

Gradient Attention Balance Network: Mitigating Face Recognition Racial Bias via Gradient Attention

Linzhi Huang¹Mei Wang¹Jiahao Liang¹Weihong Deng¹Hongzhi Shi^{1,2}Dongchao Wen^{1,2*}Yingjie Zhang^{1,2}Jian Zhao^{1,2}¹Inspur Electronic Information Industry Co., Ltd²Shandong Massive Information Technology Research Institute

{huanglinzhi_1, wangmei_1, goodlplus, dengweihong2020}@foxmail.com

{shihzh, wendongchao, zhangyj-s, zhao-jian}@inspur.com

Abstract

Although face recognition has made impressive progress in recent years, we ignore the racial bias of the recognition system when we pursue a high level of accuracy. Previous work found that for different races, face recognition networks focus on different facial regions, and the sensitive regions of darker-skinned people are much smaller. Based on this discovery, we propose a new de-bias method based on gradient attention, called Gradient Attention Balance Network (GABN). Specifically, we use the gradient attention map (GAM) of the face recognition network to track the sensitive facial regions and make the GAMs of different races tend to be consistent through adversarial learning. This method mitigates the bias by making the network focus on similar facial regions. In addition, we also use masks to erase the Top- N sensitive facial regions, forcing the network to allocate its attention to a larger facial region. This method expands the sensitive region of darker-skinned people and further reduces the gap between GAM of darker-skinned people and GAM of Caucasians. Extensive experiments show that GABN successfully mitigates racial bias in face recognition and learns more balanced performance for people of different races.

1. Introduction

With the development of deep learning [14, 19, 27, 30–32, 37], the accuracy of the face recognition (FR) algorithm has reached a high level [5, 15, 35]. However, with the emergence of more and more applications based on face recogni-

*Corresponding author.

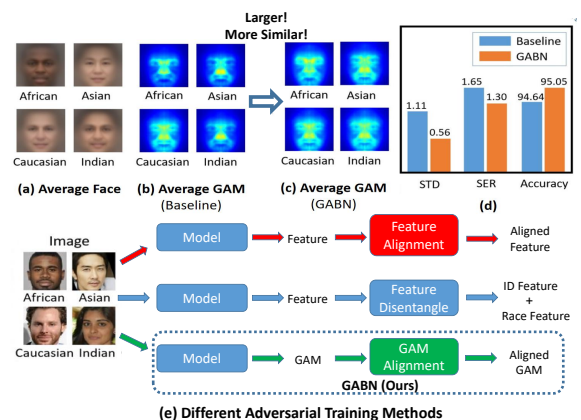


Figure 1. (a) Average faces of four races in the RFW dataset. (b) Average GAM generated by Baseline (ArcFace-34). (c) Average GAM generated by our Gradient Attention Balance Network (GABN). (d) Results on RFW dataset: We calculated standard deviation (STD) and skewed error ratio (SER) with the average accuracy of 4 races. STD and SER are used as the fairness criterion (specifically introduced in the SEC 4.1). Compared with the baseline of ArcFace-34 backbone, GABN improves the overall accuracy of face recognition and mitigates bias. GABN reduces STD from 1.11 to 0.56 and SER from 1.65 to 1.30. (e) Different adversarial training methods.

tion, its potential racial discrimination is attracting people’s attention [1, 3]. For example, Amazon’s recognition tool incorrectly matched the photos of 28 U.S. congressmen with the photos of criminals (error rate up to 39%). This unfair prediction leads to unfair treatment of different population groups. Therefore, it is very important to solve the problem

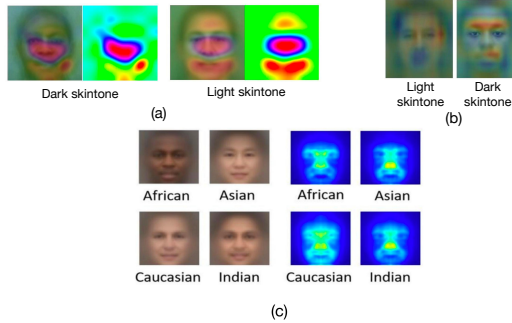


Figure 2. (a) This image is from [8]. Average Gram-CAM for dark skintone and light skintone generated by Crystalface. (b) This image is from [23]. Average Gram-CAM for dark skintone and light skintone generated by ArcFace. (c) Ours. Average GAM generated by ArcFace.

of racial bias in the face recognition system.

Previous studies [6, 13, 17, 25, 33, 43, 44] have shown that the bias of face recognition comes from data and algorithms. Since the commonly used public large-scale datasets [15, 18, 21, 24, 26], such as CASIA-WebFace [42], VGGFace2 [4], and MSCeleb-1M [11] are collected from the Internet, they inevitably encode race, gender, and age biases. Some studies [22, 36, 38, 39] propose new and more balanced datasets. However, some studies [7–9, 36, 38, 39] have proved that the racial bias of the models trained with such balanced datasets cannot be eliminated completely. Therefore, we also need to mitigate this bias through algorithms. Wang et al. [36, 39] propose the adaptive margin loss functions based on reinforcement-learning and meta-learning. Both methods use additional small fair datasets to guide the network to select the best margin for each race. Gong et al. [9] propose a network that generates disentangled representations. This method mitigates bias by disentangling identity features and other demographic attribute features. They [10] also propose a method based on adaptive convolution kernel (GAC). GAC mitigates racial bias by introducing an additional group adaptive classifier. It is worth noting that the previous racial de-bias algorithms based on adversarial training mainly have two ways: 1) The method of feature alignment. 2) The method of feature disentangle [7, 9], as shown in Fig. 1 (e). Both of these methods delete racial features, resulting in the accuracy of face recognition lower than baseline. However, our method is to align the GAM of different races through adversarial training. In addition, our method not only improves fairness, but also improves model accuracy.

Some recent studies [8, 23] have used Grad-cam [28] for visualization and found that face recognition networks attend to different facial regions, depending on the race, as shown in Fig. 2. For darker-skinned people, the face recog-




Figure 3. Comparison of confidence curves of four races (BUPT-Balancedface dataset). Confidence refers to the prediction probability value of the network for the category of the input image. (a) Confidence curve of Baseline (ArcFace-34). There is a gap in the confidence of different races. (b) Confidence curve of GABN. The gap between different races has narrowed.

nition network will pay more attention to a local region, such as the region near the nose for Africans. For Caucasians, the face recognition network focuses on a larger region. Inspired by Wang et al. [34], we designed an image-level map, called gradient attention map (GAM), as shown in Fig. 2. Grad-cam calculates the map based on the last convolution layer of the network and prefers to highlight the decision-making region. However, GAM can locate sensitive region more accurately. We used GAM for visualization and found a similar phenomenon: the sensitive region of darker-skinned people is smaller than that of Caucasians, and the sensitive region of darker-skinned people tends to be a local region, as shown in Fig. 1 (b). Some studies [39] verified that darker-skinned people are more susceptible to noise and image quality than Caucasians. We think this is because face recognition networks pay more attention to local regions for darker-skinned people. This inconsistent sensitive region leads to racial bias in face recognition networks. In addition, we also found that during training, the

network’s confidence of different races is inconsistent, and this inconsistency will not completely disappear with the training, as shown in Fig. 3 (a). Confidence refers to the prediction probability value of the network for the category of the input image. We believe that this inconsistency in confidence also introduces racial bias. To solve the above problems, we propose a new de-bias algorithm based on gradient attention, called Gradient Attention Balance Network (GABN). To solve the problem of inconsistent attention region, we propose two methods: 1) GAM consistency training (GAM-CT): As shown in Fig. 4, we use the face recognition network ($Network_{ID}$) as a generator to generate the GAM of each image. In addition, we train a GAM race classification network ($Network_{GAM-CT}$) and use it as a discriminator. Then, we improve the consistency of GAM of different races through adversarial learning. 2) GAM guided sensitive facial region erasure (GAM-SFRE): We try to use a special method to expand the sensitive region of the network on darker-skinned people and reduce the gap between darker-skinned people and Caucasians, as shown in Fig. 4. Under the guidance of GAM, we erase the Top-N sensitive regions in the image by generating masks, forcing the network to allocate attention to a larger facial region. In addition, to solve the problem of inconsistent confidence, we try to add confidence balance loss to the objective function. After adding the confidence balance loss, the confidence gap between the 4 races is reduced, as shown in Fig. 3(b). It is worth noting that our confidence balance loss is to obtain balanced confidence, not to maximize confidence.

A brief comparison between the baseline (ArcFace-34 [5]) and our method (GABN) is sketched in Fig. 1. Comparing (b) and (c) in Fig. 1, we can see that our method enables the face recognition network to pay more consistent attention to different races, and successfully enables the face recognition network to allocate attention to a larger facial region. As shown in Fig. 1(d), our method not only enhances the overall accuracy but also greatly mitigates racial bias (STD, SER).

The main contributions of this work are as follows:

- We propose a new face recognition de-bias algorithm called gradient attention balance network (GABN). To our best knowledge, this is the first work to mitigate the racial bias of face recognition by improving the consistency of gradient attention map (GAM) of different races through adversarial learning (GAM-CT), which provides a new perspective to improve face recognition fairness. Unlike previous work based on feature disentangle or feature alignment, our method is based on GAM alignment.
- Face recognition networks pay more attention to local regions for darker-skinned people. To further reduce

the performance gap between darker-skinned people and Caucasians, we use GAM guided sensitive facial region erasure (GAM-SFRE) to force the attention of the network to be allocated to a larger facial region. This method reduces the overfitting of darker-skinned people and improves the fairness of the network.

- Extensive experiments on the BUPT-Globalface [36], BUPT-Balancedface [36], and RFW [38] datasets show that our gradient attention balance network (GABN) shows more balanced performance.

2. Related Work

Bias Mitigation in Face Recognition. To mitigate the statistical bias of face recognition models, a variety of de-bias techniques [7–10, 36, 38, 39] have emerged. Wang et al. [38] proposed an information maximization adaptation network to mitigate the bias through domain adaptation. After that, Wang et al. [36, 39] proposed adaptive balance networks based on reinforcement-learning and meta-learning. They introduced a small and balanced verification set to guide the selection of margin parameters for each race during training. Xu et al. [41] proposed a new loss function to mitigate the bias by improving the consistency of false positives. Gong et al. [9] proposed using adversarial learning to disentangle the demographic attribute features and identity features. Recently, they [10] proposed a group adaptive classifier based on estimating demographic attributes to further improve the method. Unlike previous work, our GABN aims to make the network have similar gradient attention to people of different races. Mitigate racial prejudice by improving the consistency of gradient attention.

De-bias Algorithm Based on Adversarial Learning (Face Recognition). Adversarial learning is widely used in the de-bias task of face recognition. Gong et al. [9] proposed to use the disentangled representation for face recognition. They used an adversarial network with four classifiers, one for face recognition and the other three for demographic attributes. Adversarial learning is used to confuse and delete demographic information in features, so as to mitigate potential bias. Recent work [7] proposed a new de-bias algorithm based on adversarial learning. They also disentangle the features of identity and demographic attributes through adversarial learning. The difference is that their work is based on the pre-trained model. Some work [2, 20, 40] has also used similar methods to eliminate sensitive information. Because demographic attribute features are an indispensable part of face identity, this method reduces the overall accuracy of face recognition. Different from previous work [7, 9], our GABN uses adversarial learning to make the network have similar gradient attention to people of different races. GABN not only mitigates the bias but also

improves the overall accuracy.

3. Methodology

3.1. Problem Statement

We visualized the images of 4 races to get the corresponding GAM. It is found that the sensitive region of darker-skinned people is small and that of Caucasians is large, as shown in Fig. 1(b). Fig. 1(b) is the average GAM of 4 races in the RFW dataset. We believe that this inconsistency in sensitive region introduces racial bias. In addition, when we trained the network, we found that the network’s confidence in the 4 racial groups was inconsistent, as shown in Fig. 3(a). Therefore, we propose two methods to solve the problem of inconsistent attention region: 1) We use GAM consistency training (GAM-CT) to improve the similarity of GAM. 2) We use GAM guided sensitive facial region erasure (GAM-SFRE) to force the GAM of darker-skinned people to tend to a larger facial region. We also propose to use confidence balance loss to solve the problem of inconsistent confidence. Specific details will be introduced in this section. The detail procedure of GABN is summarized in Algorithm 1.

3.2. Gradient Attention Map (GAM)

To know which facial regions are mainly concerned by the face recognition network ($Network_{ID}$), we use GAM to accurately locate the sensitive facial regions. As shown in Fig. 2, GAM is the image-level map. Each value of GAM accurately represents the sensitivity of $Network_{ID}$ to the corresponding pixels in the face image. As shown in Fig. 4(steps 1 and 2), forward propagation is used to obtain the probability value P_i of each class (identity). Then, calculate the relative magnitude T_{GAM} of the maximum probability value P_{max} predicted by $Network_{ID}$ and the probability value P_i of other classes. The formula is as follows:

$$T_{GAM} = \frac{\sum_{i=1}^N (P_{max} - P_i)}{N} \quad (1)$$

where N represents the number of identity categories. Because any perturbation will affect P_i and P_{max} , we use T_{GAM} as the objective function for backward propagation to obtain the gradient value and the GAM. The formula is as follows:

$$GAM_{x,y} = \max(\nabla_{x,y}(abs(T_{GAM}))) \quad (2)$$

where x and y represent pixel coordinates, ∇ represents the gradient value, the function $max(\cdot)$ calculates the maximum value along channel axis and the function $abs(\cdot)$ obtains the absolute value of each pixel. It is worth noting that Grad-cam [28] is based on the last convolution layer of the network, which tends to highlight the decision-making

region. GAM generates a map at the image-level to locate sensitive region. The positioning accuracy of GAM will be higher, but the attention distribution will be sparse, as shown in Fig. 2.

3.3. GAM Consistency Training (GAM-CT)

We use adversarial learning to improve the consistency of GAM among different races, as shown in Fig. 4(steps 3 and 4). Unlike previous works based on feature disentangle or feature alignment, our method is based on GAM alignment. Firstly, a GAM race classification network ($Network_{GAM-CT}$) is trained to classify the races of GAM. The input of $Network_{GAM-CT}$ is the GAM of each image, and the output is the race corresponding to each GAM. $Network_{GAM-CT}$ consists of resnet-18 [12] and fully connected layer. The number of input channels of the first convolution layer of resnet-18 [12] is changed to 1 (the dimension of GAM is $W \times H \times 1$). The loss function L_{cls} of $Network_{GAM-CT}$ is the standard cross entropy loss function:

$$L_{cls} = -\frac{1}{N} \cdot \sum_{i=1}^N \log \frac{e^{W_{y_i}^T x_i + b_{y_i}}}{\sum_{j=1}^n e^{W_j^T x_i + b_j}} \quad (3)$$

where x_i denotes the deep feature of the i_{th} sample, belonging to the $y_{i_{th}}$ class. W_j denotes the j_{th} column of the weight W and b_j is the bias term. The batch size and the class number are N and n . Then, we use $Network_{GAM-CT}$ as discriminator and face recognition network ($Network_{ID}$) as generator for adversarial training. $Network_{GAM-CT}$ provides an adversarial loss L_{adv} for face recognition network ($Network_{ID}$). L_{adv} will make $Network_{ID}$ generate GAM in which $Network_{GAM-CT}$ cannot distinguish race. It is well known that a uniform distribution has the highest entropy and presents the most randomness. If an optimal classifier operating on GAM always produces a posterior probability of $\frac{1}{N}$ for all categories in the racial attribute, it means that the GAM generated by the $Network_{ID}$ has a consistent sensitive region of different races. N represents the number of race. Therefore, we define the adversarial loss L_{adv} as:

$$L_{adv} = -\sum_{n=1}^N \left(\frac{1}{N} \cdot \left(\log \frac{e^{Network_{GAM-CT}(GAM)_n}}{\sum_{j=1}^N e^{Network_{GAM-CT}(GAM)_j}} \right) \right) \quad (4)$$

where N is the number of races. GAM-CT enables the face recognition network to pay more consistent attention to different races. See supplementary materials for specific algorithm process.

3.4. GAM Guided Sensitive Facial Region Erasure (GAM-SFRE)

To further improve fairness, we use GAM guided sensitive facial region erasure (GAM-SFRE) to expand the facial




Figure 4. An illustration of our method. 1) $Network_{ID}$ forward propagation to get T_{GAM} . 2) $Network_{ID}$ backward propagation to obtain gradient attention map (GAM). 3) We input GAM into GAM race classification network ($Network_{GAM-CT}$) for training, so that the network can distinguish GAM of different races. 4) We use $Network_{GAM-CT}$ as a discriminator. It provides an adversarial loss L_{adv} for face recognition network ($Network_{ID}$). L_{adv} will make $Network_{ID}$ generate GAM in which $Network_{GAM-CT}$ cannot distinguish race. 5) According to GAM, calculate the Top-N sensitive regions, and use mask to erase the image. Then we input the erased image into the $Network_{ID}$ for training. 6) We use L_{conf} , L_{ID} and L_{adv} for backward propagation and update the parameters of $Network_{ID}$.

region of darker-skinned people, as shown in Fig. 4 (step 5). First, we search each pixel of GAM and sort them in descending order according to the attention value of each pixel. Then, we select N^{mask} pixels with the highest attention value as the central point C_n^{mask} of the mask. The mask we use is a rectangular block with its height less than h^{mask} and its width less than w^{mask} . Finally, we overlay the mask to the corresponding position of the original image, and then input the image into the network for training. See supplementary materials for specific algorithm process.

3.5. Objective Function

As described in SEC 3.1, we believe that this inconsistency in confidence is also one of the factors leading to racial bias. Therefore, based on the objective function of ArcFace [5], we add a confidence balance loss to reduce this inconsistency. The details are as follows: 1) The normalized features and weights are inputted into the loss function of ArcFace (additive angular margin penalty) to get the probability value P_i of each class, and get the maximum probability value P_{max} from it, which is called confidence. 2) Calculate the number of samples with P_{max} less than the threshold $T_{confidence}$ in the current batch ($n_{confidence}$). 3)

Get confidence balanced loss function:

$$L_{conf} = \frac{n_{batch} - n_{confidence}}{n_{batch}} \quad (5)$$

where n_{batch} represents the number of samples in current batch. 4) Finally, we combine L_{conf} , L_{ID} and adversarial loss L_{adv} in GAM-CT to obtain the final objective function L_{Final} . Here, L_{conf} and L_{adv} are taken as the penalty coefficient.

$$L_{ID} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s(\cos(\theta_{y_i} + m))}}{e^{s(\cos(\theta_{y_i} + m))} + \sum_{j=1}^n e^{s(\cos(\theta_j))}} \quad (6)$$

$$K = L_{conf} + L_{adv} \quad (7)$$

$$L_{Final} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s(\cos(\theta_{y_i} + m))}}{e^{s(\cos(\theta_{y_i} + m))} + \sum_{j=1}^n e^{s(\cos(\theta_j + K))}} \quad (8)$$

We update the parameters of face recognition network ($Network_{ID}$) by minimizing L_{Final} .

Table 1. Verification performance (%) of protocol on RFW with SOTA methods ([BUPT-Globalface, ResNet34, Loss-ArcFace]).

Methods	Caucasian	Indian	Asian	African	Avg	Fairness	
						STD	SER
Baseline (ArcFace) [5]	97.37	95.68	94.55	93.87	95.37	1.53	2.33
M-RBN [36]	97.03	95.58	94.40	95.18	95.55	1.10	1.89
RL-RBN [36]	97.08	95.63	95.57	94.87	95.79	0.93	1.76
MBN [39]	96.87	96.20	95.63	95.00	95.93	0.80	1.60
GABN (Ours)	97.21	96.06	95.73	95.51	96.13	0.75	1.60

Table 2. Verification performance (%) of protocol on RFW with SOTA methods ([BUPT-Balancedface, ResNet34, Loss-ArcFace]).

Methods	Caucasian	Indian	Asian	African	Avg	Fairness	
						STD	SER
Baseline (ArcFace) [5]	96.18	94.67	93.72	93.98	94.64	1.11	1.65
ACNN [16]	96.12	94.55	93.67	94.00	94.58	1.08	1.63
PFE [29]	96.38	94.60	94.27	95.17	95.11	0.93	1.58
DebFace [9]	95.95	94.78	94.33	93.67	94.68	0.96	1.56
GAC [10]	96.02	94.22	94.10	94.12	94.62	0.93	1.48
RL-RBN [36]	96.27	94.68	94.82	95.00	94.64	0.73	1.43
MBN [39]	96.25	95.32	94.85	95.38	95.45	0.58	1.37
GABN (Ours)	95.78	95.21	94.51	94.71	95.05	0.56	1.30

4. Experiments

4.1. Experimental Setting

Dataset: Our bias study uses BUPT-Balancedface [36] and BUPT-Globalface datasets [36] for training and RFW [38] datasets for testing. BUPT-Balancedface consists of faces of 4 races: Caucasians, Asians, Indians, and Africans. This dataset contains 1.3 million images of 28K celebrities, with about 7K identities per race. BUPT-Globalface contains 2 million images of 38K celebrities, whose racial distribution is roughly the same as the real distribution of the world’s population. RFW consists of faces of 4 races: Caucasians, Asians, Indians and Africans. Each subset of RFW contains about 10K images of 3K identities.

Evaluation Protocol. The evaluation protocol we used are consistent with the previous work [9, 10, 36, 39]. We calculated standard deviation (STD) and skewed error ratio (SER) with the average accuracy of 4 races. STD and SER are used as the fairness criterion. STD reflects the amount of dispersion of accuracies of different races. SER is computed by the ratio of the highest error rate to the lowest error rate among 4 races. $SER = \frac{\max Error_g}{\min Error_g}$, where g in *Caucasian, Indian, Asian, African*.

Implementation details: For preprocessing, we use 5 facial landmarks for similarity transformation, then crop and resize the images to 112×112 pixels. Each pixel ([0, 255]) in RGB images are normalized by subtracting 127.5

and then being divided by 128. Consistent with the previous work [9, 10, 36, 39, 41], the CNN architecture we use is ResNet-34 [12]. The CNN model is trained on 3 GPUs (NVIDIA GeForce 1080Ti). The models are trained by the SGD algorithm, with momentum 0.9 and weight decay $5e-4$. We set the scale and margin as $s = 64$ and $m = 0.35$ according to the common settings in the previous work [41]. The confidence threshold $T_{confidence}$ is set to 70%. On BUPT-Balancedface, the batch size is set to 200, the learning rate starts from 0.1 and is divided by 10 at 9, 15, 20 epochs. The training process is finished at 30 epochs. On BUPT-Globalface, the batch size is set to 250, we divide the learning rate at 9, 15, 20 epochs and finish at 40 epochs. Our GAM race classifier uses ResNet-18 [12], in which the number of input channels of the first convolution layer of ResNet-18 is changed to 1. The GAM race classifier is trained by the SGD algorithm, with momentum 0.9 and weight decay $5e-4$. Its learning rate starts from 0.1 and is divided by 10 at 9, 15, 20 epochs. It should be noted that the face recognition network and GAM race classifier are trained in turn.

4.2. Comparisons with SOTA Methods

Consistent with the previous work [9, 10, 16, 29, 36, 39], we train a ResNet-34 model on BUPT-Globalface, use the loss function of ArcFace [5], and then report the results on RFW protocol, as shown in Table 1. Compared with the SOTA results, GABN reduces STD to 0.75 and SER to 1.60.

Algorithm 1 GABN (GAM-CT + GAM-SFRE)

Require:

Original facial images I_O .
Gradient attention map GAM .
Max number of epochs E .
Number of network updates per epoch B .
Mask count N .

Ensure:

Face recognition network $Network_{ID}$.
GAM race classification network $Network_{GAM-CT}$.

For e **in** E **For** b **in** B

FP $Network_{ID}$ to obtain P_{max} and P_i ;
Use P_{max} to get L_{conf} ;
Use formula 1 in the paper to get T_{GAM} ;
BP $Network_{ID}$ to obtain GAM (formula 2);
Get the erased image I_E guided by GAM :
 $n = 0$;
For n **in** N
 $i, j =$ coordinate of the n th largest value;
 $H_t = \text{Rand}(1, H^{mask})$;
 $W_t = \text{Rand}(1, W^{mask})$;
 Use a rectangle with H_t and W_t to erase I_O ;
 Training $Network_{GAM-CT}$;
 BP to optimize $Network_{GAM-CT}$ (formula 3);
 Training $Network_{ID}$;
 FP $Network_{GAM-CT}$ to get L_{adv} (formula 4);
 Use L_{id}, L_{conf} and L_{adv} to get L_{final} ;
 Use L_{final} to optimize $Network_{ID}$.

We also used the same method to conduct experiments on BUPT-Balancedface, as shown in Table 2. Compared with the SOTA results, GABN reduces STD to 0.56. Compared with DebFace [9], which is also based on adversarial learning, our method not only obtains lower STD and SER but also improves the overall accuracy of the model. The above results show that our method has achieved competitive performance on both balanced and unbalanced datasets.

4.3. Ablation Study

To investigate the efficacy of our method, we conduct a lot of ablation studies. The setup of ablation study is consistent with that described in SEC 4.1. We use the framework of ArcFace-34, then train Baseline and GABN on BUPT-Balancedface dataset, and test on RFW dataset.

Effect of the confidence balanced loss. We use $T_{confidence}$ to control the confidence threshold. Table 3 reports the results of different $T_{confidence}$. $T_{confidence}$ is not the higher the better. The best results are obtained when we set $T_{confidence}$ to 70%. This is in line with our expectations, because our confidence balanced loss is to obtain a more balanced confidence, not to maximize the confidence.




Figure 5. (a) Comparison of average GAM of 4 races of Baseline and GABN. We convert GAM from gray to color. Red represents greater attention.

We use this setting in later experiment.

Effect of the GAM-SFRE. To study the effectiveness of GAM-SFRE module, We compare our GAM-SFRE with random erasure (RE). We keep the settings of RE and GAM-SFRE consistent and conduct experiments. The result of our method is better than random erasure, as shown in Table 4. We explain the reasons as follows: Although random erasure increases the diversity of samples, most of the random erasure regions are insensitive, which makes the performance improvement of the network limited. However, our method is based on GAM guidance for erasure. Each erased region is a sensitive region, forcing the network to focus on a larger facial region, so as to effectively reduce overfitting. We also compared RE and GAM-SFRE through visualization, as shown in SEC 4.4.

Effect of the GAM-CT. To study the effectiveness of GAM-CT module, we use GAM-CT for ablation study, as shown in Table. 4. When we used GAM-CT, the STD decreased to 0.63.

We can get better results than Baseline when we use GAM-CT or GAM-SFRE. This shows that these components can effectively improve the fairness of the model. When we use these modules at the same time, we can get the performance of SOTA. We summarize the best results of each module in Table 4.

4.4. Visualization and Analysis

To prove that our method improves the consistency of GAM of different races, we visualize GAM. We use the method introduced in SEC 3 to obtain the GAM of each image in the RFW dataset and calculate the average GAM of 4 races. As shown in Fig. 5, in the average GAM of the 4 races obtained by Baseline, the sensitive region of Caucasians is large and that of darker-skinned people is small. However, the sensitive region of GABN has higher consistency and larger sensitive region. This shows that our GABN has successfully reduced the attention gap between different races. In addition, we also visualize the GAM of random erasure (RE) and GAM-SFRE, as shown in Fig. 6. RE can expand sensitive region, but the effect is very small. GAM-SFRE can expand the region to the whole face, and the effect is

Table 3. Verification performance (%) of different threshold of confidence balanced loss ($T_{confidence}$).

Methods	$T_{confidence}$	Caucasian	Indian	Asian	African	Avg	Fairness	
							STD	SER
Baseline	-	96.18	94.67	93.72	93.98	94.64	1.11	1.65
+ confidence balanced loss	60%	95.21	94.65	93.73	94.05	94.44	0.66	1.31
	70%	96.23	95.41	94.66	95.15	95.36	0.65	1.42
	75%	96.13	95.06	94.45	94.83	95.12	0.72	1.43
	80%	95.61	94.36	93.58	93.88	94.36	0.89	1.46

Table 4. Ablation of different components on the RFW protocol. **+Random erase** is the result of using random erase (RE).

Methods	Caucasian	Indian	Asian	African	Avg	Fairness	
						STD	SER
Baseline (ArcFace-34)	96.18	94.67	93.72	93.98	94.64	1.11	1.65
+ Random Erase	96.20	95.40	94.57	94.57	95.18	0.78	1.43
+ confidence balance loss	96.23	95.41	94.66	95.15	95.36	0.65	1.42
GAM-CT	95.93	95.20	94.41	94.90	95.11	0.63	1.37
GAM-SFRE	96.16	95.43	94.63	95.08	95.32	0.64	1.39
GABN (GAM)	95.78	95.21	94.51	94.71	95.05	0.56	1.30




Figure 6. GAM comparison diagram of random erase (RE) and GAM-SFRE.

remarkable.

4.5. Limitation Analysis

As shown in Fig. 4, GABN needs 2 backward propagation to train 1 step. The first backward propagation generates GAM, and the second backward propagation updates the parameters. In addition, we need to train a GAM race classification network ($Network_{GAM-CT}$). We train on 3 GPUs (NVIDIA Geforce 1080TI), which takes about 2.2 times longer than baseline (ArcFace-34). However, we think the training time is within the acceptable range.

5. Conclusion

In this paper, we propose a de-bias network (GABN) based on gradient attention and adversarial learning to reduce racial bias. GABN improves the consistency of gradient attention map (GAM) of different races through ad-

versarial learning. In addition, GABN also uses GAM guided sensitive region erasure to expand the sensitive region of darker-skinned people, further reducing the performance gap between darker-skinned people and Caucasians. Comprehensive experiments show the effectiveness of our GABN.

Acknowledgements. This work was supported by China Postdoctoral Science Foundation under Grant 2022M720517.

References

- [1] Mohsan Alvi, Andrew Zisserman, and Christoffer Nellaker. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 1
- [2] Blaž Bortolato, Marija Ivanovska, Peter Rot, Janez Križaj, Philipp Terhörst, Naser Damer, Peter Peer, and Vitomir Štruc. Learning privacy-enhancing face representations through feature disentanglement. In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 495–502. IEEE, 2020. 3
- [3] 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, 2018. 1
- [4] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 67–74. IEEE, 2018. 2

- [5] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [1](#), [3](#), [5](#), [6](#)
- [6] Prithviraj Dhar, Ankan Bansal, Carlos D Castillo, Joshua Gleason, P Jonathon Phillips, and Rama Chellappa. How are attributes expressed in face dcnn? In *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 85–92. IEEE, 2020. [2](#)
- [7] Prithviraj Dhar, Joshua Gleason, Aniket Roy, Carlos D Castillo, and Rama Chellappa. Pass: Protected attribute suppression system for mitigating bias in face recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15087–15096, 2021. [2](#), [3](#)
- [8] Prithviraj Dhar, Joshua Gleason, Aniket Roy, Carlos D Castillo, P Jonathon Phillips, and Rama Chellappa. Distill and de-bias: Mitigating bias in face recognition using knowledge distillation. *arXiv preprint arXiv:2112.09786*, 2021. [2](#), [3](#)
- [9] Sixue Gong, Xiaoming Liu, and Anil K Jain. Jointly debiasing face recognition and demographic attribute estimation. In *European conference on computer vision*, pages 330–347. Springer, 2020. [2](#), [3](#), [6](#), [7](#)
- [10] Sixue Gong, Xiaoming Liu, and Anil K Jain. Mitigating face recognition bias via group adaptive classifier. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3414–3424, 2021. [2](#), [3](#), [6](#)
- [11] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European conference on computer vision*, pages 87–102. Springer, 2016. [2](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#), [6](#)
- [13] Matthew Q Hill, Connor J Parde, Carlos D Castillo, Y Ivette Colon, Rajeep Ranjan, Jun-Cheng Chen, Volker Blanz, and Alice J O’Toole. Deep convolutional neural networks in the face of caricature. *Nature Machine Intelligence*, 1(11):522–529, 2019. [2](#)
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. [1](#)
- [15] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in ‘Real-Life’ Images: detection, alignment, and recognition*, 2008. [1](#), [2](#)
- [16] Di Kang, Debarun Dhar, and Antoni Chan. Incorporating side information by adaptive convolution. *Advances in Neural Information Processing Systems*, 30, 2017. [6](#)
- [17] Brendan F Klare, Mark J Burge, Joshua C Klontz, Richard W Vorder Bruegge, and Anil K Jain. Face recognition performance: Role of demographic information. *IEEE Transactions on Information Forensics and Security*, 7(6):1789–1801, 2012. [2](#)
- [18] Brendan F Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, and Anil K Jain. Pushing the frontiers of unconstrained face detection and recognition: Iarpa janus benchmark a. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1931–1939, 2015. [2](#)
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. [1](#)
- [20] Ang Li, Jiayi Guo, Huanrui Yang, and Yiran Chen. Deepobfuscator: Adversarial training framework for privacy-preserving image classification. *arXiv preprint arXiv:1909.04126*, 2(3), 2019. [3](#)
- [21] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *2018 international conference on biometrics (ICB)*, pages 158–165. IEEE, 2018. [2](#)
- [22] Michele Merler, Nalini Ratha, Rogerio S Feris, and John R Smith. Diversity in faces. *arXiv preprint arXiv:1901.10436*, 2019. [2](#)
- [23] Shruti Nagpal, Maneet Singh, Richa Singh, and Mayank Vatsa. Deep learning for face recognition: Pride or prejudiced? *arXiv preprint arXiv:1904.01219*, 2019. [2](#)
- [24] P Jonathon Phillips, J Ross Beveridge, Bruce A Draper, Geof Givens, Alice J O’Toole, David Bolme, Joseph Dunlop, Yui Man Lui, Hassan Sahibzada, and Samuel Weimer. The good, the bad, and the ugly face challenge problem. *Image and Vision Computing*, 30(3):177–185, 2012. [2](#)
- [25] P Jonathon Phillips, Patrick Grother, Ross Micheals, Duane M Blackburn, Elham Tabassi, and Mike Bone. Face recognition vendor test 2002. In *2003 IEEE International SOI Conference. Proceedings (Cat. No. 03CH37443)*, page 44. IEEE, 2003. [2](#)
- [26] Joseph P Robinson, Gennady Livitz, Yann Henon, Can Qin, Yun Fu, and Samson Timoner. Face recognition: too bias, or not too bias? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–1, 2020. [2](#)
- [27] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. [1](#)
- [28] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. [2](#), [4](#)
- [29] Yichun Shi and Anil K Jain. Probabilistic face embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6902–6911, 2019. [6](#)
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [1](#)

- [31] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-verification. *Advances in neural information processing systems*, 27, 2014. [1](#)
- [32] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. [1](#)
- [33] Philipp Terhörst, Daniel Fährmann, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Beyond identity: What information is stored in biometric face templates? In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2020. [2](#)
- [34] Chengrui Wang and Weihong Deng. Representative forgery mining for fake face detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14923–14932, 2021. [2](#)
- [35] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5265–5274, 2018. [1](#)
- [36] Mei Wang and Weihong Deng. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9319–9328, 2020. [2](#), [3](#), [6](#)
- [37] Mei Wang and Weihong Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021. [1](#)
- [38] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 692–702, 2019. [2](#), [3](#), [6](#)
- [39] Mei Wang, Yaobin Zhang, and Weihong Deng. Meta balanced network for fair face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2021. [2](#), [3](#), [6](#)
- [40] Zhenyu Wu, Zhangyang Wang, Zhaowen Wang, and Hailin Jin. Towards privacy-preserving visual recognition via adversarial training: A pilot study. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 606–624, 2018. [3](#)
- [41] Xingkun Xu, Yuge Huang, Pengcheng Shen, Shaoxin Li, Jilin Li, Feiyue Huang, Yong Li, and Zhen Cui. Consistent instance false positive improves fairness in face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 578–586, 2021. [3](#), [6](#)
- [42] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014. [2](#)
- [43] Xi Yin, Xiang Yu, Kihyuk Sohn, Xiaoming Liu, and Manmohan Chandraker. Feature transfer learning for face recognition with under-represented data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5704–5713, 2019. [2](#)
- [44] Xiao Zhang, Zhiyuan Fang, Yandong Wen, Zhifeng Li, and Yu Qiao. Range loss for deep face recognition with long-tailed training data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5409–5418, 2017. [2](#)