



THE UNIVERSITY *of* EDINBURGH

## Edinburgh Research Explorer

### BMFT

Achieving fairness via bias-based weight masking fine-tuning

**Citation for published version:**

Xue, Y, Yan, J, Dutt, R, Haider, F, Liu, J, McDonagh, S & Tsafaris, SA 2024, BMFT: Achieving fairness via bias-based weight masking fine-tuning. in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Fairness of AI in Medical Imaging 2024 Workshop, Marrakesh, Morocco, 10/10/24. <https://doi.org/10.48550/arXiv.2408.06890>

**Digital Object Identifier (DOI):**

[10.48550/arXiv.2408.06890](https://doi.org/10.48550/arXiv.2408.06890)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Peer reviewed version

**Published In:**

International Conference on Medical Image Computing and Computer-Assisted Intervention

**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# BMFT: Achieving Fairness via Bias-based Weight Masking Fine-tuning

Yuyang Xue<sup>1\*</sup>, Junyu Yan<sup>1\*</sup>, Raman Dutt<sup>1,2</sup>, Fasih Haider<sup>1</sup>, Jingshuai Liu<sup>1</sup>,  
Steven McDonagh<sup>1</sup>, and Sotirios A. Tsafaris<sup>1</sup>

<sup>1</sup> School of Engineering, The University of Edinburgh, Edinburgh, EH9 3FG, UK  
{yuyang.xue, junyu.yan, fasih.haider, jliu11, s.mcdonagh, s.tsafaris}@ed.ac.uk

<sup>2</sup> School of Informatics, The University of Edinburgh, Edinburgh, EH8 9AB, UK  
{r.dutt}@sms.ed.ac.uk

**Abstract.** Developing models with robust group fairness properties is paramount, particularly in ethically sensitive domains such as medical diagnosis. Recent approaches to achieving fairness in machine learning require a substantial amount of training data and depend on model re-training, which may not be practical in real-world scenarios. To mitigate these challenges, we propose Bias-based Weight Masking Fine-Tuning (BMFT), a novel *post-processing* method that enhances the fairness of a trained model in significantly fewer epochs without requiring access to the original training data. BMFT produces a mask over model parameters, which efficiently identifies the weights contributing the most towards biased predictions. Furthermore, we propose a two-step debiasing strategy, wherein the feature extractor undergoes initial fine-tuning on the identified bias-influenced weights, succeeded by a fine-tuning phase on a reinitialised classification layer to uphold discriminative performance. Extensive experiments across four dermatological datasets and two sensitive attributes demonstrate that BMFT outperforms existing state-of-the-art (SOTA) techniques in both diagnostic accuracy and fairness metrics. Our findings underscore the efficacy and robustness of BMFT in advancing fairness across various out-of-distribution (OOD) settings. Our code is available at: <https://github.com/vios-s/BMFT>

**Keywords:** Algorithm Fairness · Bias Identification · Bias Removal.

## 1 Introduction

Machine learning is known to exhibit biases from various sources, such as human judgement, inherent algorithmic predispositions, and representation bias [22]. The latter emerges when data from minority ethnicities, genders, and age groups are under-represented, leading to unfair predictions by machine learning systems. This can result in serious issues, including direct or indirect discrimination [24]. It is important to ensure an equitable and responsible application of AI to not undermine the trustworthiness and acceptance of AI solutions by end users [32].

---


\* The two authors share the equal contribution.

Xu et al. [34] identified three key strategies to mitigate bias in medical imaging: pre-processing, e.g. sample re-distribution [27]; in-processing, including learning of disentangled representations [2,19]; and *post-processing*, involving weight pruning [21,33]. While pre- and in-processing are effective in producing fair predictions, a challenge remains: *Access to original medical training data, post-deployment in the real-world, is often restricted*. Mao et al. [20] showed that empirical risk minimisation (ERM) on imbalanced data can effectively capture features for classification within the representation domain. This suggests that post-processing can overcome bias with only a small amount of external data, e.g. by fine-tuning or pruning weights, thus eliminating the need for complete model retraining on the large-scale original training dataset.

Weights learnt with ERM encode two types of features: *core* features that contribute to the target task and *bias* features that have captured spurious correlations or irrelevant information in the data and could therefore exhibit bias during decision-making. It is well understood that these two distinct types of features are entangled in a complex manner [17,35]. Current weight pruning strategies [33,15] aim to identify and remove weights that contribute to bias, based on certain heuristics. However, this can cause substantial performance degradation mainly due to the complex entanglement between *core* features and bias features, especially when a model is trained from random initialisation [30]. It is well known that fine-tuning from a pre-trained initialisation leads to faster convergence and better performance [28]. However, if fine-tuning is indiscriminately applied to all weights in pursuit of fairness, there is the potential of updating *core* features which have previously captured essential discriminative information, therefore risking predictive performance power, even if model fairness gets improved. Moreover, fine-tuning all model parameters with limited data is prone to overfitting [8]. A potential solution is to fine-tune a subset of model weights to lessen the effect of biased feature on the decision boundary [17].

We propose Bias-based Weight Masking Fine-Tuning (BMFT), a fast, retrain-free, post-processing debiasing approach designed to enhance fairness while preserving performance, which utilises a small external dataset and eliminates the need for original training data. Different from other mask-based methods that implement iterative re-training on the original data [7], *BMFT* firstly generates a parameter update mask which highlights the model parameters contributing to bias. Then, a two-step process fine-tunes the feature extractor on mask-selected weights for feature debiasing, which is followed by classification layer fine-tuning to efficiently integrate core features and improve classification performance. This strategy can lead to remarkable performance improvements in terms of prediction fairness and accuracy within significantly fewer training steps.

Our **contributions** can be summarised as follows: **(1)** We propose a fast mask generation technique to filter weights contributing to bias with a small external dataset and without needing the original training data. **(2)** We introduce a two-step fine-tuning strategy that is model-agnostic and contributes to bias mitigation without sacrificing the model’s performance. **(3)** We conduct exten-



**Fig. 1.** *BMFT* is a masked-based fine-tuning post-processing approach. (1) A weight mask is generated by calculating weight importance to the bias and loss functions. (2) The masked weights of the feature extractor  $E(\cdot)$  are fine-tuned to debias, followed by fine-tuning a reinitialised classification layer  $C(\cdot)$  to maintain model performance.

sive experiments across different datasets and sensitive attributes to reveal the superiority of our proposed approach compared to SOTAs.

## 2 Method

Bias in model prediction typically originates from two key sources: (1) the *entanglement* between noisy, harmful (spurious or irrelevant) features and, useful, *core* features within the feature extractor; and (2) the *incorrect composition* of representation in the classification layer, leading to the core features information loss [20]. **The main idea** of *BMFT* is to start with debiasing the feature extractor, by fine-tuning only the targeted weights that are disproportionately instigating bias. Subsequently, as the first process may affect how the core features are integrated by the classification layer for prediction, *BMFT* then fine-tunes a reinitialised classification layer.

### 2.1 Preliminaries

Consider an external dataset of  $N$  subjects,  $\mathcal{D}_e = \{\mathbf{x}_i, \mathbf{y}_i, \mathbf{a}_i\}_{i=1}^N$ , where  $\mathbf{x}_i$ ,  $\mathbf{y}_i$ , and  $\mathbf{a}_i$  are the input image, the corresponding label, and the sensitive attribute, respectively. For simplicity, we assume binary targets and attributes (i.e.  $\mathbf{y}_i, \mathbf{a}_i \in \{0, 1\}$ ). We will assume a compositional construction of a biased model  $f_\theta(\cdot)$ , consisting of a feature extractor  $E(\cdot)$  and a classification head  $C(\cdot)$ , pre-trained by the original training data  $\mathcal{D}_{\text{org}}$ . Our goal is to debias the model  $f_\theta(\cdot)$  w.r.t. fairness metrics such as Demographic Parity (DP) [9] or Equalised Odds (EOdds) [13] by fine-tuning with significantly fewer epochs, using  $\mathcal{D}_e$ ; and further enhance the generalisability of the model on various OOD test datasets  $\mathcal{D}_t$ , which share the same sensitive attribute  $\mathbf{a}_i$ . To avoid attribute bias from the *external dataset*  $\mathcal{D}_e$ , we sample a *group-balanced subset*  $\mathcal{D}_r$  from  $\mathcal{D}_e$ . This ensures that each attribute group  $\mathbf{a}_i$  has the same amount of data. The proportion of positive to negative labelled images is adjusted to be the same in each  $\mathbf{a}_i$  group, which prevents bias due to different label distributions. Implementing these two

strategies prepares the model for debiasing by providing equal representation across attribute groups and equal ratios of label types. Choosing proportionally over numerical equality in  $\mathcal{D}_r$  helps when there are few positive or negative labels in  $\mathcal{D}_e$ , thus avoiding a reduction in the overall sample size.

## 2.2 Bias-Importance-Based Mask Generation

To identify the effective parameters for updating, we evaluated the importance of each parameter with respect to the designed bias influence function and loss function. The top  $K$  parameters, that are highly correlated with bias and yet provide only minimal contribution to prediction, were selected as the masking weights. We adopt weighted Binary Cross Entropy (WBCE) as loss function since the dermoscopic dataset contains label imbalance. Together with the bias influence function  $\mathcal{B}$ , the differentiable proxy of EOdds [13], is given by:

$$\mathcal{L}_{\text{WBCE}} = \sum_i \left( -\frac{N_n}{N_n + N_p} \mathbf{y}_i \log(p_i) - \frac{N_p}{N_n + N_p} (1 - \mathbf{y}_i) \log(1 - p_i) \right), \quad (1)$$

$$\mathcal{B}(f_\theta, \mathbf{x}_i, \mathbf{y}_i, \mathbf{a}_i) = \left| \frac{\sum_{i=1}^N f_\theta(\mathbf{x}_i)(1 - \mathbf{a}_i)\mathbf{y}_i}{\sum_{i=1}^N (1 - \mathbf{a}_i)\mathbf{y}_i} - \frac{\sum_{i=1}^N f_\theta(\mathbf{x}_i)\mathbf{a}_i\mathbf{y}_i}{\sum_{i=1}^N \mathbf{a}_i\mathbf{y}_i} \right|, \quad (2)$$

where  $p_i$  is the prediction probability,  $N_p$  is the number of images with positive label,  $N_n$  the number of images with negative label. For a pre-trained model  $f_\theta(\cdot)$  with effectively learnt parameters  $\theta^*$ , the sensitivity for each parameter  $\theta$  can be obtained by the second-order derivative of loss close to the minimum [10]. This value can be interpreted as the importance of the parameter for the prediction. The diagonal of the Fisher Information Matrix (FIM) is equivalent to the second-order derivative of the loss [26], which provides an efficient computation through first-order derivatives. The first-order expression for the FIM is provided by:

$$\mathcal{I}(f_\theta, \mathcal{D}_e) = \mathbb{E} \left[ \left( \frac{\partial \log f_\theta(\mathcal{D}_e|\theta)}{\partial \theta} \right) \left( \frac{\partial \log f_\theta(\mathcal{D}_e|\theta)}{\partial \theta} \right)^\top \Big|_\theta \right]. \quad (3)$$

We employ Eq. 3 to quantify the importance of each weight that contributes to the prediction. The importance of each weight for the bias can be computed using the second derivative of the bias influence function (that is, Eq. 2). Noting that the bias influence function is a linear function of the prediction, and thus replaces the loss function with the bias influence function during backpropagation, the importance of a parameter w.r.t. the bias can also be derived by Eq. 3. We denote the weight mask as  $M_i$ , where  $i$  is the weight index.

To ensure that the weight importance for bias or prediction is not skewed in favour of a particular subgroup, in practice, we use the group-balanced dataset  $\mathcal{D}_r$  to calculate the importance. To select top  $K\%$  weights in the mask  $M_i$ , we use the bias importance  $\mathcal{I}_{i,b}$  and the loss importance  $\mathcal{I}_{i,l}$  of each parameter as:

$$M_i = \begin{cases} 1, & \text{if } \mathcal{I}_{i,b}/\mathcal{I}_{i,l} > \alpha \\ 0, & \text{otherwise} \end{cases}, \quad (4)$$

where  $\mathcal{I}_{i,b}$  and  $\mathcal{I}_{i,l}$  denote the importance of each parameter in relation to bias and loss, respectively. The hyperparameter  $\alpha$  distinguishes the significance of the weight to bias-induced features as opposed to core features, calculated by using the weight selection ratio  $K$ , defined as the top  $K$ -th percentile of the ratios of importance of the parameters. An illustrative example of a mask is provided in the supplementary material.

### 2.3 Impair-Repair: Two-Step Fine-Tuning

To disentangle core features from biased features, whilst preserving predictive capability, a two-stage fine-tuning process is structured in an ‘‘impair-repair’’ manner. Firstly, we fine-tune the selected masked parameters within the feature extractor  $E(\cdot)$  using the group-balanced dataset  $\mathcal{D}_r$ , guided by an objective function (Eq. 2). Next, we fine-tune the entire reinitialised classification layer  $C(\cdot)$ . The weights identified by the mask in the feature extractor are closely linked to attribute bias and play a smaller role in predicting the class. This relationship helps us remove the unfavourable influence of bias contained in the model, and yet crucially also retain information embedded in core features.

We highlight that the pre-trained model has the capability to sufficiently extract core features. Further, eliminating the influence of (only) bias-effected weights suggests moving the model decision boundary away from bias features. Moreover, fine-tuning is prone to overfitting on a small dataset. These observations suggest that small epochs are sufficient to realise the proposed debiasing fine-tuning strategy.

An unbiased model, trained with a balanced dataset, will have the majority of its classification layer weights being zero, due to minimisation of irrelevant feature impact on class prediction [17]. To this end, we first reinitialise the weights and biases in the classification layer  $C(\cdot)$  to zero to enable fast convergence to the optimal unbiased state and then fine-tune the layer on  $\mathcal{D}_r$  using a loss which combines the WBCE loss  $\mathcal{L}_{\text{WBCE}}$  and fairness constraints (i.e. Eq. 2), as:

$$\mathcal{L} = \mathcal{L}_{\text{WBCE}}(f_\theta) + \beta\mathcal{B}(f_\theta), \quad (5)$$

where the fairness constraint  $\mathcal{B}(f_\theta)$  modulates the emphasis of the features, and the hyperparameter  $\beta$  establishes an efficiently identified core trade-off between class prediction accuracy and fairness. The improved feature extractor retains core features without being influenced by attribute bias, allowing the model to efficiently identify and recover its predictive performance in limited epochs. To balance bias reduction with efficiency, an equal number of epochs is used for both the ‘‘impair’’ and ‘‘repair’’ stages. Our empirical findings indicate that only 10% of the total original training epochs are adequate for successful fine-tuning.

## 3 Experiments and Results

### 3.1 Dataset and Data-processing

**ISIC Challenge Training Dataset.** The International Skin Imaging Collaboration (ISIC) challenge dataset is a collection of dermoscopic images of melanoma

classification, complete with diagnostic labels and metadata. A combination of the 2017 [4], 2019 [31,5] and 2020 [29] ISIC challenge data is used as the training dataset (46,938 images), while the 2018 [3] data is reserved as the *external dataset*  $\mathcal{D}_e$  (9,925 images, excluding any overlap with the training data). We select 4,024 images for skin tone and 9,084 for gender from  $\mathcal{D}_e$  to build *group-balanced subset*  $\mathcal{D}_r$ . We maintain the same skin tone annotation as [1] for training data. All images are pre-processed with centre-cropping and resizing to size  $256 \times 256$ . We consider two available sensitive attributes (skin tone and gender).

**Test Datasets.** Our study employs four different OOD datasets with skin images for melanoma detection. Fitzpatrick-17k [11] contains 16,577 images across six skin tone levels, which we group into light (1-3) and dark (4-6) categories. The PAD-UFES-20 [23] includes 2,298 images, with detailed metadata such as gender, ethnicity, and skin tone. Interactive Atlas of Dermoscopy (Atlas) [18] provides 1,011 images with gender, and DDI [6] offers 656 images with age and skin information. These datasets expand diversity, aiding in evaluating our model’s generalisability across attributions. Details are in the supplementary material.

### 3.2 Implementation

We conducted experiments in PyTorch [25] using NVIDIA A100 40GB GPUs. We adopt different variants of ImageNet pretrained ResNet architecture (ResNet-18, 34, 50, and 101) [14]. The learning rates for impair and repair processes were set at 0.001 and 0.003 respectively, using Adam optimiser [16], with a batch size of 64. We perform a hyperparameters sweep and selected  $K = 50$ , and  $\beta$  was selected as 0.02 for the ResNet50 model (details on hyperparameter tuning are shown in supplementary materials). To address class-imbalance, we adopt the WBCE loss function  $\mathcal{L}_{WBCE}$ . Furthermore, we use accuracy (ACC), area under the receiver operating characteristic curve (AUC) as primary performance metrics, and equalised odds (EOdds [13]) as fairness metric, as done previously [7,12,33]. Our post fine-tuning method requires just an additional five epochs, which is 10% of the initial 50-epoch training duration. Code will be released soon.

### 3.3 Results

We compare with baselines and SOTA models which we describe below. We pre-trained the ResNet model on the training dataset  $\mathcal{D}_{org}$  as a *Baseline*. Performance is evaluated using established metrics. Full fine-tuning (*Full FT*), a widely adopted fine-tuning practice, adjusts all weights on the external dataset  $\mathcal{D}_e$  with the loss function specified in Eq. 5. *FairPrune* [33] improves fairness by pruning parameters and prioritises subgroup accuracy; however, the optimal hyperparameter configuration is specific to model and dataset. *LLFT* [20] fine-tunes only the last layer of a deep classification model to promote fairness. Similarly, *Diff-Bias* [21] fine-tunes on an external dataset, using a bias-aware loss function to steer network optimisation. We highlight that a good fairness score does not constitute the sole criterion for success and that a model should

**Table 1.** Comparison of debiasing methods on skin tone and gender attributes, for a ResNet-50 model. We re-implemented SOTA methods (denoted with \*) or used publicly available code (<sup>†</sup>). *FairPrune* had difficulty in classifying the Atlas dataset.

Methods	Skin Tone									Gender					
	Fizpatric17k			PAD-UFES-20			DDI			Atlas			PAD-UFES-20		
	ACC <sup>†</sup>	AUC <sup>†</sup>	EOdds <sup>↓</sup>	ACC <sup>†</sup>	AUC <sup>†</sup>	EOdds <sup>↓</sup>	ACC <sup>†</sup>	AUC <sup>†</sup>	EOdds <sup>↓</sup>	ACC <sup>†</sup>	AUC <sup>†</sup>	EOdds <sup>↓</sup>	ACC <sup>†</sup>	AUC <sup>†</sup>	EOdds <sup>↓</sup>
Baseline	0.719	0.521	0.1745	0.858	0.506	0.0412	0.744	0.646	0.1288	0.574	0.614	0.0348	0.858	0.506	0.0167
Full FT	0.703	0.522	0.1510	0.725	0.501	<b>0.0321</b>	0.727	0.618	0.2059	0.690	0.655	0.0303	0.725	0.501	0.0730
FairPrune [33]*	0.763	0.547	0.0120	0.698	0.492	0.1237	0.737	0.536	<b>0.0024</b>	\	\	\	0.634	0.499	0.1456
LLFT [20] <sup>†</sup>	0.579	0.513	0.1446	0.694	0.492	0.0726	0.646	0.561	0.1197	0.724	0.589	0.0042	<b>0.964</b>	0.482	<b>0.0007</b>
Diff-Bias [21] <sup>†</sup>	0.716	0.507	0.1144	0.884	0.501	0.0599	0.737	0.629	0.0741	<b>0.769</b>	0.687	0.0383	0.878	0.504	0.0461
BMFT (Ours)	<b>0.865</b>	<b>0.637</b>	<b>0.0055</b>	<b>0.902</b>	<b>0.507</b>	0.0657	<b>0.759</b>	<b>0.736</b>	0.0754	0.752	<b>0.876</b>	<b>0.0037</b>	0.893	<b>0.512</b>	0.0365

also afford accurate classification behaviour, represented here by high ACC and AUC. The results presented here correspond to the ResNet-50 model and analogous trends were observed across the investigated models, with further details present in the supplementary materials.

**Fine-tuning can help achieve fairness.** We start by evaluating the behaviour of fine-tuning methods. In terms of fairness, we observed significant improvements both in AUC and EOdds across the majority of test datasets when employing fine-tuning-based methods (Table 1). This indicates that fine-tuning can improve the accuracy gaps between two attribute groups and improve fairness metrics. Fine-tuning with fairness constraints, i.e. *LLFT* and *Diff-Bias*, exhibit lower EOdds compared to using a WBCE loss, i.e. *Full FT*. This illustrates that incorporating the bias influence function into the loss has measurable benefits.

**Pruning is not always the best option.** Model pruning, a popular technique for model simplification towards reducing bias, does not always lead to fairer outcomes, c.f. fine-tuning methods. Our results demonstrate that pruning achieves low AUC across most test datasets, even under-performing the baseline. For example, the performance of *FairPrune* drops drastically under the Atlas dataset, as detailed in Table 1, regardless of hyperparameter values. It is “easier” to obtain good fairness by finding less useful features for classification. This leads to poor classification results in sensitive groups. We conjecture that such performance declines are due to pruning which removes neurons containing *both* bias and core features that are important for classification, therefore causing undesirable information loss. Balancing between fairness and classification performance through pruning requires extensive iterations and is heavily dependent on hyperparameters and datasets. Instead, our method considers both classification capability and fairness to force the selection of core features.

**Masked fine-tuning utility.** Traditional fine-tuning techniques struggle to achieve a balance between classification performance and fairness. *Full FT* and *Diff-Bias* improve all metrics only in the Atlas dataset for the gender attribute (see Table 1). This illustrates that fine-tuning all layers guides model behaviour towards the distribution of the external dataset, therefore necessitating similar distributions between external and test datasets. Moreover, this may affect weights that are useful for prediction in lieu of fairness, leading to degradation in classification performance. *LLFT* shows improvement in EOdds, but performs



**Table 2.** Comparison of the random mask, manual masks, and our proposed mask generation method on both skin and gender attributes, using ResNet-50 backbone.

Methods	DDI (Skin Tone)			Atlas (Gender)		
	ACC $\uparrow$	AUC $\uparrow$	EOdds $\downarrow$	ACC $\uparrow$	AUC $\uparrow$	EOdds $\downarrow$
Random	0.733	0.581	0.1194	0.741	0.629	0.0105
Manual Top Layers	0.749	0.670	0.0948	0.753	0.644	0.0169
Manual BN Layers	0.741	0.635	0.1116	<b>0.757</b>	0.661	0.0052
BMFT (Ours)	<b>0.759</b>	<b>0.736</b>	<b>0.0754</b>	0.752	<b>0.876</b>	<b>0.0037</b>

poorly in terms of AUC (0.589 in Atlas), showing that fine-tuning the last layer from scratch requires all core features that are well captured and isolated from the bias features; otherwise, fairness and prediction irrelevant features may be selected by the classification layer for achieving fairness. The proposed method, *BMFT*, manages to find an ideal balance and exhibits the best performance in terms of ACC, AUC and EOdds on most test datasets.

**Different bias, different difficulty.** Comparing *Baseline* results for gender and skin tone attributes in the PAD-UFES-20 dataset (see Table 1), the pre-trained model demonstrates less inherent bias for gender, indicated by a lower EOdds. Furthermore, Table 1 shows that *BMFT* achieves more prominent EOdds and AUC improvements for the gender attribute c.f. the skin tone. This finding suggests that attributes with less inherent bias in the pre-trained model are easier for bias mitigation strategies to rectify. Moreover, the presence of less inherent bias suggests that core and biased features are less entangled, which simplifies the task of maintaining model performance during mitigation process.

**Mask: better than manual, superior to random.** We compare our mask with random masking, manual shallow convolutional (Top) masking and batch normalisation (BN) masking in Table 2. The latter two options are motivated by analysis of the mask distribution, where we saw that mask-selected weights exist mainly in shallow convolutional layers and BN. Compared to random, the results align with our hypothesis that shallow convolutional layers and BN layers contribute to spurious correlation. Our method outperforms the baselines by  $\sim 0.2$  in AUC, illustrating the effectiveness of our fast mask generation process. With masking, spurious-related weights can be fine-tuned, and core-feature-related weights will remain intact, leading to improved fairness and prediction efficacy.

## 4 Conclusion

Our study advances bias mitigation in discriminative models for dermatological disease. The proposed *BMFT* distinguishes the core features from biased features, enhancing fairness without sacrificing classification performance. This two-step impair-repair fine-tuning approach can effectively reduce bias within minimal epochs, with only 10% of the original training effort, providing an efficient solution when computing resources or data access are limited. Our method is agnostic to the choice of pre-trained models. Our future work will explore a

broader range of tasks and multiple attribute-label pairs. The challenge of obtaining diverse, well-labelled public datasets with comprehensive meta-data or sensitive attributes remains a limitation on our work.

**Acknowledgments.** Y. Xue and J. Yan thank additional financial support from the School of Engineering, the University of Edinburgh. J. Yan also thanks the support of Advanced Care Research Center. S.A. Tsafaris also acknowledges the support of Canon Medical and the Royal Academy of Engineering and the Research Chairs and Senior Research Fellowships scheme (grant RCSR1819\8\25), of the UK's Engineering and Physical Sciences Research Council (EPSRC) (grant EP/X017680/1) and the National Institutes of Health (NIH) grant 7R01HL148788-03. We thank Dr. Edward Moroshko for his help and support.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Bevan, P., Atapour-Abarghouei, A.: Skin deep unlearning: Artefact and instrument debiasing in the context of melanoma classification
2. Bissoto, A., Valle, E., Avila, S.: Debiasing skin lesion datasets and models? not so fast. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 740–741 (2020)
3. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., et al.: Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC). arXiv preprint arXiv:1902.03368 (2019)
4. Codella, N.C., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., Dusza, S.W., Kalloo, A., 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). pp. 168–172. IEEE (2018)
5. Combalia, M., Codella, N.C., Rotemberg, V., Helba, B., Vilaplana, V., et al.: Bcn20000: Dermoscopic lesions in the wild. arXiv preprint arXiv:1908.02288 (2019)
6. Daneshjou, R., Vodrahalli, K., Novoa, R.A., Jenkins, M., Liang, W., Rotemberg, V., Ko, J., Swetter, S.M., et al.: Disparities in dermatology ai performance on a diverse, curated clinical image set. *Science advances* **8**(31), 6147 (2022)
7. Dutt, R., Bohdal, O., Tsaftaris, S.A., Hospedales, T.: FairTune: Optimizing parameter efficient fine tuning for fairness in medical image analysis. arXiv preprint arXiv:2310.05055 (2023)
8. Dutt, R., Ericsson, L., Sanchez, P., Tsaftaris, S.A., Hospedales, T.: Parameter-efficient fine-tuning for medical image analysis: The missed opportunity. arXiv preprint arXiv:2305.08252 (2023)
9. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. pp. 214–226 (2012)
10. Foster, J., Schoepf, S., Brintrup, A.: Fast machine unlearning without retraining through selective synaptic dampening. arXiv preprint arXiv:2308.07707 (2023)
11. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1820–1828 (2021)
12. Han, S.S., Kim, M.S., Lim, W., Park, G.H., Park, I., Chang, S.E.: Classification of the clinical images for benign and malignant cutaneous tumors using a deep learning algorithm. *Journal of Investigative Dermatology* **138**(7), 1529–1538 (2018)
13. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems* **29** (2016)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
15. Huang, Y.Y., Chiuwanara, V., Lin, C.H., Kuo, P.C.: Mitigating bias in mri-based alzheimer’s disease classifiers through pruning of deep neural networks. In: Workshop on Clinical Image-Based Procedures. pp. 163–171. Springer (2023)
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
17. Le, P.Q., Schlötterer, J., Seifert, C.: Is last layer re-training truly sufficient for robustness to spurious correlations? arXiv preprint arXiv:2308.00473 (2023)

18. Lio, P.A., Nghiem, P.: Interactive atlas of dermoscopy. *Journal of the American Academy of Dermatology* **50**(5), 807–808 (2004)
19. Liu, X., Thermos, S., Chartsias, A., O’Neil, A., Tsaftaris, S.A.: Disentangled representations for domain-generalized cardiac segmentation. In: *Statistical Atlases and Computational Models of the Heart. M&Ms and EMIDEC Challenges: MICCAI 2020*, October 4, 2020, Revised Selected Papers 11. pp. 187–195. Springer (2021)
20. Mao, Y., Deng, Z., Yao, H., Kawaguchi, K., Zou, J.: Last-layer fairness fine-tuning is simple and effective for neural networks. *arXiv preprint arXiv:2304.03935* (2023)
21. Marcinkevics, R., Ozkan, E., Vogt, J.E.: Debiasing deep chest x-ray classifiers using intra-and post-processing methods. In: *Machine Learning for Healthcare Conference*. pp. 504–536. PMLR (2022)
22. Mehrabi, N., Morstatter, F., Saxena, N., Galstyan, A.: A survey on bias and fairness in machine learning. *arxiv* 2019. *arXiv preprint arXiv:1908.09635* (2019)
23. Pacheco, A.G., Lima, G.R., Salomao, A.S., Krohling, B., Biral, I.P., de Angelo, G.G., et al.: Pad-ufes-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones. *Data in Brief* **32**, 106221 (2020)
24. Pagano, T.P., Loureiro, R.B., et al.: Bias and unfairness in machine learning models: a systematic review on datasets, tools, fairness metrics, and identification and mitigation methods. *Big data and cognitive computing* **7**(1), 15 (2023)
25. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., et al.: PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019)
26. Pawitan, Y.: *In all likelihood: statistical modelling and inference using likelihood*. Oxford University Press (2001)
27. Puyol-Antón, E., Ruijsink, B., Piechnik, S.K., Neubauer, S., et al.: Fairness in cardiac MR image analysis: an investigation of bias due to data imbalance in deep learning based segmentation. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021*, September 27–October 1, 2021, Proceedings, Part III 24. pp. 413–423. Springer (2021)
28. Raghu, M., Zhang, C., et al.: Transfusion: Understanding transfer learning for medical imaging. *Advances in Neural Information Processing Systems* **32** (2019)
29. Rotemberg, V., Kurtansky, N., Betz-Stablein, B., Caffery, L., Chousakos, E., Codella, N., et al.: A patient-centric dataset of images and metadata for identifying melanomas using clinical context. *Scientific data* **8**(1), 34 (2021)
30. Tran, C., Fioretto, F., et al.: Pruning has a disparate impact on model accuracy. *Advances in Neural Information Processing Systems* **35**, 17652–17664 (2022)
31. Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data* **5**(1), 1–9 (2018)
32. Winkler, J.K., Fink, C., Toberer, F., Enk, A., Deinlein, T., Hofmann-Wellenhof, R., et al.: Association between surgical skin markings in dermoscopic images and diagnostic performance of a deep learning convolutional neural network for melanoma recognition. *JAMA dermatology* **155**(10), 1135–1141 (2019)
33. Wu, Y., Zeng, D., Xu, X., Hu, J.: FairPrune: Achieving fairness through pruning for dermatological disease diagnosis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 743–753. Springer (2022)
34. Xu, Z., Li, J., Yao, Q., Li, H., Zhou, S.K.: Fairness in medical image analysis and healthcare: A literature survey. *Authorea Preprints* (2023)
35. Ye, H., Zou, J., Zhang, L.: Freeze then train: Towards provable representation learning under spurious correlations and feature noise. In: *International Conference on Artificial Intelligence and Statistics*. pp. 8968–8990. PMLR (2023)

# BMFT: Achieving Fairness via Bias-based Weight Masking Finetuning

## — Supplementary Material

Yuyang Xue, Junyu Yan, Raman Dutt, Fasih Haider, Jingshuai Liu, Steven McDonagh, and Sotirios A. Tsaftaris


School of Engineering, The University of Edinburgh, Edinburgh, EH9 3FG, UK

**Table S1.** The specifications for all datasets we used. Note that PAD-UFES-20 is extremely imbalanced in target. “\” means the dataset does not contain this attribute.

Dataset	Images	Benign	Malignant	Target	Skin Tone	Gender
ISIC-training	46,938	43,142	3,796	1:11	6:1	1:1
ISIC-external	9,925	8,818	1,107	1:8	4:1	1:1
Fitzpatrick	16,576	14,419	2,157	1:7	2:1	\
PAD-UFES-20	1,495	1,443	52	1:28	19:1	1:1
DDI	657	486	171	1:3	2:1	\
Atlas	1,012	760	252	1:3	\	1:1

**Table S2.** Searching range for hyperparameters.

Hyperparameters	Search Range
impair learning rate	$[1 \times 10^{-2}, 1 \times 10^{-3}, 3 \times 10^{-3}, 5 \times 10^{-4}]$
repair learning rate	$[1 \times 10^{-2}, 1 \times 10^{-3}, 3 \times 10^{-3}, 5 \times 10^{-4}]$
$K$	[50, 70, 90, 95]
$\beta$	[0.01, 0.05, 0.02, 0.01]



**Fig. S1.** Mask generation starts by assessing each parameter for bias importance (a) and loss importance (b) (shown in heatmap, the higher the brighter). The ideal weights are those with high bias relevance and low loss impact. The final mask (c) is created using Eq. ??, with blue dotted boxes in examples showing how these two are combined.

**Table S3.** Comparison of post-process bias-removal methods on skin tone and gender attributions. Note that “\” means the method cannot classify the test data very well.

Methods	Architecture	Skin Tone						Gender								
		Fizpatrick		PAD-UFES-20		DDI		Atlas		PAD-UFES-20						
		ACC $\uparrow$	Eodds $\downarrow$	ACC $\uparrow$	Eodds $\downarrow$	ACC $\uparrow$	Eodds $\downarrow$	ACC $\uparrow$	Eodds $\downarrow$	ACC $\uparrow$	Eodds $\downarrow$	ACC $\uparrow$	Eodds $\downarrow$			
Baseline	ResNet-18	0.715	0.554	0.1414	0.685	0.522	0.0404	0.713	0.598	0.1441	0.688	0.414	0.0263	0.685	0.522	0.0610
	ResNet-34	0.709	0.541	0.1098	0.805	0.519	0.0467	0.715	0.579	0.0912	0.531	0.507	0.0288	0.805	0.519	0.0741
	ResNet-50	0.719	0.521	0.1745	0.858	0.506	0.0412	0.744	0.646	0.1288	0.574	0.614	0.0348	0.858	0.506	0.0167
	ResNet-101	0.717	0.520	0.1396	0.880	0.522	0.0543	0.698	0.569	0.1397	0.729	0.662	0.0674	0.880	0.522	0.1027
Full FT	ResNet-18	0.816	0.566	0.1021	0.892	0.504	0.0572	0.755	0.741	0.0422	0.751	0.661	0.0662	0.892	0.504	0.0166
	ResNet-34	0.749	0.536	0.1284	0.935	0.482	0.0910	0.733	0.577	0.0247	0.770	0.685	0.0271	0.935	0.482	0.0053
	ResNet-50	0.703	0.522	0.1510	0.725	0.501	0.0321	0.727	0.618	0.2059	0.690	0.655	0.0303	0.725	0.501	0.0730
	ResNet-101	0.737	0.541	0.1346	0.863	0.524	0.1239	0.726	0.614	0.1060	0.682	0.667	0.0098	0.863	0.524	0.0105
FairPrune	ResNet-18	0.575	0.488	0.1322	0.932	0.510	0.0273	0.739	0.620	0.0056	\	\	\	0.860	0.515	0.0688
	ResNet-34	0.762	0.468	0.0800	0.851	0.493	0.0365	0.743	0.705	0.0148	\	\	\	0.312	0.500	0.0569
	ResNet-50	0.763	0.547	0.0120	0.698	0.492	0.1237	0.737	0.536	0.0024	\	\	\	0.634	0.499	0.1456
	ResNet-101	0.688	0.518	0.2027	0.338	0.505	0.1027	0.727	0.476	0.0291	\	\	\	0.246	0.489	0.1649
LLFT	ResNet-18	0.708	0.545	0.1680	0.446	0.486	0.1946	0.699	0.583	0.1094	0.749	0.375	0.0013	0.964	0.482	0.0007
	ResNet-34	0.661	0.534	0.1072	0.763	0.501	0.0482	0.675	0.574	0.1713	0.746	0.374	0.0025	0.964	0.482	0.0006
	ResNet-50	0.579	0.513	0.1446	0.694	0.492	0.0726	0.646	0.561	0.1197	0.724	0.389	0.0041	0.964	0.482	0.0007
	ResNet-101	0.764	0.533	0.1199	0.733	0.505	0.0348	0.698	0.597	0.1633	0.750	0.625	0.0030	0.951	0.482	0.0035
Diff-Bias	ResNet-18	0.830	0.559	0.0764	0.834	0.506	0.0251	0.740	0.870	0.0041	0.768	0.683	0.0131	0.868	0.511	0.0657
	ResNet-34	0.770	0.542	0.0897	0.919	0.502	0.0498	0.740	0.634	0.0393	0.749	0.656	0.0354	0.918	0.495	0.0022
	ResNet-50	0.716	0.507	0.1144	0.884	0.501	0.0599	0.737	0.629	0.0741	0.769	0.687	0.0383	0.878	0.504	0.0461
	ResNet-101	0.763	0.533	0.0958	0.892	0.523	0.0791	0.705	0.574	0.1627	0.752	0.678	0.0223	0.887	0.531	0.0261
BMFT (Ours)	ResNet-18	0.862	0.640	0.0121	0.944	0.523	0.0364	0.749	0.873	0.0040	0.756	0.806	0.0086	0.862	0.507	0.0049
	ResNet-34	0.865	0.705	0.0033	0.924	0.505	0.0241	0.742	0.655	0.0333	0.778	0.729	0.0020	0.913	0.512	0.0468
	ResNet-50	0.865	0.637	0.0055	0.902	0.507	0.0657	0.759	0.736	0.0754	0.752	0.876	0.0037	0.893	0.512	0.0365
	ResNet-101	0.829	0.560	0.0666	0.912	0.526	0.0621	0.742	0.871	0.0082	0.755	0.752	0.0010	0.878	0.518	0.0404