Durham Research Online

Deposited in DRO:

19 October 2021

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Yucer, S. and Tekras, F. and Al Moubayed, N. and Breckon, T.P. (2022) 'Measuring Hidden Bias within Face Recognition via Racial Phenotypes.', Proc. Winter Conference on Applications of Computer Vision Waikoloa, HI, 03-08 Jan 2022.

Further information on publisher's website:

<https://doi.org/10.1109/WACV51458.2022.00326>

Publisher's copyright statement:

© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Measuring Hidden Bias within Face Recognition via Racial Phenotypes

Seyma Yucer¹, Furkan Tektas³, Noura Al Moubayed¹ and Toby P. Breckon^{1,2}
Department of {Computer Science¹, Engineering²}, Durham University, Durham, UK
BuboAI³, Middlesbrough, UK

Abstract

Recent work reports disparate performance for intersectional racial groups across face recognition tasks: face verification and identification. However, the definition of those racial groups has a significant impact on the underlying findings of such racial bias analysis. Previous studies define these groups based on either demographic information (e.g. African, Asian etc.) or skin tone (e.g. lighter or darker skins). The use of such sensitive or broad group definitions has disadvantages for bias investigation and subsequent counter-bias solutions design. By contrast, this study introduces an alternative racial bias analysis methodology via facial phenotype attributes for face recognition. We use the set of observable characteristics of an individual face where a race-related facial phenotype is hence specific to the human face and correlated to the racial profile of the subject. We propose categorical test cases to investigate the individual influence of those attributes on bias within face recognition tasks. We compare our phenotype-based grouping methodology with previous grouping strategies and show that phenotype-based groupings uncover hidden bias without reliance upon any potentially protected attributes or ill-defined grouping strategies. Furthermore, we contribute corresponding phenotype attribute category labels for two face recognition tasks: RFW for face verification and VGGFace2 (test set) for face identification.

1. Introduction

An increasing number of automated face recognition systems have been deployed by companies, nonprofits and governments to make autonomous decisions for millions of users [1]. Such wide-scale adoption within real-world scenarios brings with it valid concerns on the potential abuse of face recognition due to the presence of data and algorithmic bias [2, 3]. The most common issue pertaining to such bias arises in racial groups [4]. Subsequently, the research community have been focused on methods that rely on demographic or skin type group annotations drawn from public face recognition benchmark datasets [5, 6]. This provides algorithmic performance on such predefined groupings to

measure bias. However, current grouping annotations and related bias evaluation strategies may lead to unintended negative implications- each of which we now detail to illustrate our motivation clearly.

Ambiguous Definition of Race: The historical and biological definitions of race vary and racial context is not fixed over time [7]. Such ambiguity becomes more problematic for the face recognition literature, as many researchers do not provide any related background about the details of their racial categorisation design process [8]. However, racial groupings are critical to the effective evolution of face recognition methodologies as they often represent the all-important means of quantitative evaluation. As in any recognition task, poorly defined groupings result in skewed mean and standard deviation measures of relative performance due to the ill-posed boundary conditions on membership of each group that can cause a given an example to justifiably transit from one group to another.

Privacy of Protected Attributes: Auditing benchmark datasets can cause potential privacy and consent violations [9] for dataset subjects. For example, exposing demographic origin may enhance the representations of a group under threat, leading to the potential for racial profiling and associated targeting [10]. As information of racial or ethnic origin is sensitive [11], researchers should either avoid revealing such sensitive data or provide an appropriate context for use [9].

Confined Groupings: Skin or racial grouping strategies such as binary {*light* vs. *dark*; *black* vs. *white*} for evaluating racial bias limits the scope of any study as they fail to capture the whole aspect of the bias problem where it needs to consider both multi-racial or less stereotypical members of such groups instead [12, 13] use Fitzpatrick skin type groupings to evaluate racial bias, but one such skin-tone based racial grouping contains multidimensional traits including nose, hair type, eye, and lips [14]. Leveraging all such traits together instead brings improved interpretations and derivations to address racial bias.

Racial Appearance Bias: Maddox [15] explains racial appearance bias as a negative disposition toward phenotypic variations in facial appearance. He also [16] discusses how race-conscious social policies may fail to address racial bi-

Study	Dataset Name	Release Year	Racial Grouping Strategy	Number of Image	Source
[17]	MORPH	2006	Caucasian, Hispanic, Asian, or African American	55K	Public Data
[18]	UTK Face	2017	Asian, Black, Indian, Others (like Hispanic, Latino, Middle Eastern) and White	20K	MORPH, CACD Web
[19]	RFW	2019	African, Asian, Caucasian, Indian	45K	MS-Celeb [6]
[20]	BUPT-Balanced	2020	African, Asian, Caucasian, Indian	1.3M	MS-Celeb [6]
[21]	Fair Face Challenge	2020	Light skin-toned (Fitzpatrick I-III), Dark skin-toned (Fitzpatrick IV-VI)	152K	IJB-C [22] and public domain images
[23]	Casual Conversations	2021	Fitzpatrick Skin Types	45K*	Vendor data
[24]	FairFace	2021	Black, East Asian, Indian, Latino, Middle Eastern, Southeast Asian, and White	108K	Flickr, Twitter Newspapers, Web
Ours	VGGFace2 [5]	2018 (2021)	Fitzpatrick Skin Types, Nose Shape, Eyelid Type, Lip Shape, Hair Type	3.3M	Google Image Search
Ours	RFW [19]	2019 (2021)	Fitzpatrick Skin Types, Nose Shape, Eyelid Type, Lip Shape, Hair Type	45K	MS-Celeb [6]

Table 1. Publicly available face datasets for different types of facial analysis tasks and their grouping strategies to address racial bias.

*Casual Conversations dataset provides videos.

ases in the treatment and outcomes of disadvantaged groups. Many studies show that individuals with more stereotypical racial appearance suffer poorer outcomes than those with less stereotypical appearance for their race [16, 25, 26]. On the other hand, a better understanding of the role of phenotypic variation complements solutions for both racial bias [15]. By way of phenotype, we mean the set of observable characteristics of an individual face where a race-related facial phenotype is hence specific to the human face and correlated to the racial profile of the subject.

Accordingly, we propose using race-related facial (phenotype) characteristics within face recognition to investigate racial bias. We categorise representative racial characteristics on the face and explore the impact of each characteristic phenotype attribute: skin types, eyelid type, nose shape, lips shape, hair colour and hair type. We audit these attributes for two different publicly available face datasets: VGGFace2 (test set) and RFW. We assess the impact of both attribute-based and subgroup-based evaluations on racial bias of face recognition tasks. We utilise two different training setups for face verification to compare performance disparities between imbalance and racially balanced training datasets. We compare our phenotype-based evaluation strategy with race or skin type based grouping evaluation. We show that our strategy provides a more elaborate perception of bias without revealing any potentially protected or ill-defined information.

This study presents a new evaluation strategy using facial phenotype attributes to investigate and measure racial bias with greater granularity within face recognition tasks. In this paper, our key contributions are as follows:

- we propose a new evaluation strategy that uses facial phenotype attributes rather than race labels to measure racial bias within both attribute-based and subgroup-based performance of state-of-the-art face recognition algorithms.
- we contribute additional facial phenotype attribute labelling for the VGGFace2 (face identification) and RFW (face verification) benchmark face datasets.
- we uncover the potentially hidden source of bias within the evaluation of racial groups, which is supported by quantitative evidence.

2. Related Work

Automated facial recognition encompasses two different tasks: face identification and face verification. For both tasks, studies present approaches that achieve overall performance on public benchmark datasets [22, 27] whilst the racial diversity within these datasets is often limited, biased and overlooked [28]. Consequently, numerous studies audit publicly available face datasets to demonstrate the dataset bias in face recognition [19, 29, 30]. However, the group definitions in use vary, meaning that a lack of consensus makes it significantly harder to tackle bias collaboratively due to an inconsistent problem definition across the literature.

We show leading benchmark face datasets with their grouping strategies in Table 1 to tackle racial bias. As a subset of MS-Celeb-1M, the RFW dataset [19] measures the racial performance of face verification on four different racial groups: *{African, Asian, Indian, Caucasian}*. FairFace [29] is another dataset drawn from the YFCC-100M Flickr dataset, providing additional group labels, *{Middle*

East, Latino} to evaluate bias on wider groupings. UTK-Face [31] is a large-scale face dataset with five different ethnicity categories for a variety of tasks, such as face detection, age estimation, age progression/regression, etc. More recently, The Casual Conversations Dataset [23] yielded from vendor data contains 45K videos with corresponding Fitzpatrick skin type labels. We subsequently categorise studies in the literature according to grouping strategies they adopt and explain each category below.

Racial Groupings: Although the definition of race carries a large amount of complexity and ambiguity, an increasing number of studies adopt various racial groupings and show performance disparities among them [4, 32]. The underlying reasons for such disparate results are summarised into two categories by [33]. Both the distribution of data on pre-defined racial groups and how we measure the bias play a major role in the results [33]. However, the majority of racial bias research rarely contains underlying details about how racial groups are determined or how racial bias evaluation metrics are designed [8]. Furthermore, [34] showed that non-explicit racial factors (accessories, hairstyles or facial anomalies) conflates with explicit racial factors (skin tone, nose shape or eye shape) and both factors strongly affect the recognition performance. He discusses the need to investigate each factor in order to have robust, fair, and explainable face recognition solutions. Such needs contradict the use of racial groups as they remain too narrow to have elaborate explanations [35].

Skin Type Groupings: Moreover, various studies [12, 21, 23] measure the racial bias in face recognition using either the Fitzpatrick Skin Types [36] or binary skin groups instead of using racial groups. Skin type grouping labels are yielded mostly from crowd-sourcing [30], neural network-based classifiers [37, 38] or professional human annotators. Merler [30] presents additional human-interpretable quantitative measures of intrinsic facial features along with subjective annotations. Although [39] show that Fitzpatrick Skin Types classification from uncontrolled imagery is a challenging task, [38] achieves high correlation with ground-truth labels using consistent reference points in human-level annotation interface.

Another issue is the impact of skin tone on racial bias. Cook [40] shows that the measure of skin reflectance on binary skin groups had the greatest net effect on the average biometric performance of face recognition. On the contrary, [41] claims it is not the case when it comes to continuous Fitzpatrick Skin Type groupings. Furthermore, skin tone is not enough for analysing racial bias as there is no clear evidence that skin tone is the primary driver for disparate false match rates [41]. Accordingly, many studies [42, 43] suggest looking for other race-related facial attributes including lip, eye, face shape in order to measure racial bias within face recognition.

In order to address these aforementioned issues, we propose a phenotype-based evaluation strategy for racial bias within face recognition. We provide facial phenotype attributes for the public benchmark datasets VGGFace2 and RFW in Table 1. We explain the phenotype-based attribute category selection process followed by our annotation framework in Section 4 and 5, respectively. We elaborate on our findings by measuring two state-of-the-art algorithms performance using imbalance and racially balanced training sets. Firstly, we analyse bias for each phenotype attribute category (attribute-based). Secondly, we produce different appearance-based joint distributions of face subjects and assess algorithm performance on subject groupings (subject-based). We provide related experiments with details in Section 6.

3. Ethical Considerations

Intent: This work intends to provide a novel racial bias analysis methodology via facial phenotype attributes for face recognition. The proposed strategy avoids the need for researchers to use potentially protected or ill-defined subject attributes and instead introduces racial phenotype attributes to explore racial bias in face recognition.

Denotation of Facial Phenotypes: We denote race-related phenotype attributes according to the studies of [44, 45] to have descriptive naming whilst avoiding causing any unintended offence to individuals.

Use of VGGFace2 and RFW: We conduct our experiments on two different face datasets which are publicly available for research use only. The reader is directed to the original source publication and the associated research organisation for access to these datasets. We make available supplementary facial attribute labels for these datasets in order to facilitate the use of our proposed evaluation strategy by other researchers, with the aim of furthering our stated intent above.

4. Racial Phenotypes on Face Images

In this section, we explain the categorisation of facial phenotype attributes for face recognition. Quine [46] presents three possible definitions of the race concept: a genetic variation between humans, morphological attributes, and genetically determined psychological characteristics. These morphological attributes are the primary interest for resolving racial bias in face recognition. For morphological attributes, studies [47, 48] focus on the impact of human phenotype characteristics over race estimation. They categorise the attributes by considering biological traits. The study of Shades of Race [44] investigates the marginal effects of phenotypic characteristics including skin tone, lips, nose, hair and body type on racial categorisation. Zhuang [49] considers 21 anthropometric measurements such as face width, length, nose breadth and length, eye corner

Figure 1. The distribution of facial phenotype attributes of RFW (left) and VGGFace2 Test (right) datasets.

points. He finds statistically significant differences in facial measurements between four racial/ethnic groups, which are {*Caucasian, Hispanic, African, other (mainly Asian)*}.

We adopt such groupings and measurements for face recognition by considering two limitations. Firstly, effectively evaluating face recognition tasks requires tight cropped (e.g. 112×112 px) low-quality images containing occlusion, shadows, and illumination variations for both the training and test stages. This makes phenotype attribute detection on the specific characteristics of face dataset images more difficult when compared to real-world human faces. Secondly, the broader categorisation increases the number of potential groupings, making bias evaluation inefficient for face recognition systems. Correspondingly, we decide to use 6 primary attributes that define the phenotype groupings for our study: skin type, eyelid type, nose shape, lip shape, hair type and hair colour¹. Subsequently, we have 21 different attribute categories under the 6 primary attributes as listed in Table 2.

Attribute	Categories	RFW	VGGFace2
Skin Type	Type 1 / 2 / 3 / 4 / 5 / 6	0.71	1.14
Eyelid Type	Monolid / Other	0.80	1.09
Nose Shape	Wide / Narrow	0.24	0.18
Lip Shape	Full / Small	0.28	0.63
Hair Type	Straight / Wavy / Curly / Bald	0.70	1.11
Hair Colour	Red / Blonde / Brown / Black / Grey	1.23	0.67

Table 2. Facial phenotype attributes and their categorisation based on [44] along with normalised standard deviations σ/μ .

We choose to use Fitzpatrick Skin Types [36] for skin tones as it provides more granularity, {*Type 1, Type 2, Type 3, Type 4, Type 5, Type 6*}, than binary skin-tone groupings, {*lighter skin-tone, darker skin-tone*}. The appearance of the human eye has been grouped by its position, shape and settings in many cosmetic industry guidelines [50]. However, they have either no scientific background or solid re-

lation with race. Instead, we look into epicanthal folds and check eyelid difference as it is a more distinctive attribute for racial bias [51]. We acknowledge that a single attribute category can be observed in multiple race groups. However, our main concern is identifying the most observable and convenient racial phenotype attributes on images to evaluate the bias (see Table 2).

For the appearance of the nose, we use two categories, wide and narrow, by examining the nasal breadth [49]. Hair texture is labelled into eight categories using the frequency of twists, waves, and curve diameter metric by [52]. Here we utilise eight categories and group them into three main hair texture types: straight, wavy, curly, in addition to bald. Despite being the most artificially manipulable attribute, we retain hair colour as it is related to skin tone [53]—the categories for hair colour we use: red, grey, black, blonde, brown (see Table 2).

5. Annotation of Racial Phenotypes

Previously in Section 4, we explain how we define racial phenotype attributes and their categories. Before the annotation process, we choose the most established face recognition datasets to validate our proposed methodology. For the face verification task, we choose the RFW dataset [19] as it provides a relatively broader racial distribution of subjects where each subject contains 3-5 images. For face identification, we use the VGGFace2 closed-test set [5], which contains at least 300 images per subject. For both datasets, we design an annotation interface to make the annotation process both user-friendly and robust. We present multiple sample images of a subject to avoid incorrect annotation caused by challenging samples such as grey-scale images, facial makeup, and poor scene illumination. Each subject is presented with attribute category selectors next to a set of face images within the annotation interface. Subsequently, an experienced annotator who has experience in morphological differences among races annotates each subject using the interface.

¹We note that hair information is still present in the tightly cropped images, and it may be helpful for future automated facial analyses tasks.

We obtain 11654 subjects annotations from the RFW and VGGFace2 benchmark datasets. Each annotation took 10-20 seconds, and overall annotation took 12 days (, i.e. annotator working at a maximum of 6 hours per day with regular breaks). The result of this annotation process, the phenotype attributes distributions for the RFW and VGGFace2 benchmark datasets, are shown in Figure 1 left/right, respectively. We also present the normalised standard deviations (Coefficient of Variance), σ/μ , among attribute categories of benchmark datasets to show the level of imbalance within these categories in Table 2. For both datasets, we can observe that the dominant phenotype attribute categories are Skin Type 3, Straight Hair, Narrow Nose, Other (non-monolid) Eyes, Small Lips, which correlates to the dominant presence of Caucasian faces based on the analysis of Figure 1.

6. Experimental Results and Discussion

In this section, we analyse the performance of our phenotype-based grouping methodology for face recognition tasks. We provide a public reference implementation, dataset reference links and pre-trained models ².

6.1. Training Setups

Setup 1 (Imbalanced Training Data): We train ArcFace [54] with a ResNet100 [55] on the VGGFace2 benchmark datasets that contains 8631 subjects where subject distribution is racially imbalanced. Here, our specific choice of VGGFace2 is due to investigate the impact of imbalanced training data that includes data bias on our proposed evaluation strategy.

Setup 2 (Racially Balanced Training Data): We use a ResNet34 [55] backbone architecture with the Softmax loss [56] trained on the BUPT-Balanced benchmark dataset [20] that contains 28000 face subjects. The BUPT-Balanced has racially balanced distributions among four groups {*African, Asian, Indian, Caucasian*} with 7000 face subjects each. The primary purpose of Setup 2 is to assess the impact of a racially balanced training dataset on results over the bias using our proposed phenotype-based methodology. We compare how much a racially balanced training dataset improved the performance difference compared to Setup 1.

6.2. Face Verification

Face verification, also known as one-to-one verification, is the task of comparing two different facial images to estimate whether they belong to the same individual subject. We follow two pairing strategies to explore the impact of a single attribute (attribute-based) and appearance-based facial groups (subgroup-based) on the evaluation performance of face verification.

Attribute Name	Setup 1 Accuracy %	Setup 2 Accuracy %
Blonde Hair	97.02	96.63
Red Hair*	96.33	96.83
Type 2	96.22	95.83
Gray Hair	94.85	95.83
Bald	94.75	95.70
Wavy Hair	94.32	95.50
Brown Hair	94.25	94.83
Type 6	93.77	94.77
Narrow Nose	92.92	94.77
Type 5	92.15	94.38
Curly Hair	92.02	93.63
Small Lips	91.92	94.98
Type 3	91.72	93.77
Type 1*	91.31	89.51
Straight Hair	91.25	94.32
Wide Nose	90.68	91.02
Full Lips	89.98	93.23
Type 4	89.90	93.55
Other Eye	89.88	93.75
Black Hair	89.88	91.42
Monolid Eye	88.27	89.73
σ	2.44	2.06
σ^*	2.39	1.77

Table 3. Attribute-based face verification performance of RFW. σ represents the standard deviation of all attribute category accuracies, including red hair and type 1, σ^* represents excluding standard deviation.

Attribute-based pairing: Firstly, we generate pairs from images containing the same attribute category—for example, facial images from people who all have monolid eyes. Consequently, we compare individual attributes performance using both training setups for face verification.

For attribute-based face verification, we randomly select 20k positive and 20k negative pairs from all possible pairs of each attribute. We calculate the cosine similarity of feature encoding of all selected negative and positive pairs to obtain the most challenging pairs. Subsequently, we select the most similar 3000 pairs from the negative samples and the least similar 3000 pairs from the positive samples for each attribute category in Table 3. Since the Type 1 category of the skin type attribute and red hair category of hair colour attribute do not have enough samples to generate 6000 pairs, we instead produce 602 pairs (301 positive, 301 negative) for Type 1, 1200 (600 positives, 600 negative) pairs for red hair.

In this way, we measure each face attributes accuracy using on face verification performance. We use both training setups to show how much standard deviation (σ) changes between balanced and imbalanced training data. We present the attribute-based sample groups in Table 3 with a standard deviation of accuracies excluding red hair and Type

²<https://github.com/seymayucer/FacialPhenotypes>

Skin	Lips	Eye	Nose	Hair Type	Ratio (%)	Accuracy (%)	Skin	Lips	Eye	Nose	Hair Type	Ratio (%)	Accuracy (%)
{1, 2}	Small	Other	Narrow	Straight	3.82	96.53	{3, 4}	Full	Monolid	Wide	Straight	1.55	91.63
{3, 4}	Small	Other	Narrow	Straight	7.43	96.45	{1, 2}	Small	Other	Narrow	Bald	0.28	91.29
{3, 4}	Small	Other	Narrow	Wavy	3.67	96.11	{5, 6}	Full	Other	Narrow	Curly	1.97	91.23
{1, 2}	Small	Other	Wide	Straight	3.03	95.63	{3, 4}	Small	Other	Wide	Bald	1.68	91.01
{1, 2}	Small	Other	Narrow	Wavy	1.64	95.62	{1, 2}	Full	Other	Narrow	Wavy	0.27	90.74
{1, 2}	Full	Other	Narrow	Straight	0.70	95.59	{3, 4}	Small	Monolid	Wide	Wavy	0.96	90.17
{3, 4}	Full	Other	Narrow	Straight	3.59	95.28	{1, 2}	Small	Other	Wide	Bald	0.46	89.78
{3, 4}	Full	Other	Wide	Straight	4.47	94.98	{5, 6}	Small	Other	Narrow	Curly	0.81	89.50
{3, 4}	Small	Other	Wide	Wavy	2.95	94.92	{3, 4}	Small	Monolid	Narrow	Wavy	1.20	89.35
{3, 4}	Small	Other	Wide	Straight	8.83	94.92	{5, 6}	Full	Other	Wide	Curly	13.09	89.18
{1, 2}	Full	Other	Wide	Straight	0.33	94.87	{3, 4}	Full	Other	Wide	Bald	0.80	86.02
{1, 2}	Small	Other	Wide	Wavy	0.72	94.56	{5, 6}	Small	Other	Wide	Bald	0.99	85.90
{3, 4}	Small	Other	Wide	Curly	0.51	93.89	{3, 4}	Full	Other	Wide	Curly	0.46	85.38
{3, 4}	Full	Other	Wide	Wavy	1.90	93.41	{3, 4}	Small	Monolid	Narrow	Bald	0.32	84.10
{3, 4}	Full	Other	Narrow	Wavy	1.94	93.10	{5, 6}	Small	Other	Narrow	Bald	0.30	82.81
{3, 4}	Small	Other	Narrow	Bald	0.68	92.50	{3, 4}	Small	Monolid	Wide	Bald	0.52	82.67
{3, 4}	Small	Other	Narrow	Curly	0.31	92.45	{3, 4}	Full	Monolid	Narrow	Wavy	0.43	82.04
{5, 6}	Small	Other	Wide	Curly	2.81	92.23	{5, 6}	Full	Other	Narrow	Bald	0.53	81.24
{3, 4}	Small	Monolid	Wide	Straight	6.59	91.93	{1, 2}	Small	Monolid	Narrow	Straight	0.47	81.04
{3, 4}	Full	Monolid	Narrow	Straight	1.81	91.78	{3, 4}	Full	Monolid	Wide	Wavy	0.27	79.47
{5, 6}	Full	Other	Wide	Bald	3.62	91.74	{5, 6}	Full	Other	Wide	Wavy	0.32	78.94
{3, 4}	Small	Monolid	Narrow	Straight	7.95	91.70							
σ												5.07	

Table 4. Subgroup-based face verification performance of RFW using training setup 1, sorted by descending order of accuracy.

1 attribute accuracies (σ^*) and including them (σ). It is clear from Table 3 that for both setup 1 (imbalanced training data) and setup 2 (racially balanced training data), accuracy is lower for monolid eyes, black hair, full lips, and wide nose than the other eye, blonde hair, and small lips, and narrow nose respectively. We also do find a slight correlation between darker skin tones and higher false matching rates when we pair from the same attribute categories (Supplementary Table S1). Moreover, although the imbalanced training setup results a bigger performance difference (σ), the amount of difference between two setups is small, meaning that a racially balanced dataset distribution is not enough to overcome performance bias.

Additionally, NIST [4] suggests providing false matching rates of pairing combinations between each grouping in the dataset as it is necessary for real-world scenarios. Therefore, we pair each attribute category with all other attribute categories to assess cross-attribute pairing performance. Subsequently, we evaluate false matching rates between any attribute category pair combination in Figure 2. We randomly generate 10000 pairs for each category pairings; in total, we have 441 (21×21) pairings. For example, each cross-attribute pairings means 10000 pairs between

blonde hair - monolid eye, type 3 - wide nose or wavy hair - full lips etc. As a result of this, we clearly show that Type 5, Type 6 and monolid eyes pairings have higher false matching rates among all attribute categories in Figure 2 using training setup 1. Consequently, the impact of the dark skin tones on performance increases for cross-attribute pairings compared to the attribute-based pairings.

Subgroup-based pairing: Secondly, we create various subgroups with different phenotypic attribute combinations in the dataset. For example, one such subgroup consists of subjects with skin type 3, monolid eyes, straight hair, wide nose, and small lips. Our main purpose of such pairing is to show the effects of single attribute changes over a group-for instance, what would change when only skin gets darker, but other attributes remain the same?

Furthermore, we generate all possible subgroups with different phenotypic attribute category combinations to investigate subgroup-based performances. However, we need to limit the number of subgroups such that we can present our results efficiently. We first remove the hair colour attribute as it is the easiest race-relevant attribute that individuals can readily modify via styling. Consequently, we merge skin types into three groups and show them as {1,2}

Figure 2. False matching rates (FMR) of cross-attribute based pairings for 21 attribute categories using training setup 1. Each cell depicts FMR on a logarithmic scale which is $\log_{10}(\text{FMR})$ with lower negative values (close to zero) encoding superior false match rates.

for Type 1 and Type 2, $\{3,4\}$ for Type 3 and Type 4, and $\{5,6\}$ for Type 5 and Type 6. Lastly, we remove subgroups with a few or even no samples in the test set, which comprises 3% of all samples. In Table 4, we show the performance of each subgroup with its proportion in the original test dataset. To evaluate the performance, we generate 6000 pairs (3k positive and 3k negative) from all possible pairs of subgroups that have enough samples. For the rest, we generate an equal number of negative and positive pairs as much as availability facilitates. From our observation of

Table 4, we can conclude that groups who have one of the attributes like wide nose, full lips, and monolid eye type always have less accuracy than the other groups with a narrow nose, small lips and other eye (when rest of the attributes are same). Furthermore, whilst the average accuracy of the subgroups with Type $\{5,6\}$ skin type is 86.97%, subgroups with Type $\{1,2\}$ skin type is 92.56%, but this notably includes other attributes effects.

Moreover, the number of subgroup variations with darker skin tones are much smaller than lighter tones which

Figure 3. Accuracy variations for three grouping strategies. Standard deviation of the groupings reflects the amount of measured bias. Racial groupings {*African, Asian, Caucasian, Indian*} accuracies are obtained from [20]. Binary skin tones {*lighter skin-tone, darker skin-tone*} are the average accuracy of Type 1-3 and Type 4-6 skin tones, respectively.

causes many different evaluation and analysis problems. It lacks sufficient interpretation in the test phase; there are minorities in the global populous with dark skin and monolid eyes or any other less common variations. Benchmark datasets do not contain enough representations for such minority groups. An improved evaluation dataset would be one that is able to cover more phenotype combinations such that its distribution is an unbiased representation of the global populous.

Lastly, we estimate such disparities among different grouping strategies using training setup 2. We take racial groupings {*African, Asian, Indian, Caucasian*} and binary skin tone groupings {*lighter skin-tone, darker skin-tone*} as they are very common grouping strategies in the literature. We compare them with our phenotype-based grouping strategy. In Figure 3, we show that how accuracy and the standard deviation differs between sub-groups in three different strategies. Higher variation reveals hidden bias, which may be missed in narrow, erroneous racial or binary skin tones grouping strategies. The phenotype-based grouping strategy brings a more granular observation of the variability in performance (i.e. higher standard deviation) and hence a more resolute measure of performance bias.

6.3. Face Identification

Face identification as a one-to-many verification is the task of searching for a face across a facial database. There are two scenarios for face identification applications based on whether a queried face is enrolled in a database or not. Open-set identification assumes the database does not necessarily contain the queried face, while closed-set identification always looks for a match in the database. In this study, we apply closed-set identification using the test set

Attribute	Ratio (%)	Acc (%)	Attribute	Ratio (%)	Acc (%)
Bald	2.80	97.49	Type 6	1.60	96.25
Grey Hair	12.60	97.47	Wide Nose	43.60	96.19
Red Hair	0.80	97.10	Type 3	42.80	96.13
Type 5	3.60	96.87	Brown Hair	34.20	96.05
Type 4	9.60	96.75	Curly Hair	4.40	95.93
Small Lips	72.60	96.56	Wavy Hair	31.00	95.92
Type 2	39.20	96.43	Monolid Eye	11.00	95.73
Black Hair	29.80	96.43	Blonde Hair	22.60	95.52
Straight Hair	61.80	96.35	Full Lips	27.40	95.36
Other Eye	89.00	96.29	Type 1	3.20	92.90
Narrow Nose	56.40	96.26			
σ		0.93			

Table 5. Face identification performance on VGGFace2 test set using standard linear SVM and features from training setup 1, sorted by descending order of accuracy.

of the VGGFace2 benchmark dataset on the originally proposed protocol [5] and we extract the image features using training setup 1 [54]. We apply a 5-fold train-test split where we sample 50 images from each subject as the test set and use the rest as the training set. We train a standard linear SVM on the extracted feature representations and predict the identities for test samples. Our results are shown in Table 5 where we can observe that the standard deviation (σ) is much smaller when compared to the earlier attribute-based face verification results of Table 3. It shows that the closed-set face identification does not have the same level of bias correlation as we find for face verification. However, in this experiment, we are unable to have the same proportion for each attribute, and we did not measure open-set face identification. As suggested in [4], future work should design and apply open-set tests for face identification on better-distributed benchmark datasets to measure bias extensively.

7. Conclusion

We propose a new evaluation strategy using facial phenotype attributes to assess racial bias in face recognition tasks. We elaborate experimental results to show the impact of each phenotype attributes using two different training setups, including imbalanced and racially balanced training sets. We also provide different pairing strategies for face verification to draw attention to the importance of pairing for comprehensive evaluation. We observe apparent performance differences between race-related phenotype attribute categories and subgroups for both training setups. However, we also uncover more considerable performance disparities among phenotype attributes than racial groups. We demonstrate that phenotype-based evaluation strategy reveals racial bias comprehensively whilst avoiding exposing potentially protected or ill-defined attributes. Future work will focus on improving facial appearance variations using generative models to provide more balanced and realistic test scenario distributions.

References

- [1] Yassin Kortli, Maher Jridi, Ayman Al Falou, and Mohamed Atri. Face recognition systems: A survey. *Sensors*, 2020.
- [2] Raul Vicente Garcia, Lukasz Wandzik, Louisa Grabner, and Joerg Krueger. The harms of demographic bias in deep face recognition research. In *2019 International Conference on Biometrics (ICB)*. IEEE, 2019.
- [3] Nisha Srinivas, Karl Ricanek, Dana Michalski, David S Bolme, and Michael King. Face recognition algorithm bias: Performance differences on images of children and adults. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [4] Patrick Grother, Mei Ngan, and Kayee Hanaoka. *Face Recognition Vendor Test (FVRT): Part 3, Demographic Effects*. National Institute of Standards and Technology, 2019.
- [5] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *IEEE International Conference on Automatic Face & Gesture Recognition*, 2018.
- [6] 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*. Springer, 2016.
- [7] Jayne Chong-Soon Lee. Navigating the topology of race, 1994.
- [8] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R Brubaker. How we’ve taught algorithms to see identity: Constructing race and gender in image databases for facial analysis. *ACM on Human-Computer Interaction*, 2020.
- [9] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. Saving face: Investigating the ethical concerns of facial recognition auditing. In *AAAI/ACM Conference on AI, Ethics, and Society*, 2020.
- [10] Paul Mozur. One month, 500,000 face scans: How china is using ai to profile a minority. *The New York Times*, 2019.
- [11] Bob Hepple. The new single equality act in britain. *The Equal Rights Review*, 2010.
- [12] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 2018.
- [13] Margaret Mitchell, Dylan Baker, Nyalleng Moorosi, Emily Denton, Ben Hutchinson, Alex Hanna, Timnit Gebru, and Jamie Morgenstern. Diversity and inclusion metrics in subset selection. In *AAAI/ACM Conference on AI, Ethics, and Society*, 2020.
- [14] Wendy D Roth. The multiple dimensions of race. *Ethnic and Racial Studies*, 2016.
- [15] Keith B Maddox. Perspectives on racial phenotypicality bias. *Personality and Social Psychology Review*, 2004.
- [16] Keith B Maddox and Jennifer M Perry. Racial appearance bias: Improving evidence-based policies to address racial disparities. *Policy Insights from the Behavioral and Brain Sciences*, 2018.
- [17] Karl Ricanek and Tamirat Tesafaye. Morph: A longitudinal image database of normal adult age-progression. In *7th International Conference on Automatic Face and Gesture Recognition (FG06)*. IEEE, 2006.
- [18] Zhifei Zhang, Yang Song, and H. Qi. Age progression/regression by conditional adversarial autoencoder. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [19] 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 *IEEE International Conference on Computer Vision*, 2019.
- [20] Mei Wang and Weihong Deng. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [21] Tomáš Sixta, Julio Junior, CS Jacques, Pau Buch-Cardona, Eduard Vazquez, and Sergio Escalera. Fairface challenge at eccv 2020: Analyzing bias in face recognition. *arXiv preprint arXiv:2009.07838*, 2020.
- [22] 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 *IEEE International Conference on Biometrics*, 2018.
- [23] Caner Hazirbas, Joanna Bitton, Brian Dolhansky, Jacqueline Pan, Albert Gordo, and Cristian Canton Ferrer. Casual conversations: A dataset for measuring fairness in ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [24] Kimmo Kärkkäinen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age. *arXiv preprint arXiv:1908.04913*, 2019.
- [25] Allison L Skinner and Gandalf Nicolas. Looking black or looking back? using phenotype and ancestry to make racial categorizations. *Journal of Experimental Social Psychology*, 2015.
- [26] Kimberly Barsamian Kahn and Paul G Davies. Differentially dangerous? phenotypic racial stereotypicality increases implicit bias among ingroup and outgroup members. *Group Processes & Intergroup Relations*, 2011.
- [27] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, University of Massachusetts, Amherst, 2007.
- [28] Erjin Zhou, Zhimin Cao, and Qi Yin. Naive-deep face recognition: Touching the limit of lfw benchmark or not? *arXiv preprint arXiv:1501.04690*, 2015.

- [29] Kimmo Karkkainen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021.
- [30] Michele Merler, Nalini Ratha, Rogerio S Feris, and John R Smith. Diversity in faces. *arXiv preprint arXiv:1901.10436*, 2019.
- [31] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [32] 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*, 2020.
- [33] Jacqueline G Cavazos, P Jonathon Phillips, Carlos D Castillo, and Alice J O’Toole. Accuracy comparison across face recognition algorithms: Where are we on measuring race bias? *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2020.
- [34] Philipp Terhörst, Jan Niklas Kolf, Marco Huber, Florian Kirchbuchner, Naser Damer, Aythami Morales, Julian Fierrez, and Arjan Kuijper. A comprehensive study on face recognition biases beyond demographics. *arXiv preprint arXiv:2103.01592*, 2021.
- [35] Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Krones, Meredith Ringel Morris, Jennifer Wortman Vaughan, Duncan Wadsworth, and Hanna Wallach. Designing disaggregated evaluations of ai systems: Choices, considerations, and tradeoffs. *arXiv preprint arXiv:2103.06076*, 2021.
- [36] Thomas B Fitzpatrick. The validity and practicality of sun-reactive skin types i through vi. *Archives of dermatology*, 1988.
- [37] Philipp Terhörst, Daniel Fährmann, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Maad-face: A massively annotated attribute dataset for face images. *arXiv preprint arXiv:2012.01030*, 2020.
- [38] KS Krishnapriya, Michael C King, and Kevin W Bowyer. Analysis of manual and automated skin tone assignments for face recognition applications. *arXiv preprint arXiv:2104.14685*, 2021.
- [39] John J Howard, Yevgeniy B Sirotnin, Jerry L Tipton, and Arun R Vemury. Reliability and validity of image-based and self-reported skin phenotype metrics. *arXiv preprint arXiv:2106.11240*, 2021.
- [40] Cynthia M Cook, John J Howard, Yevgeniy B Sirotnin, Jerry L Tipton, and Arun R Vemury. 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*, 2019.
- [41] KS Krishnapriya, Vítor Albiero, Kushal Vangara, Michael C King, and Kevin W Bowyer. Issues related to face recognition accuracy varying based on race and skin tone. *IEEE Transactions on Technology and Society*, 2020.
- [42] Vidya Muthukumar, Tejaswini Pedapati, Nalini Ratha, Prasanna Sattigeri, Chai-Wah Wu, Brian Kingsbury, Abhishek Kumar, Samuel Thomas, Aleksandra Mojsilovic, and Kush R Varshney. Understanding unequal gender classification accuracy from face images. *arXiv preprint arXiv:1812.00099*, 2018.
- [43] Vidya Muthukumar. Color-theoretic experiments to understand unequal gender classification accuracy from face images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [44] Cynthia Feliciano. Shades of race: How phenotype and observer characteristics shape racial classification. *American Behavioral Scientist*, 2016.
- [45] Abdulla Fakhro, Hyung Woo Yim, Yong Kyu Kim, and Anh H Nguyen. The evolution of looks and expectations of asian eyelid and eye appearance. In *Seminars in plastic surgery*. Thieme Medical Publishers, 2015.
- [46] Willard V Quine. Three grades of modal involvement. In *XIth International Congress of Philosophy*, 1953.
- [47] Neven Sesardic. Race: a social destruction of a biological concept. *Biology & Philosophy*, 2010.
- [48] Stephen Ousley, Richard Jantz, and Donna Freid. Understanding race and human variation: why forensic anthropologists are good at identifying race. *American Journal of Physical Anthropology*, 2009.
- [49] Ziqing Zhuang, Douglas Landsittel, Stacey Benson, Raymond Roberge, and Ronald Shaffer. Facial anthropometric differences among gender, ethnicity, and age groups. *Annals of occupational hygiene*, 2010.
- [50] Theiab Alzahrani, Waleed Al-Nuaimy, and Baidaa Al-Bander. Integrated multi-model face shape and eye attributes identification for hair style and eyelashes recommendation. *Computation*, 2021.
- [51] Yoonho Lee, Euitae Lee, and Won Jin Park. Anchor epicanthoplasty combined with out-fold type double eyelidplasty for asians: do we have to make an additional scar to correct the asian epicanthal fold? *Plastic and reconstructive surgery*, 2000.
- [52] Roland De La Mettrie, Didier Saint-Léger, Genevieve Loussouarn, Annelise Garcel, Crystal Porter, and André Langaney. Shape variability and classification of human hair: a worldwide approach. *Human biology*, 2007.
- [53] Jonathan L Rees. Genetics of hair and skin color. *Annual review of genetics*, 2003.
- [54] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE International Conference on Computer Vision and Pattern Recognition*, 2019.
- [55] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [56] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.

Measuring Hidden Bias within Face Recognition via Racial Phenotypes

- Supplementary Material

Seyma Yucer¹, Furkan Tektas³, Noura Al Moubayed¹ and Toby P. Breckon^{1,2}
Department of {Computer Science¹, Engineering²}, Durham University, Durham, UK
BuboAI³, Middlesbrough, UK

8. Supplementary Material

8.1. Attribute-based Face Verification

We present attribute-based face verification scores including False Non-Match Rate (FNMR), False match rate (FMR) and F1 score in the Table S1. We use the same pairings and protocol [27] presented in the Section 6.2 for Table 3.

Whilst F1 scores are correlated with Table 3 accuracies, for the imbalanced training setup 1, the false matching ratio is higher on attributes like Monolid Eye, Type 6/5/4/3, Wide Nose, Full Lips than the different categories under the same attribute. Moreover, we observe that the balanced training setup 2 improves the FMR while increasing the FNMR for the attribute categories with higher accuracies and F1 scores.

Attribute Name	Setup 1			Setup 2		
	F1	FNMR	FMR	F1	FNMR	FMR
Blonde Hair	96.85	1.40	3.83	96.04	2.53	3.13
Red Hair	96.60	2.83	4.00	96.48	2.83	2.83
Type 2	95.98	3.10	3.90	95.25	3.77	3.90
Bald	95.32	5.00	5.37	95.44	3.23	4.97
Gray Hair	95.00	3.70	5.87	95.93	2.53	5.13
Brown Hair	94.46	6.40	4.63	94.08	5.43	4.43
Type 6	94.42	4.10	7.30	95.01	5.53	4.67
Wavy Hair	93.96	3.27	7.53	95.42	4.97	3.77
Narrow Nose	93.05	6.77	7.20	94.29	4.40	5.87
Type 5	92.72	4.07	10.33	94.45	5.63	5.47
Curly Hair	92.51	5.47	9.67	93.58	6.87	5.70
Small Lips	92.36	5.80	8.37	94.29	5.03	4.70
Type 1	92.08	5.99	6.45	90.14	6.67	9.41
Type 3	91.80	8.63	7.73	93.59	5.93	6.13
Straight Hair	91.19	9.17	6.87	93.97	4.30	6.53
Other Eye	91.16	7.23	7.27	93.76	7.43	4.47
Wide Nose	90.99	7.23	7.43	89.78	7.10	5.27
Full Lips	90.73	6.60	10.17	93.43	7.13	5.77
Type 4	90.45	8.30	8.53	93.50	5.70	6.93
Black Hair	90.12	7.77	8.73	90.50	6.83	5.83
Monolid Eye	88.84	9.53	13.03	90.62	8.47	6.93

Table S1. Attribute-based face verification F1, FNMR, TMR scores of RFW dataset on both training setups.