

Detecting Melanoma Fairly: Skin Tone Detection and Debiasing for Skin Lesion Classification

Peter J. Bevan¹ and Amir Atapour-Abarghouei²

¹ School of Computing, Newcastle University, UK peterbevan@hotmail.co.uk

² Department of Computer Science, Durham University, UK

Abstract Convolutional Neural Networks have demonstrated human-level performance in the classification of melanoma and other skin lesions, but evident performance disparities between differing skin tones should be addressed before widespread deployment. In this work, we propose an efficient yet effective algorithm for automatically labelling the skin tone of lesion images, and use this to annotate the benchmark ISIC dataset. We subsequently use these automated labels as the target for two leading bias ‘unlearning’ techniques towards mitigating skin tone bias. Our experimental results provide evidence that our skin tone detection algorithm outperforms existing solutions and that ‘unlearning’ skin tone may improve generalisation and can reduce the performance disparity between melanoma detection in lighter and darker skin tones.

1 Introduction

Convolutional Neural Networks (CNN) have demonstrated impressive performance on a variety of medical imaging tasks, one such being the classification of skin lesion images [10,2,3]. However, there are also many potential pitfalls that must be identified and mitigated before widespread deployment to prevent the replication of mistakes and systematic issues on a massive scale. For example, an issue that is commonly raised in the existing literature is skin tone bias in lesion classification tasks. Groh et al. [9] provide a compiled dataset of clinical lesions with human annotated Fitzpatrick skin type [7] labels, and show that CNNs perform best at classifying skin types similar to the skin types in the training data used. We use the skin type labels in this dataset as the target for supervised debiasing methods to evaluate the effectiveness of these methods at helping melanoma classification models generalise to unseen skin types.

Once we have evaluated the effectiveness of the debiasing methods using human labelled skin tone labels, we look to automate the pipeline further, since human annotated labels are expensive and impractical to gather in practice. We use a novel variation on the skin tone labelling algorithm presented in [16] to annotate the ISIC data and subsequently use these generated labels as the target for a debiasing head, towards creating a fully automated solution to improving the generalisation of models to images of individuals from differing ethnic origins.

In summary, our primary contributions towards the discussed issues are:




Figure 1: Visualisation of the Fitzpatrick 6 point scale [7], widely accepted as the gold standard amongst dermatologists [4].

- *Skin tone detection* - We propose an effective skin tone detection algorithm inspired by [16] (Section 4.2), the results of which can be used as labels for skin tone bias removal.
- *Skin tone debiasing* - We assess the effectiveness of leading debiasing methods [15,1] for skin tone bias removal in melanoma classification, and implement these using automated labels as the target for debiasing (Sections 4.1 and 4.3).

Code is available at <https://github.com/pbevan1/Detecting-Melanoma-Fairly>.

2 Related work

Groh et al. [9] illustrate that CNNs perform better at classifying images with similar skin tones to those the model was trained on. Performance is, therefore, likely to be poor for patients with darker skin tones when the training data is predominantly images of light-skinned patients, which is the case with many of the current commonly-used dermoscopic training datasets such as the ISIC archive data [19,5]. While melanoma incidence is much lower among the black population (1.0 per 100,000 compared to 23.5 per 100,000 for whites), 10-year melanoma-specific survival is lower for black patients (73%) than white patients (88%) or other races (85%) [6], and so it is of heightened importance to classify lesions in patients of colour correctly.

One way to ensure a more even classification performance across skin tones is to re-balance the training data by collecting more high-quality images of lesions on skin of colour, but the low incidence of melanoma in darker skin means this could be a slow process over many years. During the time that unbalanced data continues to be an issue, a robust automated method for removing skin tone bias from the model pipeline could potentially help models to operate with increased fairness across skin tones.

3 Methods

3.1 Debiasing methods

In this work, two leading debiasing techniques within the literature are used, namely ‘Learning Not To Learn’ (LNTL) [15] and ‘Turning a Blind Eye’ (TABE)

[1]. Both are often referred to as ‘unlearning’ techniques because of their ability to remove bias from the feature representation of a network by minimising the mutual information between the feature embedding and the unwanted bias. Generic schematics of both ‘Learning Not to Learn’ and ‘Turning a Blind Eye’ are shown in Figure 2.




Figure 2: ‘Learning Not to Learn’ architecture (left) and ‘Turning a Blind Eye’ architecture (right). f is implemented as a convolutional architecture such as ResNeXt or EfficientNet in this work. ‘fc’ denotes a fully connected layer.

Learning Not to Learn ‘Learning Not to Learn’ (LNTL) [15] introduces a secondary regularisation loss in combination with a gradient reversal layer [8] to remove a target bias from the feature representation of a CNN during training.

The input image, x , is passed into a CNN feature extractor, $f: x \rightarrow \mathbb{R}^K$, where K is the dimension of the embedded feature.

The extracted feature embedding is then passed in parallel into the primary classification head $g: \mathbb{R}^K \rightarrow \mathcal{Y}$ and the secondary bias classification head $h: \mathbb{R}^K \rightarrow \mathcal{B}$. \mathcal{Y} denotes the set of possible lesion classes and \mathcal{B} denotes the target bias classes.

Formulated as a minimax game, h minimises cross-entropy, learning to classify bias from the extracted features, whilst f maximises cross-entropy, restraining h from predicting the bias, and also minimises negative conditional entropy, reducing the mutual information between the feature representation and the bias. The gradient reversal layer between h and f is used as an additional step to remove information relating to the target bias from the feature representation by multiplying the gradient of the secondary classification loss by a negative scalar during backpropagation, further facilitating the feature extraction network, f , to ‘unlearn’ the targeted bias, $b(x)$. On completion of training, f extracts a feature embedding absent of bias information, g uses this feature embedding to perform an unbiased primary classification, and the performance of h has deteriorated because of the resulting lack of bias signal in the feature embedding.

Turning a Blind Eye ‘Turning a Blind Eye’ (TABE) also removes unwanted bias using a secondary classifier, θ_m , m being the m -th bias to be removed. The TABE secondary classifier identifies bias in the feature representation θ_{repr} by

minimising a secondary classification loss, \mathcal{L}_s , and also a secondary confusion loss [22], $\mathcal{L}_{\text{conf}}$, which pushes θ_{repr} towards invariance to the identified bias. The losses are minimised in separate steps since they oppose one another: \mathcal{L}_s is minimised alone, followed by the primary classification loss, \mathcal{L}_p , together with $\mathcal{L}_{\text{conf}}$. The confusion loss calculates the cross entropy between a uniform distribution and the output predicted bias.

As suggested in [15], TABE can also apply gradient reversal (GR) to the secondary classification loss, and is referred to as ‘CLGR’ in this work.

3.2 Skin tone detection

We calculate the individual typology angle (ITA) of the healthy skin in each image to approximate skin tone [16,9], given by:

$$ITA = \arctan\left(\frac{L - 50}{b}\right) \times \frac{180}{\pi}, \quad (1)$$

where L and b are obtained by converting RGB pixel values to the CIELAB colour space. We propose a simpler and more efficient method for isolating healthy skin than the segmentation method used in [16,9]. Across all skin tones, lesions and blemishes are mostly darker than the surrounding skin. Consequently, to select a non-diseased patch of skin, we take 8 samples of 20×20 pixels from around the edges of each image and use the sample with the highest ITA value (lightest skin tone) as the estimated skin tone. The idea behind replacing segmentation with this method is to reduce the impact of variable lighting conditions on the skin tone estimation by selecting the lightest sample rather than the entire healthy skin area. This method is also quicker and more efficient than segmentation methods due to its simplicity.

Eq. 2 shows the thresholds set out in [9], which are taken from [16] and modified to fit the Fitzpatrick 6 point scale [7] (see Figure 1). We use these thresholds in our skin tone labelling algorithm.

$$Fitzpatrick(ITA) = \begin{cases} 1 & ITA > 55 \\ 2 & 55 \geq ITA > 41 \\ 3 & 41 \geq ITA > 28 \\ 4 & 28 \geq ITA > 19 \\ 5 & 19 \geq ITA > 10 \\ 6 & 10 \geq ITA \end{cases} \quad (2)$$

We pre-process each image using black-hat morphology to remove hair, preventing dark pixels from hairs skewing the calculation. This hair removal is purely for skin tone detection and the original images are used for training the debiased classification models. It is clear that even with large lesions with hard-to-define borders, our method is highly likely to select a sample of healthy skin.

3.3 Data

Training data A compilation of clinical skin condition images with human annotated Fitzpatrick skin types [7], called the ‘Fitzpatrick17k’ dataset [9], is




Figure 3: Left of each pair shows ISIC input images, right of each pair shows the placement of the 20×20 pixel samples on images with hair removed. Green square indicates chosen sample based on lightest calculated tone. This sampling method eliminates the need for segmentation.

used for training to demonstrate the effectiveness of unlearning for skin tone debiasing, and to evaluate our automated skin tone labelling algorithm. Of the 16,577 images, we focus on the 4,316 of these that are neoplastic (tumorous). These labels are provided by non-dermatologist annotators, so are likely to be imperfect. When attempting debiasing of ISIC data, a combination of the 2017 and 2020 challenge data [19,5] (35,574 images) is used as training data.

Test data The MClass [3] dataset is used to evaluate generalisation and provide a human benchmark. This dataset comprises a set of 100 dermoscopic images and 100 clinical images (*different* lesions), each with 20 malignant and 80 benign lesions. The human benchmark is the classification performance of 157 dermatologists on the images in the dataset. The Interactive Atlas of Dermoscopy [17], and the ASAN datasets [11] were used to further test the robustness of the models. The Atlas dataset has 1,000 lesions, with one dermoscopic and one clinical image per lesion (2,000 total), while the ASAN test dataset has 852 images, all clinical. Whilst the ISIC training data [19,5] is mostly white Western patients, the Atlas seems to have representation from a broad variety of ethnic groups, and ASAN from predominantly South Korean patients, which should allow for a good test of a model’s ability to deal with different domain shifts.

3.4 Implementation

PyTorch [18] is used to implement the models. The setup used for experimentation consists of two NVIDIA Titan RTX GPUs in parallel with a combined memory of 48 GB on an Arch Linux system with a 3.30GHz 10-core Intel CPU and 64 GB of memory. The source code is publicly released to enable reproducibility and further technical analysis.

After experimentation with EfficientNet-B3 [21], ResNet-101 [12], ResNeXt-101 [23], DenseNet [13] and Inception-v3 [20], ResNeXt-101 looked to show the best performance and so was used as the feature extractor in the debiasing experiments. All classification heads are implemented as one fully-connected layer, as in [15]. Stochastic gradient descent (SGD) is used across all models, ensuring comparability and compatibility between the baseline and debiasing networks.

4 Experimental results

4.1 Fitzpatrick17k skin tone debiasing

A CNN trained using Fitzpatrick [7] types 1 and 2 skin is shown to perform better at classifying skin conditions in types 3 and 4 than types 5 and 6 skin in [9]. We are able to reproduce these findings with our baseline ResNeXt-101 model, trained and tested on the neoplastic subset of the Fitzpatrick17k data. Our objective is to close this gap with the addition of a secondary debiasing head which uses skin type labels as its target. The CLGR configuration proves to be most effective, and is shown in Table 1. The disparity in AUC between the two groups is closed from 0.037 to 0.030, with types 3 and 4 boosted by 1.3% and types 5 and 6 boosted by 2.2%. It is important to note that due to the critical nature of the problem and the significant ramifications of false predictions in real-world applications, even small improvements are highly valuable. This experiment serves as a proof of concept for the mitigation of skin tone bias with unlearning techniques, and gives us precedent to explore this for debiasing the ISIC [19,5] or other similar datasets. Since the ISIC data does not have human annotated skin tone labels, to explore debiasing this dataset we first generate these labels with an automated skin tone labelling algorithm (see section 4.2).

4.2 Automated skin tone labelling algorithm

To validate the effectiveness of our skin tone labelling algorithm, we re-label the Fitzpatrick17k data and compare these automated labels against the human annotated skin tones to calculate accuracy, with a correct prediction being within ± 1 point on the Fitzpatrick scale [9]. Our method achieves 60.61% accuracy, in comparison to the 53.30% accuracy achieved by the algorithm presented in [9], which segments the healthy skin using a YCbCr masking algorithm. The authors of [9] improve their accuracy to 70.38% using empirically selected ITA thresholds, but we decide against using these to label the ISIC data, given that they are optimised to suit only the Fitzpatrick17k data and do not generalise.

We expect our algorithm to perform better still on the ISIC data [19,5] than the Fitzpatrick17k data [9], since the images are less noisy, meaning the assumption that the lightest patch in the image is healthy skin is less likely to be undermined by artefacts or a lightly coloured background.

Experiment	Types 3&4	Types 5&6
Baseline	0.872	0.835
CLGR	0.883	0.853

Table 1: Improving model generalisation to skin tones different to the training data [9]. All scores are **AUC**. Trained using types 1&2 skin images from the Fitzpatrick17k dataset [9], tested on types 3&4 skin and types 5&6.

Figure 4 shows the distribution of Fitzpatrick skin types in the ISIC training data, labelled by our skin tone detection algorithm. The figure shows a clear imbalance towards lighter skin tones. The relatively high number of type 6 classifications could be due to the labelling algorithm picking up on dark lighting conditions, since upon visual inspection of the dataset, it can be concluded that there is not likely to be this many type 6 skin images in the dataset. This is something that should be explored and improved in future work.

4.3 ISIC skin tone debiasing

The ISIC archive is one of the most popular publicly available melanoma training datasets, but there are no skin tone labels available, so we use our skin tone labelling algorithm to analyse the distribution of skin tones in this data as well as to further test the debiasing methods. We also use these labels as the target for the debiasing heads during training. Although these labels have low accuracy, it has been shown that deep learning is still able to learn, even in cases where labels are noisy [14]. We see a small performance improvement across the board when debiasing with the TABE [1] head, indicating that this model generalises to the test sets better than the baseline (see Table 2), including a 5.3% improvement in AUC on the ASAN test set. Performance on this dataset is of particular interest since these images are known to be from Korean patients and so represent a definitive domain shift in comparison to the predominantly Western ISIC training data. The TABE head also prompts a 14.8% increase in performance on the Atlas clinical test set [17] compared to the baseline, and all debiasing heads show noticeable improvements on the MClass dermoscopic and clinical test sets [3]. Although the origins of the Atlas and MClass clinical data are unknown, these also look to be drawn from significantly different populations to the ISIC data (containing many more examples of darker skin tones), so improvements on these test sets could be interpreted as evidence of the mitigation of skin tone bias.




Figure 4: Distribution of Fitzpatrick skin types in ISIC [19,5] training data, as labelled by our algorithm.

Our models demonstrate superior classification performance compared to the group of dermatologists from [3]. While impressive, this comparison should be taken with a grain of salt, as these dermatologists were classifying solely using images and no other information. A standard clinical encounter with each patient would likely result in better performance than this. Moreover, systems like this are not meant to replace the expertise of a dermatologist at this stage, but to augment and enhance the diagnosis and facilitate easier access to certain patients.

Experiment	Atlas		Asan	MClass	
	Dermoscopic	Clinical	Clinical	Dermoscopic	Clinical
Dermatologists	—	—	—	0.671	0.769
Baseline	0.819	0.616	0.768	0.853	0.744
LNTL	0.803	0.608	0.765	0.858	0.787
TABE	0.825	0.707	0.809	0.865	0.859
CLGR	0.820	0.641	0.740	0.918	0.771

Table 2: Comparison of skin tone debiasing techniques, with AUC used as the primary metric. Models are trained using ISIC 2020 & ISIC 2017 data [19,5].

4.4 Ablation studies

TABE [1] with and without gradient reversal has provided impressive results, but ablation of the gradient reversal layer from LNTL [15] led to degraded performance (see Table 3). Deeper secondary heads were experimented with (additional fully-connected layer), but did not have a noticeable impact on performance (see supplementary material).

Experiment	Types 3&4	Types 5&6
LNTL	0.873	0.834
LNTL*	0.867	0.829

Table 3: Ablation of gradient reversal layer from LNTL (ResNext101). Asterisk (*) indicates ablation of gradient reversal.

5 Limitations and future work

As mentioned in section 4.2, the skin tone detection algorithm has a problem with over-classifying type 6 skin which is a key limitation and should be addressed.

ITA is an imperfect method for estimating skin tone, given its sensitivity to lighting conditions, and the Fitzpatrick conversion thresholds are tight and may not generalise well. Empirical calibration of these thresholds tailored to the specific data in question may help, as is done in [9].

Further work may collect dermatologist annotated skin tone labels for dermoscopic datasets and evaluate the effectiveness of debiasing techniques using these human labels. These labels would also allow a more robust evaluation of skin tone bias in the ISIC data than we were able to provide.

Although this work provides potential methods for bias mitigation in melanoma detection, we caution against over-reliance on this or similar systems as silver bullet solutions, as this could further lead to the root cause of the problem (imbalance and bias within the data) being overlooked. We encourage a multi-faceted approach to solving the problem going forward. Further work may also look to do a deeper analysis into the debiasing methods to confirm that the improved generalisation is a result of mitigation of the targeted bias.

6 Conclusion

This work has provided evidence that the skin tone bias shown in [9] can be at least partially mitigated by using skin tone as the target for a secondary debiasing head. We have also presented an effective variation of Kinyanjui et al.’s skin tone detection algorithm [16], and used this to label ISIC data. We have used these labels to unlearn skin tone when training on ISIC data and demonstrated some improvements in generalisation, especially when using a ‘Turning a Blind Eye’ [1] debiasing head. Given that current publicly available data in this field is mostly collected in Western countries, generalisation and bias removal tools such as these may be important in ensuring these models can be deployed to less represented locations as soon as possible in a fair and safe manner.

References

1. Mohsan Alvi, Andrew Zisserman, and Christoffer Nellåker. Turning a Blind Eye: Explicit Removal of Biases and Variation from Deep Neural Network Embeddings. In Laura Leal-Taixé and Stefan Roth, editors, *Computer Vision – ECCV 2018 Workshops*, volume 11129, pages 556–572. Springer International Publishing, Cham, 2019. 2, 3, 7, 8, 9
2. Titus J. Brinker, Achim Hekler, Alexander H. Enk, Joachim Klode, Axel Hauschild, Carola Berking, Bastian Schilling, Sebastian Haferkamp, Dirk Schadendorf, Stefan Fröhling, Jochen S. Utikal, Christof von Kalle, and Collaborators. A convolutional neural network trained with dermoscopic images performed on par with 145 dermatologists in a clinical melanoma image classification task. *European Journal of Cancer (Oxford, England: 1990)*, 111:148–154, April 2019. 1
3. Titus J. Brinker and et al. Hekler, Achim. Comparing artificial intelligence algorithms to 157 German dermatologists: The melanoma classification benchmark. *European Journal of Cancer*, 111:30–37, April 2019. 1, 5, 7, 8

4. Joy Buolamwini and Timnit Gebru. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Conference on Fairness, Accountability and Transparency*, pages 77–91. PMLR, January 2018. [2](#)
5. Noel C. F. Codella, David Gutman, M. Emre Celebi, Brian Helba, Michael A. Marchetti, Stephen W. Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, and Allan Halpern. 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, April 2018. [2](#), [5](#), [6](#), [7](#), [8](#), [13](#)
6. Karen Kadela Collins, Ryan C. Fields, Dadrie Baptiste, Ying Liu, Jeffrey Moley, and Donna B. Jeffe. Racial Differences in Survival after Surgical Treatment for Melanoma. *Annals of Surgical Oncology*, 18(10):2925–2936, October 2011. [2](#)
7. Thomas B. Fitzpatrick. The Validity and Practicality of Sun-Reactive Skin Types I Through VI. *Archives of Dermatology*, 124(6):869–871, June 1988. [1](#), [2](#), [4](#), [6](#)
8. Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-Adversarial Training of Neural Networks. In Gabriela Csurka, editor, *Domain Adaptation in Computer Vision Applications*, pages 189–209. Springer International Publishing, Cham, 2017. [3](#)
9. Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset. *arXiv:2104.09957 [cs]*, April 2021. [1](#), [2](#), [4](#), [6](#), [9](#), [12](#)
10. H.A. Haenssle, C. Fink, R. Schneiderbauer, F. Toberer, T. Buhl, A. Blum, A. Kallou, et al. 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. [1](#)
11. Seung Seog Han, Myoung Shin Kim, Woohyung Lim, Gyeong Hun Park, Ilwoo Park, and Sung Eun Chang. Classification of the Clinical Images for Benign and Malignant Cutaneous Tumors Using a Deep Learning Algorithm. *Journal of Investigative Dermatology*, 138(7):1529–1538, July 2018. [5](#)
12. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, Las Vegas, NV, USA, June 2016. IEEE. [5](#)
13. Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, Honolulu, HI, July 2017. IEEE. [5](#)
14. Lu Jiang, Di Huang, Mason Liu, and Weilong Yang. Beyond Synthetic Noise: Deep Learning on Controlled Noisy Labels, August 2020. *arXiv:1911.09781 [cs, stat]*. [7](#)
15. Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning Not to Learn: Training Deep Neural Networks With Biased Data. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9004–9012, Long Beach, CA, USA, June 2019. IEEE. [2](#), [3](#), [4](#), [5](#), [8](#)
16. Newton M Kinyanjui, Timothy Odonga, Celia Cintas, Noel C F Codella, Rameswar Panda, Prasanna Sattigeri, and Kush R Varshney. Estimating Skin Tone and Effects on Classification Performance in Dermatology Datasets. In *Fair ML for Health*, page 10, Vancouver, Canada, 2019. NeurIPS. [1](#), [2](#), [4](#), [9](#)

17. Peter A. Lio and Paul Nghiem. Interactive Atlas of Dermoscopy: 2000, Edra Medical Publishing and New Media. 208 pages. *Journal of the American Academy of Dermatology*, 50(5):807–808, May 2004. [5](#), [7](#)
18. Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [5](#)
19. Veronica Rotemberg, Nicholas Kurtansky, Brigid Betz-Stablein, Liam Caffery, Emmanouil Chousakos, Noel Codella, Marc Combalia, Stephen Dusza, Pascale Guitera, David Gutman, Allan Halpern, Brian Helba, Harald Kittler, Kivanc Kose, Steve Langer, Konstantinos Liopyris, Josep Malvehy, Shenara Musthaq, Jabpani Nanda, Ofer Reiter, George Shih, Alexander Stratigos, Philipp Tschandl, Jochen Weber, and H. Peter Soyer. A patient-centric dataset of images and metadata for identifying melanomas using clinical context. *Scientific Data*, 8(1):34, January 2021. [2](#), [5](#), [6](#), [7](#), [8](#), [13](#)
20. Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826. IEEE Computer Society, June 2016. [5](#)
21. Mingxing Tan and Quoc V Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *Proceedings of the 36th International Conference on Machine Learning*, 97:6105–6114, 2019. [5](#)
22. Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous Deep Transfer Across Domains and Tasks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4068–4076, December 2015. [4](#)
23. Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated Residual Transformations for Deep Neural Networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [5](#)