

## Generalizability and Application of the Skin Reflectance Estimate Based on Dichromatic Separation (SREDS)

Joseph Drahos, Richard Plesh, Keivan Bahmani, Mahesh Banavar, Stephanie Schuckers<sup>1</sup>

**Abstract:** Face recognition (FR) systems have become widely used and readily available in recent history. However, differential performance between certain demographics has been identified within popular FR models. Skin tone differences between demographics can be one of the factors contributing to the differential performance observed in face recognition models. Skin tone metrics provide an alternative to self-reported race labels when such labels are lacking or completely not available e.g. large-scale face recognition datasets. In this work, we provide a further analysis of the generalizability of the Skin Reflectance Estimate based on Dichromatic Separation (SREDS) against other skin tone metrics and provide a use case for substituting race labels for SREDS scores in a privacy-preserving learning solution. Our findings suggest that SREDS consistently creates a skin tone metric with lower variability within each subject and SREDS values can be utilized as an alternative to the self-reported race labels at minimal drop in performance. Finally, we provide a publicly available and open-source implementation of SREDS to help the research community. Available at <https://github.com/JosephDrahos/SREDS>

**Keywords:** Face Recognition, Privacy-Preserving, Feature Unlearning, Skin Reflectance.

### 1 Introduction

Face recognition systems are increasingly used as a form of biometric authentication for many new and existing systems. Research on the differential performance between demographics is an important topic of study to mitigate bias and ensure fairness [dFPM22, Ho22]. Modern facial recognition systems use deep learning pipelines to take an image of a person's face and create a unique template for that person. In such systems, the demographic information of a dataset is needed to assess or mitigate the differential performance of a particular face recognition algorithm. However, many of the large-scale datasets which have been aggregated from public images on the internet and used to train and benchmark face recognition networks lack self-reported race labels. Additionally, the large scale of such datasets makes it impractical and expensive to efficiently label demographics by human annotators. As a result, methods to automatically label a dataset can provide a valuable asset.

Our research focuses on the intersection of privacy preservation, bias mitigation, and skin tone metrics. We present our analysis of the Skin Reflectance Estimate based on Dichromatic Separation (SREDS) skin tone metric from [Ba21]. SREDS is a continuous skin

<sup>1</sup> Department of Electrical and Computer Engineering, 8 Clarkson Ave, Potsdam, NY,  
 {drahosj, pleshro, bahmank, mbanavar, sschucke} @clarkson.edu

This material is based upon work supported by the Center for Identification Technology Research and the National Science Foundation (NSF) under Grant No.1650503.

tone metric that can be used to automatically label skin tones on face datasets. Our goal is to evaluate SREDS' ability to label datasets compared to other skin tone metrics, assess the generalizability of SREDS on unseen data, and demonstrate an application of SREDS using a sensitive information removal approach when race labels are not available.

## 2 Background

### 2.1 Skin Tone Metrics

Previous methods formulated for generating a metric for subject skin tones to more accurately describe skin color are listed as follows: Fitzpatrick Skin Type (FST), Monk Skin Tone (MST) Scale, Individual Typology Angle (ITA), and Relative Skin Reflectance (RSR) [Fi88, Mo19, CCH91, Co19]. FST and its successor MST require a manual calculation from a survey of the subject, while ITA and RSR can be computed automatically via an algorithm. RSR was created to analyze skin tone for a specific dataset by fitting a Principal Component Analysis (PCA) model to the RGB space of the dataset. RSR is not resistant to changes in lighting and is specific to a particular dataset. The need for a skin tone metric that can be computed automatically and is more resistant to changes in lighting prompted the research that led to the Skin Reflectance Estimate based on a Dichromatic Separation (SREDS) [Ba21]. SREDS aims to decompose patches of skin into specular and diffuse components using the dichromatic reflectance model. A Kernel Principal Component Analysis (KPCA) is fit onto the diffuse components extracted from the dataset, resulting in a data-driven skin tone metric.

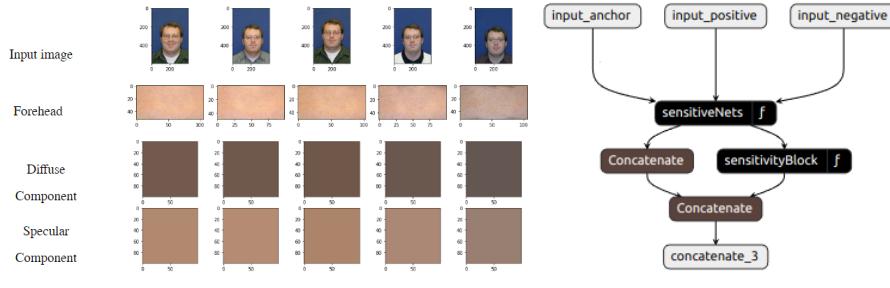
### 2.2 Bias Mitigation

The inclusion of demographic information in a dataset is to observe and attempt to eliminate the differential performance between demographic groups in FR models. Differing methods of bias mitigation have been attempted and documented at the feature, comparison, and post-comparison levels. A method of bias mitigation at the feature level is the triplet mining approach of [Se22] which used a triplet loss for discrimination-aware learning. Closely related triplets are mined based on race information to try and train a new representation that mitigates biased learning within the face embedding space of a pre-trained model. At the comparison level, a learning classifier method reduces ethnic bias by introducing group and individual fairness to the decision process at the cost of matching performance [Te20c]. At the post-comparison level, an unsupervised method of score normalization has been presented to reduce bias between ethnic groups while increasing the performance of the system [Te20b].

### 2.3 Soft Biometric Privacy Preservation

Soft biometric information such as gender, race, age, etc. is stored within the templates created from FR systems and can be extracted without the user's consent [Te20a]. Meth-

ods of privacy preservation have been studied and introduced to protect users' sensitive information. The efforts in [OR15] produced a technique that morphed the input face with another face to mask the soft biometrics while maintaining matching performance. Another technique that added a perturbing element to the initial face image that would mask sensitive information while maintaining performance is [MR17]. Information removal networks attempt to remove sensitive information from the feature embedding space of the FR deep network. These methods require complex loss functions to maintain the performance accuracy of the network while also suppressing the racial information from the learned space, as performed in [Xu18]. A method that combines the methods from [Se22, Xu18] and was used within this research is [Mo20], which attempts to maintain the inter-identity distance using triplet loss and simultaneously unlearn<sup>2</sup> the facial features used to differentiate between demographic classes.



(a) Intra-subject variation between diffuse and specular components extracted using SREDS. (b) SensitiveNets model implementation

Figure 1

### 3 Methodology

#### 3.1 Skin Tone Metrics Assessment

The skin tone metrics outlined in section 2.1 will be used as a baseline to compare the previously developed methods to the performance of the SREDS measure. Individual typology angle (ITA) is a type of colorimetric analysis designed to measure acquired tanning [CCH91]. An RGB image is converted into CIE-Lab space [CF97], as follows: (1) the 'L' component which quantifies luminance, (2) the 'a' component - absence or presence of redness, and (3) the 'b' component - yellowness. Using the 'L' and 'b' components, Pixel-wise ITA value, in degrees, can be estimated throughout an image as:

$$ITA = \frac{\arctan(L - 50)}{b} * \frac{180}{\pi}. \quad (1)$$

To find suitable skin pixels in the image, a landmark extractor based on Dlib is used to detect the forehead, left cheek, and right cheek facial regions [Ki06]. For each facial region,

<sup>2</sup> The term unlearn will be used throughout the paper in the same context as introduced in the literature [Mo20].

ITA is computed over each pixel and smoothed using an averaging filter. The mode from each region’s resulting values is averaged to result in a single skin tone estimate for a face.

Relative Skin Reflectance (RSR) is a process designed to relate the physical properties of the skin to the performance of facial recognition [Co19]. The pipeline works by removing the confounding effects of imaging artifacts on skin pixels and fitting a line in the direction of the greatest variance in the RGB color space using PCA. The resulting metric is related to the skin tone of each subject relative to the rest of the photos in the dataset. Assumptions include consistent lighting, the same acquisition camera, and constant background. As a further limitation, the metric only indicates where a subject lies regarding net skin reflectance relative to the other subjects in the dataset, rather than an absolute measure.

The process to compute SREDS begins by extracting patches of skin from the forehead, right, and left cheeks using Dlib landmarks of each face image. Using the dichromatic reflection model as a guide, Non-Negative Matrix Factorization (NNMF) is used to estimate the diffuse and specular components of the selected skin patches. KPCA is utilized on the extracted diffuse components to learn a skin tone gradient across the dataset. The averaged value of the first principal components of the extracted diffuse bases for a particular face defines that person’s SREDS score. The KPCA model used for SREDS is data-driven, so the generalizability of the KPCA model onto unseen datasets is a point of interest within this study. A full description of the extraction of SREDS is found in [Ba21].

### 3.2 Datasets

For our experiments, we selected datasets that included demographic information of subjects across race, age, gender, orientation, and lighting. We utilized CMU Multi-PIE, MEDS-II, and Morph-II datasets [SBB03, Fo11, RT06]. Multi-PIE contains 750,000 sample images from 337 subjects images under 15 viewpoints with 19 illumination conditions. We selected three viewpoints (14 0, 05 1, 05 0) where full views of the face were captured for our testing, which reduced our sample images to 150,668 from 314 subjects. MEDS-II contains only 836 sample images from 425 subjects imaged in a controlled mugshot setting. Morph-II is a dataset from a longitudinal study that contains 55,063 sample images from 13,000 subjects within a controlled setting over 5 years. While MEDS-II and Morph-II datasets include uncontrolled illumination, Multipie includes controlled illumination samples. ITA, RSR, and SREDS scores were generated for all samples of each dataset.

### 3.3 Cross-Dataset Analysis

In prior work [Ba21], the intra-subject variance was used as a metric to describe the variance of a specific subject’s skin tone score across multiple samples. The low intra-subject variance shows the metric can produce a consistent value of the same subject independent of external conditions. An example of intra-subject variation can be seen in figure 1a. We evaluate and compare the intra-subject variance across all of our evaluation datasets and compare it to other methods. In addition, we test the generalizability of the SREDS metric

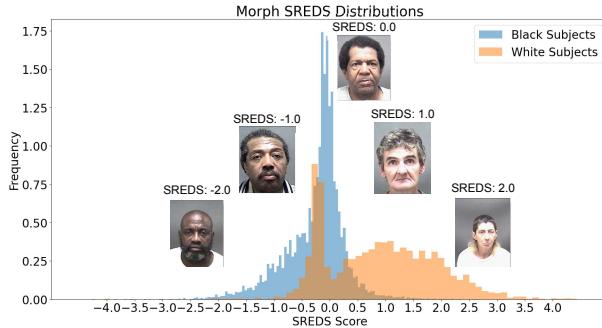


Figure 2: Distribution of SREDS Scores for Morph-II dataset separated by Race. Subjects with specific SREDS scores are shown for the range across the distribution.

by calibrating the skin tone gradient on one dataset and applying it to another, thereby testing its robustness to unseen datasets. The same experiment is run using the RSR PCA models for comparison. ITA does not have a training component and is only reported per dataset.

### 3.4 SREDS Agnostic Facial Recognition Model

To show the potential of SREDS for use in the replacement of race labels, we compared the outcome performance of SREDS versus ground-truth race labels when incorporated into SensitiveNets, a sensitive information removal network [Mo20]. SensitiveNets provides a novel privacy-preserving neural network feature representation to suppress the sensitive information of a learned space while maintaining the utility of the data. We reimplemented the sensitive removal network as our model for analysis of the suppression of race and skin tone. A diagram of our model is seen in figure 1b. SensitiveNets contains sensitive information removal dense layers added on top of a pre-trained face recognition backbone. Within our testing, we used a Resnet50 model pre-trained on VGGFace2 as the backbone, consistent with the cited literature [Mo20, Ca18, Xi]. The model's loss function requires a race classifier that acts as the sensitive information detector. The softmax probability from this detector describes the amount of racial information present within a subject's template and the goal of the loss function is to remove the sensitive race information and trend the classifier towards 50% accuracy. In our experiments, this classifier is either trained on race labels or SREDS scores binned into predetermined groups. The sensitive information removal  $\varphi$  layers are then added and trained sequentially using an adversarial approach of triplet loss and an adversarial sensitivity regularizer loss which reduces the amount of sensitive race information from the embedding space. An in-depth look at the model and loss function can be found in the SensitiveNets literature [Mo20].

## 4 Experiment Results

Our experiments were performed to analyze how the consistency of SREDS performed relative to other skin tone metrics and the outcome of replacing race labels with SREDS-generated labels in a privacy preservation method.

### 4.1 Cross-Dataset Analysis Results

We performed the cross-dataset analysis of the two skin tone metrics described in Section 2 and SREDS across the three datasets listed in Section 3. We generated ITA, RSR, and SREDS for all subjects from the mentioned datasets. As part of background normalization, RSR assumed consistent lighting, the same acquisition camera, and a constant background. Only the Multi-PIE dataset meets all conditions. However, due to the lack of constant background in MEDS-II and MORPH-II, the background normalization step was bypassed for these datasets. ITA is a non-trainable method so we collected the ITA values from each subject of each dataset. To test SREDS consistency on unseen data we used the Kernel Principal Component Analysis (KPCA) fit to one dataset’s diffuse components and used it to transform another dataset’s diffuse components. The same process was recreated using the RSR PCA models on the same datasets’ selected skin pixel values in order to compare these two methods.

Training Dataset	Testing Dataset								
	Morph-II			MEDS-II			Multi-Pie (Mugshot)		
	SREDS	RSR	ITA	SREDS	RSR	ITA	SREDS	RSR	ITA
Morph-II	0.419	0.539	0.645	0.681	0.493	N/A	0.157	0.468	N/A
MEDS-II	0.457	0.540	N/A	0.463	0.493	<b>0.448</b>	0.186	0.470	N/A
Multi-Pie (Mugshot)	<b>0.399</b>	0.538	N/A	0.674	0.493	N/A	<b>0.138</b>	0.304	0.401

Table 1: Cross dataset intra-subject variability analysis between SREDS, RSR, and ITA skin tone metrics. Bolded values are the lowest recorded intra-subject variability in that testing dataset. SREDS scores result in the least variable metric from Morph-II and Multi-Pie datasets and the second least variable metric in MEDS-II, behind ITA.

We computed the intra-subject variability of each dataset’s skin tone metrics by calculating the standard deviation of each subject’s individual skin tone measures and averaging across the dataset. The results of this analysis are seen in Table 1 and suggest that the learning-based algorithms (RSR and SREDS) perform better than ITA when evaluated on the dataset they are calibrated on. Viewing our cross-dataset results, we observe that in larger datasets (Morph, Multi-pie), SREDS outperforms both ITA and RSR even when calibrated on a different dataset, suggesting the generalizability of this approach.

### 4.2 Distribution of SREDS

To utilize SREDS by replacing race labels we needed a process to convert continuous SREDS scores into discrete labels. To understand the distribution of scores, the SREDS

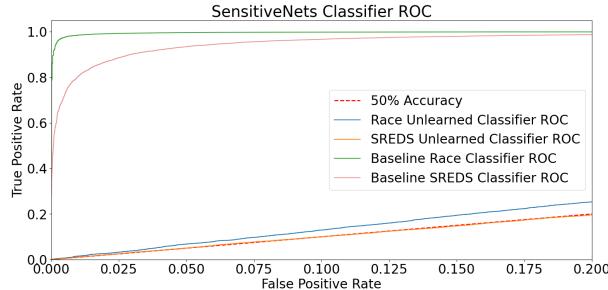


Figure 3: SensitiveNets Sensitive Information Classifier ROCs for both race labels and SREDS scores before and after training on Morph-II dataset. The goal of SensitiveNets training is for 50% classification accuracy. The unlearned classifier accuracy for both the RDM and SDM is nearly 50%, which shows SREDS scores and race labels perform similarly in this experiment.

scores across the Morph-II dataset were plotted within Figure 2. We split the dataset in half by the median SREDS score of -0.01 and binned the subjects into low and high SREDS scores to create a discrete labeling of the Morph-II dataset.

#### 4.3 Comparison of Race Labels and SREDS in Sensitive Feature Unlearning

To see the effects of SREDS scores being used in place of self-reported race labels, we implemented two SensitiveNets models. One model is trained using the black and white subject race labels from the Morph-II dataset while the second model is trained using the binned SREDS value for the same subjects.

Backbone	Classifier	Trained On	Tested On	ICA	FCA
Resnet50	Race	Race Triplets	Morph	0.985	0.47
Resnet50	SREDS	SREDS Triplets	Morph	0.937	0.48

Table 2: Sensitive Information Removal Network Experiment Results  
 ICA: Initial Classification Accuracy, FCA: Final Classification Accuracy (Goal of sensitive information removal is for FCA to be 0.50)

The first model trained on race labels and the second model trained on SREDS scores will be referred to as the Race Unlearned Model (RUM) and the SREDS Unlearned Model (SUM) respectively. An outline of this testing plan is seen in Table 2 with the initial and final classification accuracy of the SensitiveNets classifiers. For both models, the sensitive information classifier ROCs were calculated and shown in Figure 3.

The two trained SensitiveNets models matching performances are compared to the baseline Resnet50 matching performance to evaluate the results of the training on matching performance in Figure 4. The feature unlearning experiments to preserve privacy show a similar drop in performance between training with race labels and training with SREDS

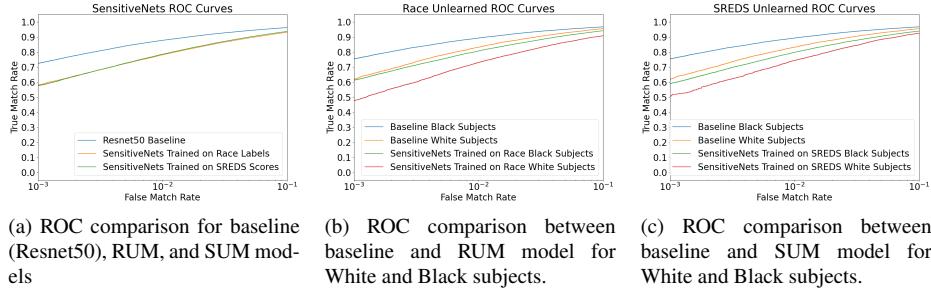


Figure 4: Comparison of biometric performance (matching) ROCs of baseline, RDM, and SDM, categorized by race labels on Morph-II dataset. Shows RDM and SDM suffer from a similar drop in matching performance when race or skin tone information is removed, respectively.

scores. The results suggest only a small (0.027) difference between the True Positive Rate (TPR) of RUM and SUM at  $10^{-3}$  False Positive Rate (FPR).

## 5 Conclusions

The existing feature unlearning methods in FR rely on large-scale and expensive-to-collect demographically-labeled datasets. Within this study, we demonstrate the ability of SREDS to mitigate this reliance by automatically extracting consistent skin tone information from face images. We have shown that SREDS outperforms other available skin tone metrics in producing continuous and less-variable skin tone estimates while generalizing well to unseen data. We have presented an application of extracted SREDS scores in the absence of race labels in a feature unlearning method and shown that SREDS could be used as a replacement.

### 5.1 Limitations and Future Work

Limitations of this work include our analysis of only black and white subjects due to the under-representation of other races in our datasets. This led to us only using two SREDS bins when categorizing our datasets to match the binary race labels. We tested using only one face matcher within our privacy-preserving method and have not seen how different networks affect our results. A limitation of using skin tone as a way to label datasets is that skin tone does not encapsulate the entirety of a self-reported race label. Skin tone is one physical characteristic that makes up race and cannot be used as an exact replacement.

Future work planned includes further analysis of the mapping of SREDS to multi-race demographic information and its use in different downstream biometric tasks, recreating our experiments with addition face matches [De22], and attempting a bias mitigation solution

using SREDS scores and evaluating using fairness metrics [dFPM22, Ho22] on an even larger scale dataset (BUPT-Globalface) [WZD21].

## References

- [Ba21] Bahmani, Keivan; Plesh, Richard; Sahu, Chinmay; Banavar, Mahesh; Schuckers, Stephanie: SREDS: A dichromatic separation based measure of skin color. In: 2021 IEEE International Workshop on Biometrics and Forensics (IWB). IEEE, pp. 1–6, 2021.
- [Ca18] Cao, Qiong; Shen, Li; Xie, Weidi; Parkhi, Omkar M.; Zisserman, Andrew: VGGFace2: A dataset for recognising faces across pose and age, 2018.
- [CCH91] Chardon, A; Cretois, I; Hourseau, C: Skin colour typology and suntanning pathways. *Int J Cosmet Sci*, 13(4):191–208, August 1991.
- [CF97] Connolly, C.; Fleiss, T.: A study of efficiency and accuracy in the transformation from RGB to CIELAB color space. *IEEE Transactions on Image Processing*, 6(7):1046–1048, 1997.
- [Co19] Cook, Cynthia M.; Howard, John J.; Sirotin, Yevgeniy B.; Tipton, Jerry L.; Vemury, Arun R.: Demographic Effects in Facial Recognition and Their Dependence on Image Acquisition: An Evaluation of Eleven Commercial Systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 1(1):32–41, 2019.
- [De22] Deng, Jiankang; Guo, Jia; Yang, Jing; Xue, Niannan; Kotsia, Irene; Zafeiriou, Stefanos: ArcFace: Additive Angular Margin Loss for Deep Face Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):5962–5979, oct 2022.
- [dFPM22] de Freitas Pereira, Tiago; Marcel, Sébastien: Fairness in Biometrics: A Figure of Merit to Assess Biometric Verification Systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 4(1):19–29, 2022.
- [Fi88] Fitzpatrick, T B: The validity and practicality of sun-reactive skin types I through VI. *Arch Dermatol*, 124(6):869–871, June 1988.
- [Fo11] Founds, Andrew; Orlans, Nick; Genevieve, Whiddon; Watson, Craig: NIST Special Database 32 - Multiple Encounter Dataset II (MEDS-II), 07 2011.
- [Ho22] Howard, John J.; Laird, Eli J.; Sirotin, Yevgeniy B.; Rubin, Rebecca E.; Tipton, Jerry L.; Vemury, Arun R.: Evaluating Proposed Fairness Models for Face Recognition Algorithms, 2022.
- [Ki06] King, Davis: Dlib C++ library, 2006.
- [Mo19] Monk, Ellis: Monk Skin Tone Scale, 2019.
- [Mo20] Morales, Aytahmi; Fierrez, Julian; Vera-Rodriguez, Ruben; Tolosana, Ruben: SensitiveNets: Learning Agnostic Representations with Application to Face Images, 2020.
- [MR17] Mirjalili, Vahid; Ross, Arun: Soft biometric privacy: Retaining biometric utility of face images while perturbing gender. In: 2017 IEEE International Joint Conference on Biometrics (IJCB). pp. 564–573, 2017.

[OR15] Othman, Asem; Ross, Arun: Privacy of Facial Soft Biometrics: Suppressing Gender But Retaining Identity. In (Agapito, Lourdes; Bronstein, Michael M.; Rother, Carsten, eds): Computer Vision - ECCV 2014 Workshops. Springer International Publishing, Cham, pp. 682–696, 2015.

[RT06] Ricanek, K.; Tesafaye, T.: MORPH: a longitudinal image database of normal adult age-progression. In: 7th International Conference on Automatic Face and Gesture Recognition (FGR06). pp. 341–345, 2006.

[SBB03] Sim, T.; Baker, S.; Bsat, M.: The CMU pose, illumination, and expression database. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1615–1618, 2003.

[Se22] Serna, Ignacio; Morales, Aythami; Fierrez, Julian; Obradovich, Nick: Sensitive loss: Improving accuracy and fairness of face representations with discrimination-aware deep learning. *Artificial Intelligence*, 305:103682, 2022.

[Te20a] Terhörst, Philipp; Fährmann, Daniel; Damer, Naser; Kirchbuchner, Florian; Kuijper, Arjan: Beyond Identity: What Information Is Stored in Biometric Face Templates?, 2020.

[Te20b] Terhörst, Philipp; Kolf, Jan Niklas; Damer, Naser; Kirchbuchner, Florian; Kuijper, Arjan: Post-Comparison Mitigation of Demographic Bias in Face Recognition Using Fair Score Normalization, 2020.

[Te20c] Terhörst, Philipp; Tran, Mai Ly; Damer, Naser; Kirchbuchner, Florian; Kuijper, Arjan: Comparison-Level Mitigation of Ethnic Bias in Face Recognition. In: 2020 8th International Workshop on Biometrics and Forensics (IWBF). pp. 1–6, 2020.

[WZD21] Wang, Mei; Zhang, Yaobin; Deng, Weihong: Meta Balanced Network for Fair Face Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.

[Xi] Xie, Weidi: Weidixie/Keras-VGGFACE2-RESNET50.

[Xu18] Xu, Depeng; Yuan, Shuhan; Zhang, Lu; Wu, Xintao: FairGAN: Fairness-aware Generative Adversarial Networks, 2018.