

Mitigating Bias in Face Recognition using Skewness-Aware Reinforcement Learning

Mei Wang, Weihong Deng*

Beijing University of Posts and Telecommunications

{wangmei1, whdeng}@bupt.edu.cn

Abstract

Racial equality is an important theme of international human rights law, but it has been largely obscured when the overall face recognition accuracy is pursued blindly. More facts indicate racial bias indeed degrades the fairness of recognition system and the error rates on non-Caucasians are usually much higher than Caucasians. To encourage fairness, we introduce the idea of adaptive margin to learn balanced performance for different races based on large margin losses. A reinforcement learning based race balance network (RL-RBN) is proposed. We formulate the process of finding the optimal margins for non-Caucasians as a Markov decision process and employ deep Q-learning to learn policies for an agent to select appropriate margin by approximating the Q-value function. Guided by the agent, the skewness of feature scatter between races can be reduced. Besides, we provide two ethnicity aware training datasets, called BUPT-Globalface and BUPT-Balancedface dataset, which can be utilized to study racial bias from both data and algorithm aspects. Extensive experiments on RFW database show that RL-RBN successfully mitigates racial bias and learns more balanced performance.

1. Introduction

Recently, with the emergence of deep convolutional neural networks (CNN) [26, 43, 46, 19, 20], the performance of face recognition (FR) [49, 45, 42] is dramatically boosted. However, as its wider and wider application, its potential for unfairness is raising alarm [8, 5, 1, 2]. For instance, according to [14], a year-long research investigation across 100 police departments revealed that African-American individuals are more likely to be stopped by law enforcement. As stated in the Universal Declaration Human Rights [7], all are equal before the law and are entitled without any discrimination to equal protection. Obviously, the development and deployment of fair and unbiased FR systems is crucial to prevent any unintended side effects and to ensure

the long-term acceptance of these FR algorithms. Previous studies [50, 54] have shown that racial bias comes on both data and algorithm aspects. Unfortunately, there are still no sufficient research efforts on the fairness of face recognition algorithms [38, 39, 13], as well as building balanced databases [50], in the literature.

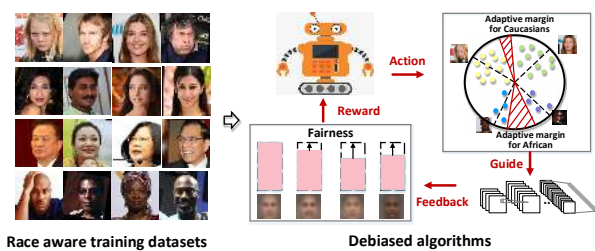


Figure 1. We provide two ethnicity aware training datasets and a debiased algorithm to reduce bias from data and algorithm aspects.

A major driver of bias in face recognition, as well as other AI tasks, is the training data. Deep face recognition networks are often trained on large-scale training datasets, such as CASIA-WebFace [52], VGGFace2 [10] and MS-Celeb-1M [17], which are typically constructed by scraping websites like Google Images. Such data collecting methods can unintentionally produce data that encode gender, ethnic and culture biases. Thus, social awareness must be brought to the building of datasets for training. In this work, we take steps to ensure such datasets are diverse and do not under represent particular groups by constructing two new training datasets, i.e. BUPT-Globalface and BUPT-Balancedface dataset. One is built up according to the population ratio of ethnicity in the world, and the other strictly balances the number of samples in ethnicity.

Another source of bias can be traced to the algorithms themselves. The state-of-the-art (SOTA) face recognition methods, such as SpheroFace [30], Cosface [48] and Arcface [12], apply a margin between classes to maximize overall prediction accuracy for the training data. If a specific group of faces appears more frequently than others in the training data, the method will optimize for those individuals because

this boosts the accuracy on the test datasets with the same bias, e.g. LFW [21, 27] and IJB-C [31]. Further, our experiments show that even with balanced training set, the feature separability of non-Caucasians is inferior to that of Caucasians. To address this problem, the algorithms must trade-off the specific requirements of margins of various groups of people, and produce more equitable recognition performance.

Since racial bias is a complex problem caused by many latent factors, including but not limited to the number, we apply reinforcement learning (RL) [34] to learn an AutoML tool on setting dynamic margins for different races. RL endows agents with the ability to perform experiments to better understand biased learning process, enabling them to learn high-level causal relationships leading to dataset-independent policy on adaptive margins. In this paper, a reinforcement learning based race-balance network (RL-RBN) is proposed. First, we use deep Q-learning to train an agent to generate adaptive margin policy for non-Caucasians through maximizing the expected sum of rewards. The rewards are designed according to the skewness of intra/inter-class distances between races. Then, we train the balanced models guided by this adaptive margin policy. Finally, RL-RBN balances the inter-class and intra-class distance among different races, and thus achieves balanced generalization ability. Besides racial bias showed in our experiments, our method can also apply to remove other demographic bias, e.g. gender and age.

Our contributions can be summarized into three aspects. 1) Two ethnicity aware training datasets are constructed and released¹ for the study on reducing racial bias. 2) A new debiased algorithm, RL-RBN, is proposed to learn adaptive margins to mitigate bias between different races. RL technique is successfully applied to learn an adaptive margin policy. 3) Extensive experiments on ethnicity aware training datasets and RFW [50] shows the effectiveness of our RL-RBN. Combining balanced training and our debiased algorithm, we obtain the fairest performance across races.

2. Related work

2.1. Racial bias in face recognition

Some studies [38, 16, 13, 39, 25] have uncovered that non-deep face recognition algorithms inherit racial bias from human and perform unequally on different races. The 2002 NIST Face Recognition Vendor Test (FRVT) is believed to be the first study that showed that non-deep FR algorithms suffer from racial bias [38]. Phillips et al. [39] utilized the images of the Face Recognition Vendor Test 2006 (FRVT 2006) to conduct cross training and matching on White and Asian races, and suggested that training and testing on different races results in severe performance drop.

¹<http://www.whdeng.cn/RFW/index.html>

Klare et al. [25] collected mug shot face images of White, Black and Hispanic from the Pinellas County Sheriff's Office (PCSO) and concluded that the Black cohorts are more difficult to recognize for all matchers. However, few efforts were made to study racial bias in deep face recognition. Recently, Wang et al. [50] contributed a test dataset called Racial Faces in-the-Wild (RFW) database, on which they validated the racial bias of four commercial APIs and four SOTA face recognition algorithms, and presented the solution using deep unsupervised domain adaptation to alleviate this bias. But there is still vacancy in training datasets which can be used to study racial bias.

2.2. Debiased algorithms

In many computer vision applications, there are some works that seek to introduce fairness into networks and mitigate data bias. These are respectively classified as unbalanced-training [41, 9, 36, 53], attribute suppression [5, 33, 37, 32] and domain adaptation [23, 22, 50, 44]. By learning the underlying latent variables in an entirely unsupervised manner, Debiasing Variational Autoencoder (DB-VAE) [6] re-weighted the importance of certain data points while training. Calmon et al. [9] transformed the given dataset into a fair dataset by constraining the conditional probability of network prediction to be similar for any two values of demographic information. SensitiveNets [35] proposed to introduce sensitive information into triplet loss. They minimized the sensitive information, while maintaining distances between positive and negative embeddings.

2.3. Deep reinforcement learning

Mimicking humans' decision making process, RL aims to enable the agent to decide the behavior from its experiences using a Markov decision process (MDP) [28]. Mnih et al. [34] combined RL learning with CNN and bridged the divide between high-dimensional sensory inputs and actions, resulting in human-level performance in Atari Games. Recently, RL has been successfully applied in compute vision. Rao et al. [40] used RL to discard the confounding frames and found the focuses of attentions for video recognition. Haque et al. [18] utilized RL to identify small, discriminative regions indicative of human identity in person identification. Liu et al. [29] used RL to learn optimal margin to improve overall FR performance. In this paper, we apply deep RL to address race balance problem in FR.

3. Ethnicity aware training datasets

A major driver of bias in face recognition is the training data. Frequently, some race groups are over-represented and others are under-represented. For example, East Asia and India together contribute just 8% of commonly-used training datasets, even though these countries represent 44% of the world's population. In order to remove this source of

bias and represent people of different regions equally, we construct two ethnicity aware training datasets, i.e. BUPT-Globalface and BUPT-Balancedface dataset. The identities in these two datasets are grouped into 4 categories, i.e. Caucasian, Indian, Asian and African, according to their races. As shown in Fig. 2, BUPT-Globalface contains 2M images from 38K celebrities in total and its racial distribution is approximately the same as real distribution of world’s population. BUPT-Balancedface dataset contains 1.3M images from 28K celebrities and is approximately race-balanced with 7K identities per race.

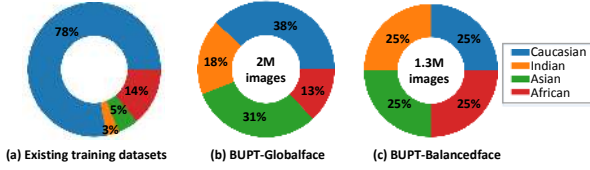


Figure 2. The percentage of different races in commonly-used training datasets, BUPT-Globalface and BUPT-Balancedface.

Similar to RFW [50], we select images of different races from MS-Celeb-1M [3] with help of the “Nationality” attribute of FreeBase celebrities [15] and Face++ API. However, due to unbalanced distribution of MS-Celeb-1M [3], we can only obtain 2K Indians and 5K Asians which are not enough to construct a large scale dataset. As we know, MS-Celeb-1M only selected the top 100K entities from one-million FreeBase celebrity list [15]. Therefore, we download the remaining images of Asians and Indians by Google according to FreeBase celebrity list, and then clean them both automatically and manually in the similar way as other FR training datasets, such as VGGface2 [10] and Megaface [24]. After obtaining enough images, we select images to construct our ethnicity aware training datasets, and remove their overlapping subjects in RFW [50].

4. Our method

4.1. Investigation and observation

To make learned features potentially separable and enhance the discrimination power, some methods proposed to incorporate a margin between classes based on Softmax. For example, Arcface [12] used an additive angular margin m :

$$L_{arc} = -\frac{1}{N} \sum_{j=1}^N \log \frac{e^{s_c(\cos(\theta_{y(j)j} + m))}}{e^{s_c(\cos(\theta_{y(j)j} + m))} + \sum_{i=1, i \neq y(j)}^n e^{s_c \cos \theta_{ij}}} \quad (1)$$

and Cosface [48] used an additive cosine margin m :

$$L_{cos} = -\frac{1}{N} \sum_{j=1}^N \log \frac{e^{s_c(\cos(\theta_{y(j)j}) - m)}}{e^{s_c(\cos(\theta_{y(j)j}) - m)} + \sum_{i=1, i \neq y(j)}^n e^{s_c \cos \theta_{ij}}} \quad (2)$$

where θ_{ij} is the angle between the weight W_i and the feature x_j . $x_j \in \mathbb{R}^d$ denotes the deep feature of the j -th sample, belonging to the $y^{(j)}$ -th class, and $W_i \in \mathbb{R}^d$ denotes the i -th column of the weight $W \in \mathbb{R}^{d \times n}$. N is the batch size, n is the number of classes, and s_c is the scale factor. Although large margin losses successfully improve feature discrimination, and get better performance on a series of face recognition benchmarks, Wang et al. [50] experimentally proved that they still fail to obtain balanced representations on different races under a uniform margin.

To get a better understanding, we train two ResNet-34 [19] models with the guidance of Arcface [12] and Softmax loss on the CASIA-Webface [52], and give the detailed angle statistics of different races in Table 1. The intra-class and inter-class angle are computed on set-1 and RFW [50]. The set-1 contains 500 identities per race randomly selected from our BUPT-Globalface dataset. Intra-class angle refers to the mean of angles between feature and the feature centre, which can be formulated as: $\Theta_{intra} = \frac{1}{N_g} \sum_{i=1}^{N_g} \frac{1}{|\mathcal{I}_i|} \sum_{x_j \in \mathcal{I}_i} \theta_{x_j, c_i}$, where N_g is the number of identities belonged to one race group, \mathcal{I}_i is the set of all images in i -th identity, and c_i is the feature centre of i -th identity computed by the mean vector of embeddings. Inter-class angle refers to the mean of minimum angles between embedding feature centres, which can be formulated as: $\Theta_{inter} = \frac{1}{N_g} \sum_{i=1}^{N_g} \min_{k=1: N_g, k \neq i} \theta_{c_k, c_i}$. In Table 1, we find that non-Caucasians can not obtain as good intra-class compactness and inter-class discrepancy as Caucasians do, especially the inter-class angle. That is to say, the performance of non-Caucasians is inferior to that of Caucasians on test set, even if a uniform margin is performed on different races when training.

loss functions		Caucasian	Indian	Asian	African
Softmax	Intra1	35.55	36.00	38.83	37.29
	Inter1	62.67	50.89	46.17	54.26
	Intra2	33.81	31.49	32.09	32.39
	Inter2	59.09	50.15	46.41	48.42
Arcface [12]	Intra1	36.70	40.00	42.78	41.50
	Inter1	67.72	59.73	55.34	62.90
	Intra2	36.22	34.87	36.15	36.18
	Inter2	65.76	59.72	56.46	58.54

Table 1. The angle statistics of different races ([CASIA [52], ResNet34, loss*]). “Intra1” and “Intra2” refer to intra-class angle on set-1 and RFW, respectively. “Inter1” and “Inter2” refer to inter-class angle on set-1 and RFW, respectively.

Considering of poorer generalization ability of non-Caucasians, we should take special care of these difficult colored faces, and thus prefer an appropriate margin with them to improve their generalization ability. Therefore, we introduce the idea of adaptive margin into race balance problem. The margin of Caucasians remains unchanged, while optimal margins are selected adaptively for each col-

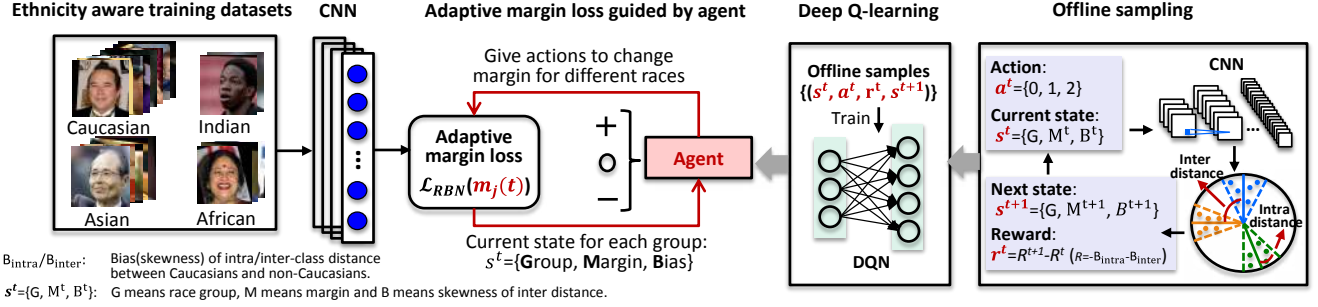


Figure 3. An illustration of our method. **Offline sampling:** We vary margin for each race group to collect some training samples , i.e. (s^t, a^t, r^t, s^{t+1}) , before training DQN. Details of (s^t, a^t, r^t, s^{t+1}) refer to Section 4.2. **Deep Q-learning network:** With these samples, DQN is trained to approximate the Q-value function, and the reward is determined by the skewness of inter/intra-class distance between races. Then, adaptive margin policy for agent can be generated according to Q-value. **Adaptive margin:** We train a race balanced network with a fixed margin for Caucasians and adaptive margins for each colored-face which changes at each training step guided by agent.

ored race in order to minimize the skewness of angles between races and learn balanced performance for different races. We replace the fixed margin m in Eqn. 1 and Eqn. 2 by a race-related and training step related parameter $m_j(t)$, where t represents the stage of the training. In fact, different races have different demands for the margins, and the demands may change during the training. The proposed adaptive margin loss function can be formulated as follows:

$$L_{RBN} = -\frac{1}{N} \sum_{j=1}^N \log \frac{e^{s_c(\cos(\theta_{y(j)} + \alpha_j(t)))}}{e^{s_c(\cos(\theta_{y(j)} + \alpha_j(t)))} + \sum_{i=1, i \neq y(j)}^n e^{s_c \cos \theta_{ij}}}$$

where, $\alpha_j(t) = \begin{cases} m, & \text{if } j \in \text{Caucasian} \\ m_j(t), & \text{otherwise} \end{cases}$ (3)

The similar modification can be made for Cosface [48]. So the key problem is to find an optimal adaptive margin policy $m_j(t)$ for each non-Caucasian group to minimize the skewness of angles between Caucasians and non-Caucasians. We make Caucasians as the benchmark (anchor) here because white subjects are overwhelmingly dominant in numbers in existing FR datasets. According to different racial distribution, our RL-RBN can change the benchmark and replace Caucasians by the most representative ethnic group.

4.2. Adaptive margin policy learning

In our method showed in Fig. 3, the problem of finding $m_j(t)$ can be formulate as a Markov Decision Process (MDP) [28], based on which we use deep Q-learning to adjust the margin at each iteration. Deep Q-learning aims to enable the agent to decide the behavior from its experiences using a MDP. At each time step $t \in [1, T]$, the agent takes an action a^t from action space \mathcal{A} according to the Q-value $Q(s^t, a)$ estimated by the deep Q-learning network (DQN) with the state s^t as input. The environment gives a reward $r(s^t, a^t)$ for this action and then the agent updates its states with s^{t+1} . According to this new state, the agent will go in-

to next time step and choose a new action. The goal of deep Q-learning is to train an agent with policy π to maximize the expected sum of rewards.

State. The state s of the MDP consists of three separate parts $\{G, M, B_{inter}\}$. $G = \{0, 1, 2\}$ represents the race group, i.e. Indian (group 0), Asian (group 1) and African (group 2). M is equivalent of the adaptive margin m . B_{inter} means the bias (skewness) of inter-class distance between g -th race group and Caucasians, which can be formulated as follows:

$$B_{inter} = |d_{inter}^g - d_{inter}^{Cau}|$$

$$\text{where, } d_{inter}^g = \frac{1}{N_g} \sum_{i=1}^{N_g} \max_{k=1: N_g, k \neq i} \cos(c_k, c_i) \quad (4)$$

where d_{inter}^g and d_{inter}^{Cau} is the inter-class distance of g -th race group and that of Caucasian, respectively. N_g is the number of identities belonged to g -th race group and c_i is the feature centre of i -th identity computed by the mean vector of embeddings. $\cos(\cdot, \cdot)$ is the cosine distance function. d_{inter}^{Cau} can be computed by the same way as d_{inter}^g . We suppose that different races have different demands for the margins, and the demands may change according to their B_{inter} . When the bias (skewness) is larger, the race group may need larger margin to improve its generalization ability, and vice versa. Therefore, we take both G and B_{inter} into consideration when designing state s . Each action will be chosen based on race group and the current skewness. Moreover, in order to make the space of states discrete, we map M and B_{inter} to discrete spaces \mathcal{M} and \mathcal{B} , where $\mathcal{M} = \{m_1, m_2, \dots, m_{n_M}\}$ and $\mathcal{B} = \{b_1, b_2, \dots, b_{n_B}\}$.

Action. The action $\mathcal{A} = \{0, 1, 2\}$ is the adjustment of margin. We define 3 types of action as ‘staying the same’ (action 0), ‘shifting to larger value’ (action 1) and ‘shifting to smaller value’ (action 2), and shifting step is set to be a constant ϵ . The optimal action taken by the agent at time step t is formulated by $a^t = \underset{a}{argmax} Q(s^t, a)$, where the

Q-value $Q(s^t, a)$ is the accumulated rewards of taking the action a . For example, at time step t , the agent chooses to take the action ‘1’ according to Q-value and state $s^t = \{0, m_2, b_1\}$, then the margin of Indian will update to $m_3 = m_2 + \epsilon$.

Reward. The reward, as a function $r(s^t, a^t)$, reflects how good the action taken by the agent is with regard to the state s^t . Since we suppose non-Caucasians should have the same generalization ability as Caucasians and the skewness between them should be minimized, we use the skewness of inter/intra-class distance between them to design the reward. The bias (skewness) of intra-class distance between Caucasians and g -th group can be formulated as:

$$B_{intra} = |d_{intra}^g - d_{intra}^{Cau}|$$

$$\text{where, } d_{intra}^g = \frac{1}{N_g} \sum_{i=1}^{N_g} \frac{1}{|\mathcal{I}_i|} \sum_{x_j \in \mathcal{I}_i} \cos(x_j, c_i) \quad (5)$$

where d_{intra}^g and d_{intra}^{Cau} is the intra-class distance of g -th race group and that of Caucasian, respectively. N_g is the number of identities belonged to g -th race group. \mathcal{I}_i is the set of all images in i -th identity, and c_i is the feature centre of i -th identity computed by the mean vector of embeddings. And the bias (skewness) of inter-class distance can be computed by Eqn. 4. When the agent takes the action a^t to adjust the margin of g -th race group, the reward of the action a^t is computed by:

$$r(s^t, a^t) = R^{t+1} - R^t$$

$$\text{where, } R = -(B_{inter} + B_{intra}) \quad (6)$$

Objective function. We choose to use deep Q-Learning [34, 51] to learn an optimal policy for agent. A two-layer fully connected network with a further hidden layer of 10 units is utilized to estimate the Q function. Each fully connected layer is followed by a ReLU activation function. The deep Q-learning network takes the state as input and produces the Q-value of all possible actions. We update the network by the minimizing the following loss function:

$$L_q = \mathbb{E}_{s^t, a^t} \left[(y^t - Q(s^t, a^t))^2 \right]$$

$$\text{where, } y^t = \mathbb{E}_{s^{t+1}} \left[r^t + \gamma \max_{a^{t+1}} Q(s^{t+1}, a^{t+1} | s^t, a^t) \right] \quad (7)$$

where $y^t - Q(s^t, a^t)$ is the temporal difference error. y^t is the target value of $Q(s^t, a^t)$. r^t is the reward of taking the action a^t , computed by Eq.6, and $\gamma \max_{a^{t+1}} Q(s^{t+1}, a^{t+1} | s^t, a^t)$ is the future reward estimated by the current deep Q-learning network with s^{t+1} .

To train DQN, we collect some offline samples in advance, i.e. $\{(s^t, a^t, r^t, s^{t+1})\}$, as input to feed network.

These offline samples are generated by a sample CNN as shown in Algorithm 1. For each non-Caucasian group, we manually adjust the margin by actions a^t and train the current sample network for one epoch with new margin. After one-epoch training, we compute the intra-class and inter-class distance of this race group and obtain the next state s^{t+1} and reward r^t . We keep doing this until all states have been traversed. Then, we train DQN by these collected samples. After that, the adaptive margin policy can be generated according to the output of trained DQN, i.e. Q-value.

Algorithm 1 Offline sampling.

Input:

The unbalanced data with four race groups.

Output:

The samples, i.e. $\{(s^t, a^t, r^t, s^{t+1})\}$, which are used for training deep Q-learning network.

- 1: **for** g **in** all groups **do**
 - 2: **for** a **in** all actions **do**
 - 3: Compute B_{inter}^t of g -th group by Eqn. 4, and obtain current state $s^t = \{G, M^t, B_{inter}^t\}$.
 - 4: Take the action a to adjust margin for g -th group.
 - 5: Train the sample network for an epoch with the updated margin M^{t+1} of g -th group; while margins of other groups remain unchanged.
 - 6: Compute B_{inter}^{t+1} of g -th group according to Eqn. 4, and obtain updated state $s^{t+1} = \{G, M^{t+1}, B_{inter}^{t+1}\}$.
 - 7: Compute reward r^t according to Eqn. 6.
 - 8: Collect the sample $\{(s^t, a^t, r^t, s^{t+1})\}$.
 - 9: **end for**
 - 10: **if** s^{t+1} is a new state that hasn't appeared **then**
 - 11: **Go** to step 2.
 - 12: **end if**
 - 13: **end for**
-

Recognition network. We utilize the policy to guide the training process of our recognition network. At each time step t , s^t is computed by recognition network according to Eqn. 4 and sent to the agent, then the trained agent will take action a^t to adjust the margins for each non-Caucasian group. Then, we update the margin of non-Caucasians in Eqn. 3, and use it to optimize the recognition network for one epoch. We keep doing this until the end of training.

5. Experiments

5.1. Experimental settings

Datasets. Because existing datasets are not race-aware except RFW [50], we use our BUPT-Globalface and BUPT-Balancedface datasets to train our models and use RFW [50] to fairly measure performance of different races. RFW [50] consists of four testing subsets, namely Caucasian, Asian,

Indian and African. Each subset contains about 10K images of 3K individuals for face verification. Moreover, in order to evaluate the generalization ability accurately, we additionally construct a validation set to compute intra-class and inter-class distance in RL instead of directly using training sets. The validation set contains 500 identities per race, and has no overlapping subjects with BUPT-Globalface, BUPT-Balancedface and RFW dataset [50].


Evaluation Protocol. Verification performance is measured by accuracy. We utilize average accuracy of four races as metric to evaluate total performance of deep models. And standard deviation (STD) and skewed error ratio (SER) are used as the fairness criterion. Standard deviation reflects the amount of dispersion of accuracies of different races. Error skewness is computed by the ratio of the highest error rate to the lowest error rate among different races. It can be formulated as $SER = \frac{\max_g Error_g}{\min_g Error_g}$ where $g \in \{\text{Caucasian, Indian, Asian, African}\}$ means race group.

Implementation details. For preprocessing, we use five facial landmarks for similarity transformation, then crop and resize the faces to 112×112 . Each pixel $([0, 255])$ in RGB images is normalized by subtracting 127.5 and then being divided by 128. MxNet [11] is utilized to implement adaptive margin loss and TensorFlow [4] is for Deep Q-learning. We use ResNet34 [19] for training recognition CNN and sampling CNN. For the case of training recognition CNN, the learning rate is started from 0.1 and decreased with a factor of 10 when errors plateau. The batch size is set as 256 for the small training dataset and 200 for the large one. We set momentum as 0.9 and weight decay as $5e - 4$. The architecture of DQN is mentioned in Section 4.2. The AdamOptimizer is used to optimize the whole network. The learning rate is set to be $1e - 4$ and the discount factor γ is set to be 0.99.

First, we collect training samples $\{(s^t, a^t, r^t, s^{t+1})\}$ for deep Q-learning. In RL-RBN(soft), the margin of Caucasians is set to be 0 similar to N-Softmax [47], and the margin of non-Caucasians varies from 0 to 0.6 in steps of 0.2 in which an additive angular margin is used similar to Arcface [12]; In RL-RBN(cos) and RL-RBN(arc), the margin of Caucasians is set to be 0.15 and 0.3, and the margin of non-Caucasians varies from 0.15 to 0.45 in steps of 0.1 and varies from 0.3 to 0.6 in steps of 0.1, respectively. After that, we feed the samples $\{(s^t, a^t, r^t, s^{t+1})\}$ to train deep Q-learning network and then generate an adaptive margin policy. Finally, the recognition CNN is trained with the guidance of this policy.

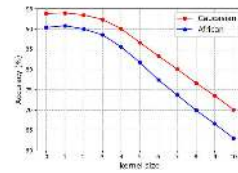
5.2. Cause of racial bias

Some papers [50, 54, 25] verified that non-Caucasians still perform poorly than Caucasians even with balanced training and faces of colored skin are inherently difficult to

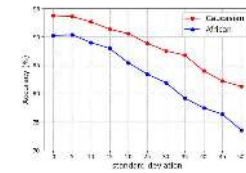


(a) Gaussian blur (b) Gaussian noise

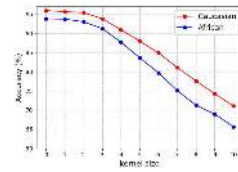
Figure 4. Examples of images degraded by applying (a) gaussian blur or (b) gaussian noise. The first columns are original images.




(a) Gaussian blur (Softmax)



(b) Gaussian noise (Softmax)



(c) Gaussian blur (Arcface)



(d) Gaussian noise (Arcface)

Figure 5. The performance of Caucasians and Africans tested on blurred and noisy RFW [50].

recognize for existing algorithms. In order to go deep into this phenomenon, we degrade images of RFW [50] by blur and noise, and observe the influence of these image degradations on the performance of Caucasians and Africans. To apply Gaussian blur, the sigma value of the Gaussian filter is 1.5 and kernel size varied from 1 to 10 in steps of one (Fig. 4). And we add noise to images by using Gaussian noise with zero mean and varying standard deviation from 5 to 50 in steps of 5. We train two ResNet-34 models with guidance of Softmax and Arcface loss [12] using BUPT-Balancedface dataset, and test them on the blurred and noisy RFW [50]. The accuracies of Caucasians and Africans are given in Fig. 5. We can see that the performance gap between Africans and Caucasians still exists even with balanced training. And both Caucasians and Africans are found to be sensitive to image blur and noise. More importantly, when blur and noise level increases, the performance gap between Caucasians and Africans widens. Therefore, we conclude that colored faces are more susceptible to noise and image quality than Caucasians. This may be one of the reasons why non-Caucasians are more difficult to recognize.

5.3. Experiments of our method

Results on simulated dataset. As we know, data bias in training set severely affects fairness of algorithms. In order to validate the effectiveness of our RL-RBN, we train our al-

Train↓	Test→ Method↓	Caucasian	Indian	Asian	African	Avg	Fairness	
							STD	SER
4 : 2 : 2 : 2	N-Softmax [47]	89.67	87.97	84.68	84.17	86.62	2.64	1.53
	RL-RBN(soft)	91.35	90.77	89.87	90.13	90.53	0.66	1.17
5 : $\frac{5}{3}$: $\frac{5}{3}$: $\frac{5}{3}$	N-Softmax [47]	89.88	88.52	85.13	83.42	86.74	2.98	1.64
	RL-RBN(soft)	90.33	90.23	88.97	89.37	89.73	0.67	1.22
6 : $\frac{4}{3}$: $\frac{4}{3}$: $\frac{4}{3}$	N-Softmax [47]	90.43	88.32	84.75	83.32	86.70	3.26	1.74
	RL-RBN(soft)	90.17	90.02	87.67	88.27	89.03	1.25	1.25
7 : 1 : 1 : 1	N-Softmax [47]	90.67	87.77	84.37	82.97	86.44	3.46	1.83
	RL-RBN(soft)	90.63	90.73	87.72	87.53	89.15	1.77	1.35

Table 2. Verification accuracy (%) on RFW [50] trained with varying racial distribution. We boldface STD (lower is better) and skewed error ratio (SER) (1 is the best) since this is the important fairness criterion.

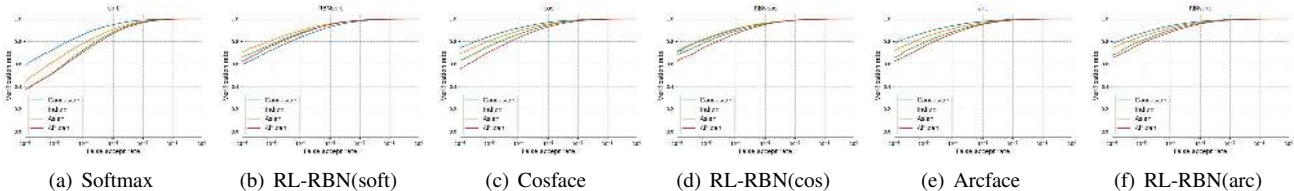


Figure 6. The ROC curves of (a) Softmax, (b) RL-RBN(soft) (c) Cosface [48], (d) RL-RBN(cos), (e) Arcface [12] and (f) RL-RBN(arc) evaluated on all pairs of RFW [50].

gorithms using training set with different racial distribution and evaluate them on RFW [50]. We randomly pick images from our BUPT-Globalface dataset to construct these training sets. Each training set contains 12K celebrities which has the similar scale with CASIA-Webface database [52], and is non-overlapping with RFW dataset. We make the number of Indians, Asians and Africans the same for simplicity, and change the ratio between Caucasians and non-Caucasians, i.e. {4:6}, {5:5}, {6:4}, {7:3}. Norm-Softmax [47], which normalizes weight and feature based on Softmax, is compared with our RL-RBN, as shown in Table 2.

From the results, we can see several important observations. First, it also shows that racial bias indeed exists in existing algorithms. For example, when the racial distribution is 4:2:2:2, the accuracy of Norm-Softmax reaches 89.69% on Caucasian testing subset, but its accuracy dramatically decreases to 84.17% on African subset. Second, the results quantitatively verify our thought that the accuracy of each race is positively correlated with its number in training set. For example, increasing ratio of Caucasians (from 2/5 to 7/10) increases its accuracy from 89.67% to 90.67% in Norm-Softmax. Also, with the change of distribution, we observe a decrease in fairness between races. Third, after adopting adaptive margin loss guided with RL, our RL-RBN(soft) significantly obtains more balanced performance than Norm-Softmax on different races. When racial distribution becomes more uneven, i.e. 7:1:1:1, our method can still perform better and decrease the SER from 1.83 to 1.35.

Results on BUPT-Globalface dataset. Training on BUPT-Globalface, we compare our RL-RBN with Softmax,

Methods	Caucasian	Indian	Asian	African	Avg	Fairness	
						STD	SER
Triplet [42]	95.80	92.77	91.03	90.47	92.52	2.40	2.27
Softmax	95.62	91.97	90.85	89.98	92.10	2.48	2.29
M-RBN(soft)	93.50	94.50	90.06	93.43	92.83	1.90	1.78
RL-RBN(soft)	94.53	95.03	94.20	94.05	94.45	0.44	1.20
Cosface [48]	96.63	94.68	93.50	92.17	94.25	1.90	2.33
M-RBN(cos)	96.15	95.73	93.43	94.76	95.02	1.21	1.70
RL-RBN(cos)	96.03	95.15	94.58	94.27	95.01	0.77	1.45
Arcface [12]	97.37	95.68	94.55	93.87	95.37	1.53	2.33
M-RBN(arc)	97.03	95.58	94.40	95.18	95.55	1.10	1.89
RL-RBN(arc)	97.08	95.63	95.57	94.87	95.79	0.93	1.76

Table 3. Verification accuracy (%) of our policy on RFW [50] ([BUPT-Globalface, ResNet34, loss*]). M-RBN is the method using different fixed margins for different races inversely proportional to their number.

Cosface [48] and Arcface [12]. The scaling parameter is set as 60 and the margin parameters are set as 0.2 and 0.3 for Cosface [48] and Arcface [12], respectively. We show the results in Table 3 and Fig. 6. First, our RL-RBN(soft) obtains more perfect performance than Softmax. It achieves about 2.35% gains for average accuracy, and STD decreases from 2.48 to 0.44. Second, we find that large margin loss, i.e. Cosface [48] and Arcface [12], can alleviate racial bias to some extent through more separate inter-class. However, racial bias cannot be eliminated completely. Third, our RL-RBN(cos) and RL-RBN(arc) can find an optimal margin for each race group and obtain more balanced performance than Cosface and Arcface. It shows the superiority of our algorithm on learning balanced features from a biased dataset.

Results on BUPT-Balancedface dataset. We also com-

Methods	Caucasian	Indian	Asian	African	Avg	Fairness	
						STD	SER
Triplet [42]	94.58	93.17	91.48	91.60	92.71	1.47	1.57
Softmax	94.18	92.82	91.23	91.42	92.41	1.38	1.51
RL-RBN(soft)	94.30	94.13	93.87	94.45	94.28	0.20	1.08
Cosface [48]	95.12	93.93	92.98	92.93	93.74	1.03	1.45
RL-RBN(cos)	95.47	95.15	94.52	95.27	95.10	0.41	1.21
Arcface [12]	96.18	94.67	93.72	93.98	94.64	1.11	1.65
RL-RBN(arc)	96.27	94.68	94.82	95.00	95.19	0.93	1.42

Table 4. Verification accuracy (%) of our policy on RFW [50] (BUPT-Balancedface, ResNet34, loss*).

pare our RL-RBN with Softmax, Cosface [48] and Arcface [12] on BUPT-Balancedface. The results are shown in Table 4. With balanced training, Softmax, Cosface and Arcface indeed obtain more balanced performance compared with trained on biased data. So training equally on all races can help to reduce racial bias to some extent. This conclusion is coincident with [50, 54]. However, even with balanced training, we see that non-Caucasians still perform poorly than Caucasians because some specific races are difficult to recognize. When combining our debiased algorithm and balanced data, we can obtain the fairest performance.

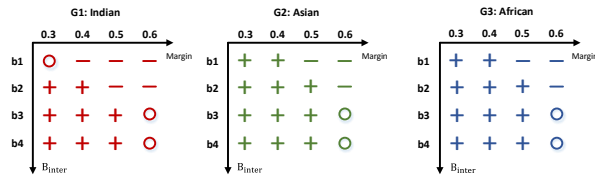


Figure 7. The adaptive margin policy of RL-RBN(arc) from the trained agent. Each symbol, i.e. ‘o’, ‘+’, ‘-’, indicates an action $a = \{0, 1, 2\}$ based on current state $s = \{G, M, B_{inter}\}$. M is mapped to discrete spaces $\{0.3, 0.4, 0.5, 0.6\}$. B_{inter} is mapped to four discrete values $b1 < b2 < b3 < b4$.

Adaptive margin policy. In our RL-RBN, the adaptive margin policy is given by a trained agent which can output an action $a = \{0, 1, 2\}$ with a state $s = \{G, M, B_{inter}\}$ as input. Here, we illustrate the adaptive margin policy of RL-RBN(arc) in Fig. 7. From the policy, we can see several important observations. First, Asian and African group have larger possibility of increasing their margin compared with Indian group. This is consistent with our theoretical analysis that we prefer stricter constraints for races which are more difficult to recognize. Based on our experiments, Asians and Africans have larger domain discrepancy with Caucasians, and perform much worse even with balanced training. Second, it is more likely for the state with larger B_{inter} to increase the margin, and vice versa. Large B_{inter} usually reflects less balanced performance between Caucasians and this race group so that a larger margin is supposed to improve the generalization ability of this group.

Distribution of margins. In Fig. 8, we illustrate the dis-

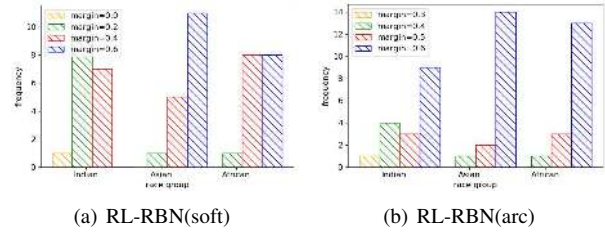


Figure 8. Distribution of margins of non-Caucasians in RL-RBN(soft) and RL-RBN(arc) training on BUPT-Balancedface.

tributions of margins of non-Caucasians in RL-RBN(soft) and RL-RBN(arc) when trained on BUPT-Balancedface dataset. Guided by the agent, the margins of Asians and Africans are indeed larger than those of Indians, especially the Asians who are the most difficult to recognize when trained with Softmax or Arcface. Moreover, RL-RBN(arc) usually selects larger margins for non-Caucasians compared with RL-RBN(soft). This is because the performance of Caucasians is set to be the anchor and the performances of other races are improved to get close to the anchor in our method. In RL-RBN(soft), the margin of Caucasians is small and fixed at 0, overlarge margins of non-Caucasians will lead to out-of-balance performance again. This result proves the robustness and adaptability of our method.

Compared with manual margin. We also compare our method with manual-margin based RBN (M-RBN). The M-RBN simply sets different fixed margins for different races inversely proportional to the number of their samples. From Table 3, we can see that the performance of Asians is always a drag on fairness in M-RBN and our method is superior to M-RBN in fairness. This is because racial bias is a complex problem in which the number is not only fact affecting out-of-balance accuracy. Although the number of Asians is much larger than that of Indians and Africans in BUPT-Globalface dataset, this group still needs a larger margin because it is the most difficult race to recognize even with balanced training.

6. Conclusion

In this paper, we provide two ethnicity aware training datasets, i.e. BUPT-Globalface and BUPT-Balancedface. Then, a reinforcement learning based race-balance network is proposed to alleviate racial bias and learn more balanced features. It introduces the Markov decision process to adaptively find optimal margins for non-Caucasians. The comprehensive experiments prove the effectiveness of RL-RBN.

7. Acknowledgments

This work was partially supported by the National Natural Science Foundation of China under Grant Nos. 61871052 and 61573068.

References

- [1] Are face recognition systems accurate? depends on your race. <https://www.technologyreview.com/s/601786>.
- [2] Facial recognition is accurate, if you're a white guy. <https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html>.
- [3] Ms-celeb-1m challenge 3: Face feature test/trillion pairs. <http://trillionpairs.deepglint.com/>.
- [4] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016.
- [5] Mohsan Alvi, Andrew Zisserman, and Christoffer Nellaker. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. *arXiv preprint arXiv:1809.02169*, 2018.
- [6] Alexander Amini, Ava Soleimany, Wilko Schwarting, Sangeeta Bhatia, and Daniela Rus. Uncovering and mitigating algorithmic bias through learned latent structure. *AIES*, 2019.
- [7] UN General Assembly. Universal declaration of human rights. *UN General Assembly*, 302(2), 1948.
- [8] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81, pages 77–91, 2018.
- [9] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. Optimized pre-processing for discrimination prevention. In *Advances in Neural Information Processing Systems*, pages 3992–4001, 2017.
- [10] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. *arXiv preprint arXiv:1710.08092*, 2017.
- [11] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*, 2015.
- [12] Jiankang Deng, Jia Guo, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *arXiv preprint arXiv:1801.07698*, 2018.
- [13] Nicholas Furl, P Jonathon Phillips, and Alice J O'Toole. Face recognition algorithms and the other-race effect: computational mechanisms for a developmental contact hypothesis. *Cognitive Science*, 26(6):797–815, 2002.
- [14] Clare Garvie. *The perpetual line-up: Unregulated police face recognition in america*. Georgetown Law, Center on Privacy & Technology, 2016.
- [15] Google. Freebase data dumps. <https://developers.google.com/freebase/data>, 2015.
- [16] Patrick J Grother, George W Quinn, and P Jonathon Phillips. Report on the evaluation of 2d still-image face recognition algorithms. *NIST interagency report*, 7709:106, 2010.
- [17] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *ECCV*, pages 87–102. Springer, 2016.
- [18] Albert Haque, Alexandre Alahi, and Li Fei-Fei. Recurrent attention models for depth-based person identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1229–1238, 2016.
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [20] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. *arXiv preprint arXiv:1709.01507*, 2017.
- [21] 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. 2008.
- [22] Meina Kan, Shiguang Shan, and Xilin Chen. Bi-shifting auto-encoder for unsupervised domain adaptation. In *ICCV*, pages 3846–3854, 2015.
- [23] Meina Kan, Juntong Wu, Shiguang Shan, and Xilin Chen. Domain adaptation for face recognition: Targetize source domain bridged by common subspace. *International Journal of Computer Vision*, 109(1-2):94–109, 2014.
- [24] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *CVPR*, pages 4873–4882, 2016.
- [25] 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.
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [27] Erik Learned-Miller, Gary B Huang, Aruni RoyChowdhury, Haoxiang Li, and Gang Hua. Labeled faces in the wild: A survey. In *Advances in face detection and facial image analysis*, pages 189–248. Springer, 2016.
- [28] Michael L Littman. Reinforcement learning improves behaviour from evaluative feedback. *Nature*, 521(7553):445, 2015.
- [29] Bingyu Liu, Weihong Deng, Yaoyao Zhong, Mei Wang, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Fair loss: Margin-aware reinforcement learning for deep face recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10052–10061, 2019.
- [30] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. SpheroFace: Deep hypersphere embedding for face recognition. In *CVPR*, volume 1, 2017.
- [31] 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.
- [32] Vahid Mirjalili, Sebastian Raschka, Anoop Namboodiri, and Arun Ross. Semi-adversarial networks: Convolutional autoencoders for imparting privacy to face images. In *2018 International Conference on Biometrics (ICB)*, pages 82–89. IEEE, 2018.
- [33] Vahid Mirjalili, Sebastian Raschka, and Arun Ross. Gender privacy: An ensemble of semi adversarial networks for confounding arbitrary gender classifiers. *arXiv preprint arXiv:1807.11936*, 2018.
- [34] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedelnd, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [35] Aythami Morales, Julian Fierrez, and Ruben Vera-Rodriguez. Sensitivenets: Learning agnostic representations with application to face recognition. *arXiv preprint arXiv:1902.00334*, 2019.
- [36] Ajinkya More. Survey of resampling techniques for improving classification performance in unbalanced datasets. *arXiv preprint arXiv:1608.06048*, 2016.
- [37] Asem Othman and Arun Ross. Privacy of facial soft biometrics: Suppressing gender but retaining identity. In *European Conference on Computer Vision*, pages 682–696. Springer, 2014.
- [38] P Jonathon Phillips, Patrick Grother, Ross Micheals, Duane M Blackburn, Elham Tabassi, and Mike Bone. Face recognition vendor test 2002. In *Analysis and Modeling of Faces and Gestures, 2003. AMFG 2003. IEEE International Workshop on*, page 44. IEEE, 2003.
- [39] P Jonathon Phillips, Fang Jiang, Abhijit Narvekar, Julianne Ayyad, and Alice J O’Toole. An other-race effect for face recognition algorithms. *ACM Transactions on Applied Perception (TAP)*, 8(2):14, 2011.
- [40] Yongming Rao, Jiwen Lu, and Jie Zhou. Attention-aware deep reinforcement learning for video face recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3931–3940, 2017.
- [41] Prasanna Sattigeri, Samuel C Hoffman, Vijil Chenthamarakshan, and Kush R Varshney. Fairness gan. *arXiv preprint arXiv:1805.09910*, 2018.
- [42] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015.
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [44] Kihyuk Sohn, Wenling Shang, Xiang Yu, and Manmohan Chandraker. Unsupervised domain adaptation for distance metric learning. 2018.
- [45] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-verification. In *NIPS*, pages 1988–1996, 2014.
- [46] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich, et al. Going deeper with convolutions. *Cvpr*, 2015.
- [47] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: 12 hypersphere embedding for face verification. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1041–1049. ACM, 2017.
- [48] 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.
- [49] Mei Wang and Weihong Deng. Deep face recognition: A survey. *arXiv preprint arXiv:1804.06655*, 2018.
- [50] 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 International Conference on Computer Vision*, pages 692–702, 2019.
- [51] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [52] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- [53] Zhi-Hua Zhou and Xu-Ying Liu. Training cost-sensitive neural networks with methods addressing the class imbalance problem. *IEEE Transactions on Knowledge & Data Engineering*, (1):63–77, 2006.
- [54] James Zou and Londa Schiebinger. Ai can be sexist and racist: it’s time to make it fair, 2018.