has strong (anti-)correlation to the melanin

Table 1: Skin tone categorization scheme.

ITA Range	Skin Tone Category	Abbreviation
$ITA > 55^\circ$	Very Light	very_lt
$48^\circ < ITA \leq 55^\circ$	Light 2	lt2
$41^\circ < ITA \leq 48^\circ$	Light 1	lt1
$34.5^\circ < ITA \leq 41^\circ$	Intermediate 2	int2
$28^\circ < ITA \leq 34.5^\circ$	Intermediate 1	int1
$19^\circ < ITA \leq 28^\circ$	Tanned 2	tan2
$10^\circ < ITA \leq 19^\circ$	Tanned 1	tan1
$ITA \leq 10^\circ$	Dark	dark

index [34], and can be simply computed from images, making it a practical method for categorizing skin color [21].

The pixels from the non-diseased region are converted to CIELab-space to obtain the L and b values. L quantifies the luminance of each pixel and b quantifies the amount of yellow in each pixel. The mean of L and b values are computed. The mean ITA value for the pixels in the non-diseased region is calculated using [21]:

$$ITA = \arctan\left(\frac{L - 50}{b}\right) \times \frac{180^\circ}{\pi}.$$

To deal with outliers, L and b values applied to the formula are kept within one standard deviation. For each image the mean ITA value is computed from the ITAs obtained.

The ITA values computed for the images are binned in a scheme similar to that employed by [35], which used 5 skin tone categories, namely: Very Light, Light, Intermediate, Tanned, and Dark. The Light, Intermediate, and Tanned categories are further subdivided into two equal ranges each for better analysis, thus giving a total of 8 ITA categories. The scheme used to categorize skin tones is summarized in Table 1.

4.2 Evaluation of classification performance across different skin tone categories.

4.2.1 Skin disease classification

To perform skin disease classification, a Densenet201 model pretrained on the ImageNet dataset is finetuned using our training data. The Densenet201 model is chosen because it was one of the best performing single models for lesion classification in the ISIC 2018 challenge [22].

The ISIC2018 dataset comprises 10,015 images which fall under one of 7 disease classes. The images are maintained at 600×450 pixels. Additional transformations such as random horizontal flipping are applied to augment the data. The samples in each batch are normalized using the mean and standard deviation computed on all samples in the dataset to ensure fast convergence during training. The data is split into training and validation data using an 80% to 20% train validation split. A weighted cross entropy loss function and an Adam optimizer are used for training. The weights for the loss function are obtained from the inverse of each disease class frequency. This loss function is chosen because it accounts for class imbalance in the dataset.

The SD-136 dataset comprises 4467 images which fall under one of 136 disease classes. The images are first resized to 450×450 pixels and then center cropped to 360×360 pixels. Additional random transformations including random horizontal flipping and random rotation by a degree selected between -90° and 90° are applied to augment the data during training. Similarly, samples in each batch are normalized using the mean and standard deviation. The data is split into training and validation data using an 80% to 20% train validation split. A weighted cross entropy loss function and an Adam optimizer are used for training.

During training, the early layers up to and including the first Dense block are frozen and all the successive layers have their weights updated. Each classification model is trained for 300 epochs with a patience of 100 epochs at which early stopping would be applied to prevent overfitting.

The performance of the skin disease classification models is measured using the normalized multi-class accuracy metric used in the 2018 ISIC challenge [23].




Figure 2: Sample images from ISIC2018 dataset (top row) and corresponding masks predicted by model (bottom row); ITA is computed on the non-diseased region which is colored black.

Table 2: Correlation of ITA with mean and median grayscale

Dataset	correlation with mean	correlation with median
ISIC2018	0.960	0.981
SD-136	0.912	0.937

4.2.2 Performance evaluation for each skin tone category

The goal is to investigate how classification models perform across the spectrum of skin tones present in the data. The motivation is to investigate if there is bias in performance of the models with respect to skin tone. The accuracy of the classification model is computed for each skin tone category in the validation data.

5 Results

The Mask R-CNN model used for segmentation on the ISIC2018 dataset yields an accuracy of 0.956, a false negative rate of 0.024, and a mean absolute error in ITA computation of 0.428 degrees. The segmentation model on the SD-136 dataset yield an accuracy of 0.802, a false negative rate of 0.076, and a mean absolute error in ITA computation of 3.572 degrees. These are all fairly good results and sufficient for further analysis. Examples with segmented mask and ITA values for both datasets are shown in Fig. 2–3. The ITA values calculated were found to be highly correlated to the mean and median grayscale values of the pixels in the non-diseased region. Table 2 summarizes the correlation values of the ITA with mean an median grayscale values.

Fig. 4 shows the distributions of the ITA values estimated from the non-diseased skin regions of the images in the entire ISIC2018 and SD-136 datasets. Both datasets are found to predominantly lie in the Light category.

On the ISIC2018 dataset, the Densenet201 model achieves an accuracy 0.869 and a balanced accuracy score of 0.814 on the validation partition after training approximately 140 epochs when early stopping occurred. This performance should be in the range of state-of-art systems as the values are similar to those shown on the ISIC 2018 Challenge leaderboard. Although the leaderboard results are computed on a separate held-out test set, we felt the performance of our model was good enough for subsequent analysis on dependence of ITA.

The model trained on SD-136 achieves an accuracy of 0.604 and a balanced accuracy score of 0.601. The benchmark model for SD-198 achieves an accuracy 0.52, as reported in [30]. However, since we

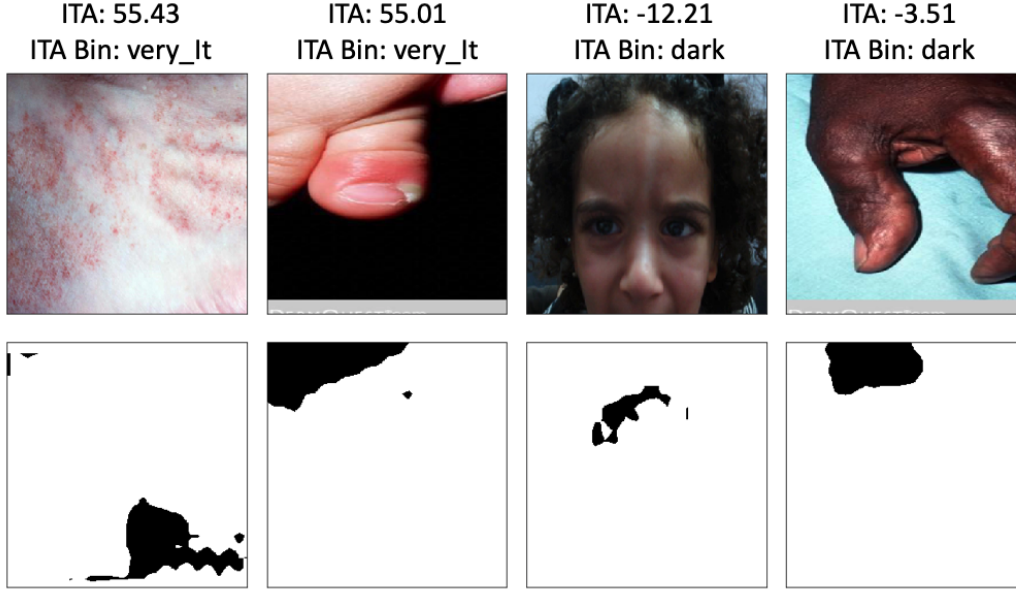


Figure 3: Sample images from SD-136 dataset (top row) and corresponding masks predicted by model (bottom row); ITA is computed on the non-diseased region which is colored black.




Figure 4: ITA skin tone distribution for (a) the entire ISIC2018 dataset, and (b) the entire SD-136 dataset. Most of the samples in the two datasets come from the lt2, lt1, and int2 bins. The SD-136 dataset has a broader distribution across skin tone categories with more samples in the darker region.

dropped the number of classes from 198 to 136, we do not have a benchmark model for comparison. Nonetheless, we are confident that we have a well-performing model.

Importantly, on evaluating the classification performance with respect to skin tone category, our results do not show a clear trend in the performance of the model. Fig. 5 plots classification accuracy versus ITA for the validation set for the two datasets. The error bars indicate the standard error estimated through ten runs with random splits. Table 3 presents the mean accuracy score per skin tone category for the two datasets. The slope of the least squares line of best fit of the mean accuracy versus the midpoint ITA value of the bin for ISIC2018 is -0.000 (per degree) with a 95% confidence interval of $(-0.001, 0.001)$, whereas that for SD-136 is -0.002 (per degree) with a 95% confidence interval of $(-0.003, -0.001)$, which indicate that there are no particular trends in both datasets.

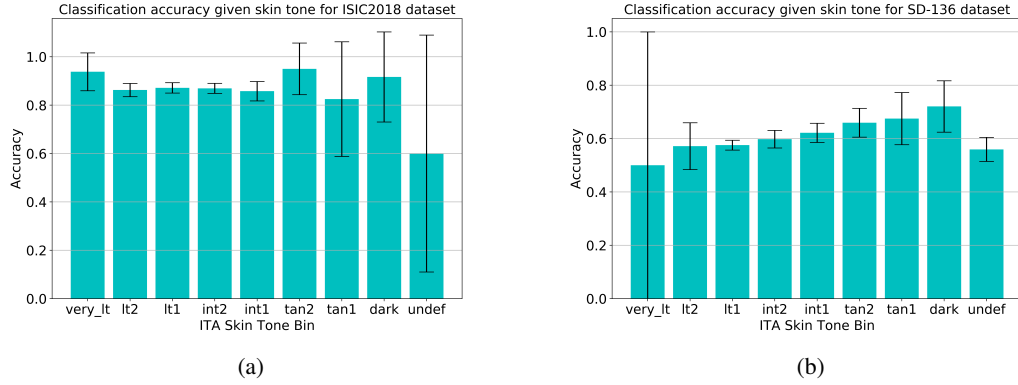


Figure 5: Classification accuracy versus ITA for (a) ISIC2018, and (b) SD-136 validation sets.

Table 3: Classification mean accuracy per skin tone. In the two datasets there is no clear trend in performance of the skin classification model across skin tone categories. The error bars indicate standard deviation across 10 random splits.

Dataset	Overall Accuracy	very_lt	lt2	lt1	int2	int1	tan2	tan1	dark
ISIC2018	0.869	0.94	0.86	0.87	0.87	0.86	0.95	0.83	0.92
SD-136	0.604	0.50	0.57	0.58	0.60	0.62	0.66	0.67	0.72

6 Discussion and Conclusion

In this work, we implemented an approach to measure approximate skin tone distributions in public dermatology image datasets using ITA as an estimator, and evaluated the performance of dermatology classification models with respect to the resultant ITA values.

The distribution of ITA values across both ISIC2018 and SD-136 datasets are consistent with under-representation of darker skin tones. The results from the evaluation of the accuracy of the skin classification model for each skin tone category in the validation data shows that there is no observable trend in the performance of the model with respect to ITA value.

Although we have not found any evidence of model performance bias under the influence of dataset bias in this particular study, further investigation is needed on datasets with more comprehensive representation.

Limitations of this work include the following. 1) ITA is an imperfect estimator of skin tone, therefore the measured distributions may contain some degree of error. 2) Ground truth segmentation masks were lacking in both datasets, requiring automatic annotation on most images, which may have introduced additional error. 3) Studies were carried out on the training dataset for ISIC2018 only. The challenge’s held-out test dataset contains images drawn from additional statistical distributions (other hospitals and geographical locations), suggesting that test set performance may be lower than observed validation performance on the training dataset.

References

- [1] Veronica Rotemberg, Allan Halpern, Stephen W. Dusza, and Noel C. F. Codella. The role of public challenges and data sets towards algorithm development, trust, and use in clinical practice. *Seminars in Cutaneous Medicine and Surgery*, 38(1):E38–E42, March 2019.
- [2] Solon Barocas and Andrew D. Selbst. Big data’s disparate impact. *California Law Review*, 104(3):671–732, June 2016.
- [3] Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Sameep Mehta, Aleksandra Mojsilović, Seema Nagar,

- Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. Think your artificial intelligence software is fair? Think again. *IEEE Software*, 36(4):76–80, July 2019.
- [4] Chante Karimkhani, Robert P. Dellavalle, Luc E. Coffeng, Carsten Flohr, Roderick J. Hay, Sinéad M. Langan, Elaine O. Nsoesie, Alize J. Ferrari, Holly E. Erskine, Jonathan I. Silverberg, Theo Vos, and Mohsen Naghavi. Global skin disease morbidity and mortality: An update from the global burden of disease study 2013. *JAMA Dermatology*, 153(5):406–412, May 2017.
 - [5] Freddie Bray, Jacques Ferlay, Isabelle Soerjomataram, Rebecca L. Siegel, Lindsey A. Torre, and Ahmedin Jemal. Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A Cancer Journal for Clinicians*, 68(6):394–424, November/December 2018.
 - [6] Centers for Disease Control and Prevention. US skin cancer costs rise from 2002 through 2011. Online, November 2014. Available: <https://www.cdc.gov/media/releases/2014/p1110-skin-cancer.html>.
 - [7] Michael A. Marchetti, Esther Chung, and Allan C. Halpern. Screening for acral lentiginous melanoma in dark-skinned individuals. *JAMA Dermatology*, 151(10):1055–1056, October 2015.
 - [8] Xiao-Cheng Wu, Melody J. Eide, Jessica King, Mona Saraiya, Youjie Huang, Charles Wiggins, Jill S. Barnholtz-Sloan, Nicolle Martin, Vilma Cokkinides, Jacqueline Miller, Pragna Patel, Donatus U. Ekwueme, and Julian Kim. Racial and ethnic variations in incidence and survival of cutaneous melanoma in the United States, 1999-2006. *Journal of the American Academy of Dermatology*, 65(5):S26.e1–S26.e13, November 2011.
 - [9] Krishnaraj Mahendraraj, Komal Sidhu, Christine S. M. Lau, Georgia J. McRoy, Ronald S. Chamberlain, and Franz O. Smith. Malignant melanoma in African-Americans: A population-based clinical outcomes study involving 1106 African-American patients from the surveillance, epidemiology, and end result (SEER) database (1988—2011). *Medicine*, 96(15):e6258, April 2017.
 - [10] M. Gohara. Skin cancer: An African perspective. *British Journal of Dermatology*, 173(Suppl. 2):17–21, July 2015.
 - [11] Roopal V. Kundu and Stavonnie Patterson. Dermatologic conditions in skin of color: Part i. special considerations for common skin disorders. *American Family Physician*, 87(12):850–856, June 2013.
 - [12] William V. Stoecker and Randy H. Moss. Editorial: Digital imaging in dermatology. *Computerized Medical Imaging and Graphics*, 16(3):145–150, May–June 1992.
 - [13] Konstantin Korotkov and Rafael Garcia. Computerized analysis of pigmented skin lesions: A review. *Artificial Intelligence in Medicine*, 56(2):69–90, October 2012.
 - [14] Mani Abedini, Qiang Chen, Noel Codella, Rahil Garnavi, Xingzhi Sun, M. Emre Celebi, Teresa Mendonca, and Jorge S. Marques. Accurate and scalable system for automatic detection of malignant melanoma. In M. Emre Celebi, Teresa Mendonca, and Jorge S. Marques, editors, *Dermoscopy Image Analysis*. CRC Press, 2015.
 - [15] M. Emre Celebi, Noel Codella, and Alan Halpern. Dermoscopy image analysis: Overview and future directions. *IEEE Journal of Biomedical and Health Informatics*, 23(2):474–478, March 2019.
 - [16] M. Emre Celebi, Noel Codella, Alan Halpern, and Dinggang Shen. Guest editorial: Skin lesion image analysis for melanoma detection. *IEEE Journal of Biomedical and Health Informatics*, 23(2):479–480, March 2019.
 - [17] Noel C. F. Codella, Quoc-Bao Nguyen, Sharath Pankanti, David A. Gutman, Brian Helba, Allan C. Halpern, and John R. Smith. Deep learning ensembles for melanoma recognition in dermoscopy images. *IBM Journal of Research and Development*, 61(4/5):5, July/September 2016.

- [18] H. A. Haenssle, C. Fink, R. Schneiderbauer, F. Toberer, T. Buhl, A. Blum, A. Kalloo, A. Ben Hadj Hassen, L. Thomas, A. Enk, and L. Uhlmann. Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. *Annals of Oncology*, 29(8):1836–1842, August 2018.
- [19] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, February 2017.
- [20] Adewole S. Adamson and Avery Smith. Machine learning and health care disparities in dermatology. *JAMA Dermatology*, 154(11):1247–1248, November 2018.
- [21] Michele Merler, Nalini Ratha, Rogerio S. Feris, and John R. Smith. Diversity in faces. arXiv:1901.10436, April 2019.
- [22] International Skin Imaging Collaboration. ISIC 2018: Skin lesion analysis towards melanoma detection, 2018. Available: <https://challenge2018.isic-archive.com/>.
- [23] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M. Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, Harald Kittler, and Allan Halpern. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC). arXiv:1902.03368, March 2019.
- [24] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. Data descriptor: The HAM10000 dataset, a large collection of multi-source dermoscopic images of common pigmented skin lesions. *Scientific Data*, 5:180161, August 2018.
- [25] Saket S. Chaturvedi, Kajol Gupta, and Prakash Prasad. Skin lesion analyser: An efficient seven-way multi-class skin cancer classification using MobileNet. arXiv:1907.03220, August 2019.
- [26] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the Conference on Fairness, Accountability and Transparency*, pages 77–91, February 2018.
- [27] Inioluwa Deborah Raji and Joy Buolamwini. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial AI products. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 429–435, January 2019.
- [28] Vidya Muthukumar. Color-theoretic experiments to understand unequal gender classification accuracy from face images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, June 2019.
- [29] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. Predictive inequity in object detection. arXiv:1902.11097, February 2019.
- [30] Xiaoxiao Sun, Jufeng Yang, Ming Sun, and Kai Wang. A benchmark for automatic visual classification of clinical skin disease images. In *Proceedings of the European Conference on Computer Vision*, pages 206–222, October 2016.
- [31] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. arXiv:1703.06870, January 2018.
- [32] Jeremiah W. Johnson. Automatic nucleus segmentation with Mask-RCNN. In *Proceedings of the Computer Vision Conference*, pages 399–407, April 2019.
- [33] Steven Eilers, Daniel Q. Bach, Rikki Gaber, Hanz Blatt, Yanina Guevara, Katie Nitsche, Roopal V. Kundu, and June K. Robinson. Accuracy of self-report in assessing Fitzpatrick skin phototypes I through VI. *JAMA Dermatology*, 149(11):1289–1294, November 2013.
- [34] Marcus Wilkes, Caradee Y. Wright, Johan L. du Plessis, and Anthony Reeder. Fitzpatrick skin type, individual typology angle, and melanin index in an African population. *JAMA Dermatology*, 151(8):902–903, August 2015.
- [35] Giuseppe R. Casale, Anna Maria Siani, Henri Diémoz, Giovanni Agnesod, Alfio V. Parisi, and Alfredo Colosimo. Extreme UV index and solar exposures at Plateau Rosà (3500 m a.s.l.) in Valle d’Aosta Region, Italy. *Science of the Total Environment*, 512–513:622–630, April 2015.