

Face Recognition Accuracy Across Demographics: Shining a Light Into the Problem

Haiyu Wu¹, Vitor Albiero¹, K. S. Krishnapriya³, Michael C. King², Kevin W. Bowyer¹

¹University of Notre Dame, ²Florida Institute of Technology, ³Valdosta State University

Abstract

We explore varying face recognition accuracy across demographic groups as a phenomenon partly caused by differences in face illumination. We observe that for a common operational scenario with controlled image acquisition, there is a large difference in face region brightness between African-American and Caucasian, and also a smaller difference between male and female. We show that impostor image pairs with both faces under-exposed, or both over-exposed, have an increased false match rate (FMR). Conversely, image pairs with strongly different face brightness have a decreased similarity measure. We propose a brightness information metric to measure variation in brightness in the face and show that face brightness that is too low or too high has reduced information in the face region, providing a cause for the lower accuracy. Based on this, for operational scenarios with controlled image acquisition, illumination should be adjusted for each individual to obtain appropriate face image brightness. This is the first work that we are aware of to explore how the level of brightness of the skin region in a pair of face images (rather than a single image) impacts face recognition accuracy, and to evaluate this as a systematic factor causing unequal accuracy across demographics. The code is at <https://github.com/HaiyuWu/FaceBrightness>

1. Introduction

Widespread adoption of facial recognition technology in recent years has brought the startling realization that accuracy can vary significantly across demographics. This phenomenon has been widely discussed in news stories [29, 34, 43, 48] and research papers [12, 16, 17, 21, 25, 33, 47, 51]. In response, facial recognition researchers [18, 30, 33, 36] have explored the fundamental questions of *why* and *how* recognition accuracy varies across demographics. This work approaches the issue of accuracy variation through the lens of image brightness, in the context of controlled image acquisition scenarios such as driver’s license, passport, ID card, etc. The contributions of this work are:



Figure 1. Impostor image pairs have, on average, higher similarity scores when both images are too bright (“over-exposed”) or too dark (“under-exposed”), due to less information available to distinguish between faces. Pairs in (a) and (b) are the same two persons, but images in (a) have appropriate brightness and in (b) are over-exposed, resulting in (b) having a higher similarity score. Similarly, Pairs in (c) and (d) are the same two persons, but images in (c) have appropriate brightness and in (d) are under-exposed, resulting in (d) having a higher similarity score.

- Using an operational dataset acquired under conditions typical of mugshot, passport or ID card photo, we show that African-Americans are more likely to have under-exposed images and Caucasians are more likely to have over-exposed images. (See Section 4.)
- We propose a face skin brightness (FSB) metric that avoids regions of the image that are not related to recognition (hair color, sunglasses, ...), to give a more recognition-relevant brightness estimate than a commercial ICAO-compliance SDK. (See Section 3.1.)
- Using ArcFace [19] and a current commercial-off-the-shelf (COTS) matcher, we show that impostor pairs that are “too dark” and “too bright” can have a FMR twice that of pairs with appropriate brightness. Also, impostor pairs that have strongly different brightness have a reduced FMR relative to pairs with appropriate brightness. (See Section 5.)
- We provide a target brightness range which supports

best accuracy for all demographic (i.e. male/female and Caucasian/African American) for both matchers. It can easily be applied in real life by using the FSB metric. (See Section 6.1.)

2. Related Work

A detailed discussion of face image quality metrics related to illumination appears in a recent survey [44]. We briefly mention selected related works considering illumination and brightness in face images in different contexts.

A large body of work was enabled by the Pose, Illumination and Expression (PIE) [28, 45] and Multi-PIE [23] datasets. These datasets acquired images with flashes at various positions around the subject in order to explore the role of lighting direction. These datasets do not support exploring the effects of the general illumination level.

Beveridge et al [15] concluded that lighting direction can explain significant drops in matching accuracy. Han et al [28] surveyed the effect of various illumination preprocessing techniques on PIE and other datasets in terms of improving recognition accuracy for a selection of pre-deep-learning recognition algorithms. De Marsico et al [38] proposed an illumination-distortion metric that reflects the degree of uniformity in face illumination, computed as the standard deviation of the brightness of a set of eight face regions. Their metric does not capture the overall level of brightness. Abaza et al [8] proposed a face image quality index based on factors such as contrast, brightness, sharpness, focus and illumination. They used gamma adjustment to create darker and brighter versions of an image. They found that lowered-brightness images have slightly reduced rank-one recognition accuracy and increased-brightness images generally have the same accuracy as the original. *Their experiment shows higher brightness can increase the recognition rate.* Later, in Table 1, we show that higher brightness must eventually result in lower accuracy. Grm et al [22] created degraded versions of probe images from the LFW dataset [31] including simulated over-exposure. Their results show that over-bright probe images have lower verification accuracy. However, based on Fig. 4 in [22], too much brightness was added, resulting in unrealistic-looking images and possibly affecting their conclusion. Terhörst et al [46] combined a face recognition model and an image-specific quality regression model to generate saliency maps in order to investigate how image quality impacts the face recognition model. They concluded that face regions with high brightness lead to degraded recognition performance.

Some papers use the average brightness divided by the highest brightness value on the face as a quality metric [35, 39]. *Taken to the extreme, this metric says that an over-bright image is higher quality.* None of the works mentioned above use their metric to characterize pairs of images compared for recognition. A very recent work [32]

considers the problem of decreased accuracy for poorly illuminated images and proposes a feature restoration network. They consider image brightness as a binary attribute of appropriate / dim, and so do not recognize an effect for too-bright images. They also do not have a metric of brightness, and work with the small Specs on Faces (SoF) dataset having images of just 112 persons.

Exploring possible bias in recognition accuracy, researchers have investigated whether persons with darker skin tone have decreased accuracy. Skin tone is typically rated on a scale from 1 (lightest) to 6 (darkest) [33, 36, 50]. In the Janus program, Mechanical Turk workers rated web-scraped, in-the-wild celebrity images [50]. Lu et al [36] analyzed accuracy across skin tone on the Janus dataset and found that the ROC curves degraded from skin tones 1 to 5, but that skin tone 6 had a better ROC than 5. In other work [33], multiple human annotators rated skin tone for a set of images from the MORPH dataset [41]. The distribution of different skin-tone pairs was compared for the middle and the high-similarity tail of the impostor distribution, with the conclusion that there is no clear evidence of increased FMR for darker skin tone. In all studies where observers rate the skin tone of a celebrity, there is the likelihood that the rating is influenced by recognizing the celebrity and the person's race, rather than the rating being based solely on face brightness in the image.

A 2022 FRVT report [24] perhaps comes the closest to connecting the brightness of an image pair to matching accuracy. The report states that "... False negatives are in large part due to one or both photographs being of poor quality, something that can be coupled with demographics. ... dark skin has more challenging dynamic range capture requirements, with underexposed facial regions resulting in reduced information available to the algorithm." However, this report does not suggest that over-exposure can have the same effect for essentially the same reason.

Distinctive elements of this work relative to the works described above include the following. Most of the above works focus on brightness of a single image rather than an image pair, and do not consider how the effects vary across demographics [8, 15, 22, 23, 38, 45]. This paper focuses on brightness of the image pair, and on how this affects accuracy across demographics. There is previous work investigating how manually-rated skin tone, in particular darker skin tones, in a pair of images is associated with recognition accuracy [33, 36]. This paper focuses on analyzing automatically-computed face brightness for an image pair, and how this affects recognition accuracy, across demographics and whether the brightness is too low or too high.




Figure 2. Region used for “Face Skin Brightness” metric. From left: original color image; using whole image to estimate brightness and its result; using whole face area to estimate brightness and its result; using FSB metric to estimate brightness and its result.

3. Face Image Brightness Distributions

There are many web-scraped, in-the-wild, celebrity image datasets: MS-Celeb [27], WebFace260M [53], VG-GFace [40] and others. With web-scraped datasets, one cannot know if any image is brightened, contrast-enhanced or otherwise manipulated. To study how face region brightness in originally-acquired images affects accuracy, we use the MORPH dataset [6, 41]. MORPH images are acquired in controlled conditions typical of mugshot, passport or ID-card photos, including nominally frontal pose, neutral expression, consistent indoor lighting and plain 18% gray background. MORPH was assembled from public records, and is widely used in face aging [42] and in study of demographic accuracy variation [11, 14, 20, 33]. The version of MORPH used contains 127,319 images: 56,245 images of 8,839 African-American males, 24,857 images of 5,929 African-American females, 35,276 images of 8,835 Caucasian males, and 10,941 images of 2,798 Caucasian females. Faces were detected and aligned using img2pose [13].

3.1. Face Skin Brightness Metrics

We compare two metrics for estimating brightness of a face image. One is the metric reported by the iFace SDK [2, 3] (version 3.13) for checking compliance with International Civil Aviation Organization (ICAO) standards [5, 37]. The other is a metric specifically developed for this work.

ICAO-compliant images have “appropriate brightness and contrast that show skin tones naturally” [49]. iFace defines brightness as a “face attribute for evaluating whether an area of the face is correctly exposed” [4]. Assigned values range from -10,000 (“too dark”) to +10,000 (“too light”), with values near 0 indicating “OK”. iFace gives decision thresholds of -5,000 for under-exposure and +5,000 for over-exposure. With the motivation that the face skin


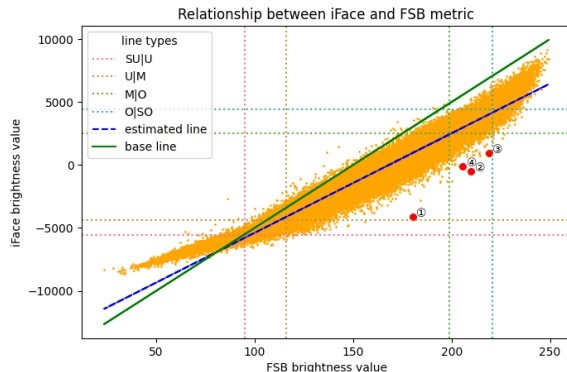


Figure 3. Comparison of iFace and FSB brightness values. ①, ②, ③, and ④ are the top outliers in each demographic group. These vertical and horizontal dotted lines are the boundaries of each brightness category: **Strongly Under-exposed (SU)**, **Under-exposed (U)**, **Middle-exposed (M)**, **Over-exposed (O)**, and **Strongly Over-exposed (SO)**. The brightness value and brightness group are shown on the top of each example, where the value’s format of each person is $FSB(group)/iFace(group)$.

region drives face matching accuracy, we introduce a face skin brightness (FSB) metric. We use the Bilateral Segmentation Network (BiSeNet) [1, 52] to segment a face image into 13 regions, such as face skin, nose, eyes, lips, etc. Previous approaches [7, 9, 18] for estimating the brightness fail to give a correct brightness estimation in the case showing in Figure 2, because beard and background strongly impact the brightness result. To alleviate the impact from the other factors, we omit the eyes, eyebrows, and lips from the calculation in order to avoid brightness variation due to sclera, eyelashes/mascara, eyebrows and lipstick. We also omit the nose region because it often contains specular highlights, and the skin region below the level of the nose because for men this often contains mustache, or beard. This resulting region used for our metric is illustrated in the last column of Figure 2. FSB is then computed as:

$$FSB_{image} = \frac{\sum_{h,w \in FS} Image_{gray}(h,w)}{N} \quad (1)$$

where FS is the selected face region, (h,w) is the pixel position in FS , and N is the number of the pixels in FS . FSB ranges from 0 (darkest) to 255 (brightest).

Figure 3 compares iFace and FSB brightness for the MORPH images. iFace estimates relatively higher brightness than FSB for under-exposed images, and relative lower brightness than FSB for the other images. To better under-



Figure 4. Examples across brightness and demographics.

stand how the metrics differ, Figure 3 shows the strongest outlier image for each demographic. These images are in the upper third of the brightness range of our FSB metric, but significantly lower in the range of the iFace metric. While we do not know how the iFace metric is calculated, we speculate that hair regions are included in its calculation.

4. Face Brightness Distributions

Figure 5 shows the iFace and FSB distributions for MORPH images, broken out by demographic. For both iFace and FSB, there is a smaller second peak at higher brightness in the female and male Caucasian distributions. We speculate this is due to specularities on the forehead in some fraction of the images. There is a narrow early peak at lower brightness on the iFace (only) distributions for African-American female and male. We speculate this is due to the iFace SDK including hair in the region it uses to estimate brightness. For both metrics, African-American males have the lowest average brightness (for FSB, mean = 138.7, $\sigma = 32.5$), followed closely by African-American females (for FSB, mean = 143.5, $\sigma = 32.6$), then a larger

gap to Caucasian males (for FSB, mean = 188.3, $\sigma = 25.6$), followed closely by Caucasian females (for FSB, mean = 191.4, $\sigma = 25$). Note that brightness metrics for African-American females and males have larger standard deviation than for Caucasian females and males.

To analyze how recognition accuracy varies for too-dark or too-bright face skin regions, we divide the brightness distribution for the whole image set into five ranges using the 5%, 15%, 85%, and 95% percentiles. We refer to the ranges as strongly under-exposed (SU), under-exposed (U), middle-exposed (M), over-exposed (O), and strongly over-exposed (SO). Figure 4 shows example images for each brightness category and demographic, illustrating that the brightness categories are broadly meaningful. Note that under- and over-exposure has strong demographic correlation. African-American females and males have substantial numbers of image pairs with (U,U) and (SU,SU) brightness, but very few with (O,O) or (SO,SO). The reverse applies to Caucasian females and males, but have substantial numbers of (O,O) and (SO,SO) but very few (U,U) or (SU,SU).

5. Accuracy Varies By Image Pair Brightness

To check if accuracy variation based on image pair brightness is matcher-dependent, we use two matchers: ArcFace [19, 26], and a recent commercial matcher, COTS-D (not named due to license restriction). The pattern of results across brightness categories and demographics is similar for the two matchers. Because the two metrics give similar results and FSB is less affected by non-skin elements, this section presents results only for FSB-based categories.

For each matcher, we compute the impostor distribution separately for each demographic, and then select the threshold corresponding to a 1-in-10,000 FMR for the Caucasian male demographic as the threshold for all demographics. This method of setting the FMR threshold follows the recent NIST report on demographic effects in face recognition accuracy [25]. Also, this method makes the cross-demographic differences in FMR more readily apparent.

Table 1 summarizes the FMRs across the brightness categories, for both matchers and the four demographic groups. There are five categories of similar-brightness impostor pairs: (SU,SU), (U,U), (M,M), (O,O) and (SO-SO); four categories of pairs with images one brightness category apart; (SU,U), (U,M), (M,O) and (O,SO); and so on. The number of image pairs varies greatly across brightness category and demographic. Categories with less than 1M image pairs are shown in gray in Table 1, to indicate lower statistical support. Our comments focus primarily on categories with larger numbers of image pairs. We comment first on how FMR varies with brightness of the image pairs, then on FMR differences across demographic groups, and last on differences between matchers.

One general result for both matchers and for all four

FSB	C Female		C Male		A-A Female		A-A Male	
	Imposters	FMR	Imposters	FMR	Imposters	FMR	Imposters	FMR
(SU,SU)	90	0.0000	1,650	0.0606	1,329,858	0.5188	10,860,733	0.0462
		0.0000		0.0000		0.0285		0.0051
(U,U)	1,222	0.0818	22,938	0.0000	5,154,158	0.4872	42,810,922	0.0440
		0.0000		0.0000		0.0339		0.0046
(M,M)	21,922,011	0.0359	273,558,170	0.0096	176,625,162	0.3273	812,147,804	0.0374
		0.0528		0.0096		0.0263		0.0034
(O,O)	3,851,462	0.0550	24,685,360	0.0121	582,234	0.2987	1,703,422	0.0919
		0.0698		0.0123		0.0477		0.0067
(SO,SO)	1,087,394	0.1014	10,487,792	0.0255	9,020	0.4324	14,986	0.3136
		0.0921		0.0166		0.1441		0.0267
(SU,U)	699	0.0000	12,447	0.0161	5,238,898	0.4919	43,133,866	0.0438
		0.0000		0.0080		0.0030		0.0048
(U,M)	331,078	0.0166	5,027,776	0.0077	60,364,176	0.3521	372,979,809	0.0338
		0.0099		0.0017		0.0241		0.0034
(M,O)	18,385,534	0.0386	164,378,978	0.0094	20,297,481	0.2047	74,433,988	0.0366
		0.0553		0.0099		0.0253		0.0033
(O,SO)	4,096,167	0.0575	32,189,078	0.0137	145,656	0.2842	321,113	0.1068
		0.0754		0.0126		0.0665		0.0106
(SU,M)	92,668	0.0183	1,356,380	0.0071	30,671,103	0.3341	187,881,881	0.0309
		0.0053		0.0010		0.0198		0.0031
(U,O)	138,810	0.0166	1,510,801	0.0052	3,468,769	0.1362	17,093,291	0.0168
		0.0079		0.0019		0.0129		0.0016
(M,SO)	9,773,081	0.0372	107,153,315	0.0091	2,537,284	0.1123	7,012,583	0.0216
		0.0524		0.0098		0.0227		0.0026

Table 1. FMR across impostor pair brightness categories for all demographic groups. For the FMR of each category, top number is ArcFace model, bottom number is COTS-D. The categories with less than 1M impostors are set in gray.

demographics is that (M,M) brightness image pairs have lower FMR compared to either (U,U) and (SU,SU) pairs, or to (O,O) and (SO,SO) pairs. For example, for Caucasian male, going from brightness category (M,M) to (SO,SO) increases ArcFace FMR from 0.0096 to 0.0255 (166%), and increases COTS-D FMR from 0.0096 to 0.0166 (73%). For Caucasian female, going from brightness (M,M) to (SO,SO) increases ArcFace FMR from 0.0359 to 0.1014 (183%), and increases COTS-D FMR from 0.0528 to 0.0921 (74%). For African-American male, going from brightness (M,M) to (SU,SU) increases ArcFace FMR from 0.0374 to 0.0462 (24%), and increases COTS-D FMR from 0.0034 to 0.0051 (50%). For African-American female, going from brightness (M,M) to (U,U) increases ArcFace FMR from 0.3273 to 0.4872 (49%), and increases COTS-D FMR from 0.0263 to 0.0339 (29%). The (SU,U) brightness pairs show similar increases in FMR for the African-American demographics, as do the (O,SO) brightness pairs for the Caucasian demographics. In general, similar brightness image pairs that are too dark or too bright have increased FMR.

Another general result is that image pairs that differ significantly in brightness have a *decreased* FMR. For ex-

ample, for Caucasian male, (U,O) pairs have an ArcFace FMR of 0.0052 and a COTS-D FMR of 0.0019, *decreases* of 45% and 80%, respectively, from the FMR for (M,M) pairs. For African-American male, (U,O) pairs have an FMR of 0.0168 and a COTS-D FMR of 0.0016, *decreases* of 55% and 53%, respectively, from the FMR for (M,M) pairs. For African-American female, (U,O) pairs have an FMR of 0.0168 and a COTS-D FMR of 0.0016, *decreases* of 58% and 51% from the FMR for (M,M) pairs.

The larger effects of image pair brightness can be seen in the impostor and genuine distributions in Figure 6. The distributions in the left plot show that image pairs that are strongly under- or over-exposed have an impostor distribution that is shifted to increased similarity scores, meaning a higher FMR, and also a lower d-prime for separation of impostor and genuine distributions. The distributions in the right plot show that image pairs that have strongly different brightness have an impostor distribution shifted to lower similarity, meaning a lower FMR, but also a genuine distribution shifted even further to lower similarity, meaning a higher FNMR and resulting in a decreased d-prime. This pattern of accuracy variation is reinforced by the impostor




Figure 5. Brightness distribution of demographic groups. For each brightness metric, it shows the brightness distributions of Caucasian females, Caucasian males, African-American females, and African-American males. The stacked version is in the Figure 1 of the Supplemental material.

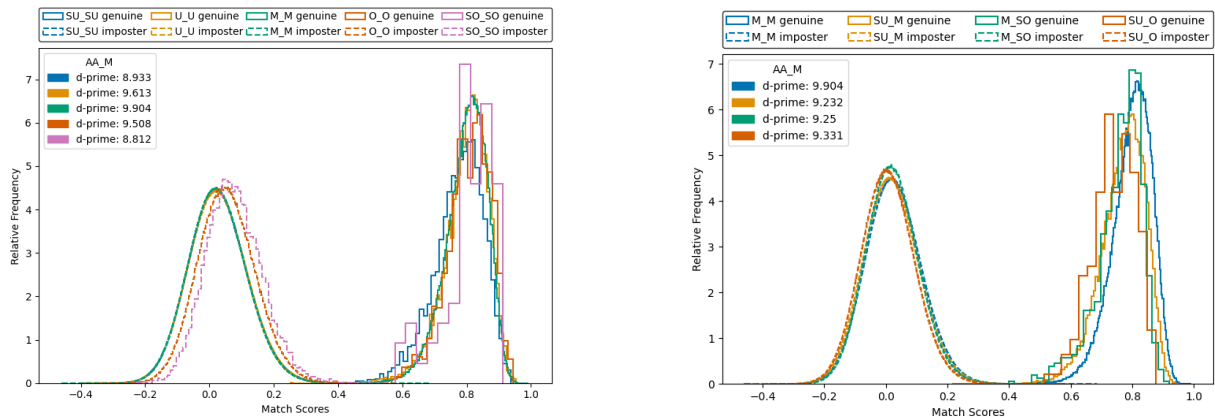


Figure 6. Genuine and impostor distributions for selected pair brightness categories for African-American Male. The African American male demographic has the most data over the broadest range of brightness categories, enabling better estimates of the genuine distribution than the other demographics. Distributions in the plot on left show that too-dark or too-bright image pairs have both worse impostor and worse genuine distributions. Distributions in the plot on the right show that image pairs with strong brightness difference have better impostor and worse genuine distributions.

and genuine distributions for other demographics shown in Figure 7. Image pairs of extreme low or high brightness result in an impostor distribution shifted to higher similarity, and image pairs of strongly different brightness result in a genuine distribution shifted to lower similarity.

FMR variation across demographics in our results mostly follow the pattern seen in the literature [25, 33]. Considering (M,M) brightness pairs, the Caucasian female ArcFace FMR is a factor of 3.7 times that for Caucasian male, and the COTS-D FMR is a factor of 5.5 times that for Caucasian male. For African-American male, the ArcFace FMR is a factor of 3.9 that for Caucasian male, but

the COTS-D FMR is 65% lower than that for Caucasian male. For African-American female, the ArcFace FMR is a factor of 34 times that for Caucasian male, and the COTS-D FMR is a factor of 2.7 times that for Caucasian male. For ArcFace, using a fixed similarity threshold to generate a FMR results in Caucasian male having the lowest FMR, similar and higher FMRs for Caucasian female and African-American male, and much higher FMR for African-American female. We will comment on the relative d-prime separation of impostor and genuine scores later.

The relative COTS-D FMR for Caucasian male versus African-American male is clearly at odds with what is seen




Figure 7. Selected impostor and genuine distributions for African American female, Caucasian male, and Caucasian female.

FSB	C Female			C Male			A-A Female			A-A Male		
	BIM	FMR	d'	BIM	FMR	d'	BIM	FMR	d'	BIM	FMR	d'
(SU,SU)	23.72	-	-	23.19	-	-	24.52	0.5188	7.5	24.49	0.0462	8.9
(U,U)	26.70	-	-	27.05	-	-	29.61	0.4872	7.9	29.69	0.044	9.6
(M,M)	35.98	0.0359	8.3	34.69	0.0096	8.9	34.53	0.3273	8.2	34.97	0.0374	9.9
(O,O)	34.87	0.0550	8.1	32.74	0.0121	8.9	33.16	-	-	33.65	0.0919	9.5
(SO,SO)	26.41	0.1014	7.9	25.26	0.0255	8.7	27.23	-	-	26.77	-	-

Table 2. Bright Information Metric for images in each brightness category with FMR.

in the literature. The black-box nature of a COTS algorithm means that we cannot offer any confident speculation of the cause for this. But we observe that there is no overlap between MORPH and the dataset used in training ArcFace, whereas the training and development data for COTS-D is unknown to us. We also observe that the COTS-D impostor and genuine distributions are distinctly not Gaussian. Similarity scores from COTS-D are in the range from 0 to 1. Across the four demographics, from 58% to 80% of the impostor scores are reported as 0, and 98% or more of the genuine scores are reported as 1. This suggests a thresholding step on the raw similarity scores before they are reported. As a result, it is meaningless to plot impostor and genuine distributions or to report d-prime values for COTS-D.

General conclusions from this analysis are as follows. One, there is target range in the middle of the observed brightness distribution that results in lowest FMR. Two, image pairs with similar brightness, darker or brighter than the target range, result in increased FMR. Three, image pairs with substantially different brightness result in decreased FMR. Thus the current standard approach in which the same matching threshold is used for all image pairs, results in some pairs having predictably higher and lower FMR based on their brightness. For controlled image acquisition, this problem can potentially be addressed by adjusting lighting to acquire images in the target range of face skin brightness.

6. Brightness, Face Information and d-prime

Brightness metrics in previous work [10, 35, 39] suggest that an image with higher brightness is of higher quality. (See Related Work.) However, the previous section shows that the FMR of similar-brightness impostor pairs is low-

est for middle brightness and increases for pairs that are too dark or too bright. This section explores how the level of variability in face skin brightness can explain the accuracy across brightness categories. We propose a Brightness Information Metric (BIM) that captures the degree of variation in face skin brightness, defined as follows:

$$BIM = \sum_{i=0}^N |B_i - \bar{B}| P(B_i) \quad (2)$$

where B_i is a brightness level, $P(B_i)$ is the probability of face skin pixels that have this brightness level, and \bar{B} is the average brightness of the face skin pixels. Note that, since the brightness values distribute in a Gaussian-like form, the BIM is positively related to the average information entropy of the selected group of images.

Table 2 summarizes the average BIM across the different brightness categories and demographics, along with FMR and the d-prime for the separation between impostor and genuine distributions. ArcFace d-prime and FMR values are shown because, as explained earlier, the COTS-D impostor and genuine distributions contain large spikes at zero and one, respectively. For each of the four demographics, the largest BIM and the lowest FMR occur for M brightness. For all demographics, the BIM decreases from (M, M) to (O,O) and decreases again to (SO,SO), and it also decreases from (M,M) to (U,U) and decreases again to (SU,SU). Across brightness categories and demographics, a lower BIM for images in an impostor pair means less information to distinguish between individuals, and so a higher similarity score and higher FMR.

FSB Sliding-bins	Range	C Female		C Male		A-A Female		A-A Male	
		BIM	d'	BIM	d'	BIM	d'	BIM	d'
(M6,M6)	145 - 185	35.42	8.40	34.38	9.15	36.00	8.36	36.70	9.93
(M7,M7)	150 - 190	35.90	8.47	34.73	9.17	36.22	8.40	36.97	9.94
(M8,M8)	155 - 195	36.30	8.55	34.99	9.20	36.35	8.40	37.09	10.00
(M9,M9)	160 - 200	36.62	8.59	35.18	9.20	36.39	8.42	37.11	10.02
(M10,M10)	165 - 205	36.85	8.54	35.29	9.18	36.29	8.44	36.96	9.99
(M11,M11)	170 - 210	36.83	8.47	35.29	9.09	36.01	8.47	36.70	9.97
(M12,M12)	175 - 215	36.66	8.49	35.11	9.01	35.69	8.47	36.34	9.96
(M13,M13)	180 - 220	36.24	8.31	34.63	8.97	35.20	8.51	35.76	9.88

Table 3. Average BIM and d-prime of each sliding-bin.

6.1. Target Brightness Range for Best Accuracy

The M brightness range in earlier sections corresponds to the middle 70% of the overall brightness distribution. Our results so far suggest it should be possible to select a narrower range that gives better accuracy. To investigate the feasibility and generality of this, we conduct a refined analysis of our M brightness range.

Table 3 summarizes BIM and d-prime for image pairs in sliding 40-brightness-level bands of our initial M brightness range. Across all four demographics, the highest BIM and d-prime are generally found for M9 and M10, with combined brightness range 160-205. The operational scenario brightness range in Figure 5 runs about 50 to 250, meaning that the target range for best accuracy is about one-fourth the observed operational range. Also, note that 28% of African-American female images, 24% of African-American male images, 62% of Caucasian male images and 59% of Caucasian female images in MORPH fall into the target range for best accuracy. Thus it appears that a portion of the demographic variation in accuracy is due to variation in quality of the images acquired.

7. Conclusions and Discussion

Using a face image dataset from a real operational scenario with controlled image acquisition, we show that the distribution of face skin brightness varies over most of the 0-255 range. While the distribution of face skin brightness has large overlap across demographics, on average, images of African-Americans have lower brightness than images of Caucasians and images of males have (slightly) lower brightness than images of females. Also, the standard deviation of the brightness metric is about one-third greater for African-Americans than for Caucasians.

To analyze how accuracy varies with image pair brightness, the face skin brightness distribution was divided into intervals representing strongly under-exposed, under-exposed, middle exposure, over-exposed and strongly over-exposed. FMR analysis from two different matchers showed very similar trends. Pairs of images that have lower

or higher face skin brightness have an increased FMR relative to image pairs with middle brightness. Also, pairs of images that differ strongly in face skin brightness have a lower FMR than image pairs with middle brightness. We used a brightness information metric that measures the variation in brightness in the face skin region to explain that the available information in the face skin region is at a maximum at a middle brightness level and declines as the face skin brightness moves toward either end of the scale.

We showed that it is possible to find a target face skin brightness range that represents the maximum BIM for all four demographics, and achieves the maximum d-prime separation of impostor and genuine distributions. This target brightness range is about one-fourth the range of brightness seen in a retrospective image dataset from an operational scenario with controlled image acquisition. This analysis suggests that operational scenarios with controlled image acquisition could maximize face recognition accuracy by adjusting the lighting on image acquisition to hit the target face skin brightness range. The analogy could be made here to image acquisition for commercial iris recognition, in which the illumination is controlled and image quality is automatically checked at the time of acquisition.

Analysis suggests that the issue of acquiring images with face skin brightness in an appropriate range for all demographics is related to the problem of varying accuracy across demographics. Consider that our analysis suggests a target face skin brightness range of 160-205. The face skin brightness distribution shows that different demographics have a very different fraction of their images falling into this range. Caucasian males and females have 62% and 59% of their images falling into this interval. But African-American males and females have just 24% and 28% of their images falling into this interval. Acquiring images of equal quality for all demographics may be an essential element to expecting equal accuracy across demographics.

References

- [1] <https://github.com/zllrunning/face-parsing.PyTorch>, last accessed on February 2021. 3

- [2] Iface sdk. <https://www.innovatrics.com/iface-sdk/>. 3
- [3] Iface sdk. <https://www.fulcrumbiometrics.com/Innovatrics-SDKs-s/206.htm>. 3
- [4] Innovatrics dot product documentation. <https://developers.innovatrics.com/digital-onboarding/technical/remote/core-server/latest/documentation/>. 3
- [5] Iso/iec 19794-5:2011 information technology – biometric data interchange formats – part 5: Face image data. <https://www.iso.org/standard/50867.html>. 3
- [6] Morph dataset. https://www.faceaginggroup.com/?page_id=1414. 3
- [7] Ayman Abaza, Mary Ann Harrison, and Thirimachos Bourlai. Quality metrics for practical face recognition. In *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*, pages 3103–3107. IEEE, 2012. 3
- [8] Ayman Abaza, Mary Ann Harrison, Thirimachos Bourlai, and Arun Ross. Design and evaluation of photometric image quality measures for effective face recognition. *IET Biometrics*, 2014. 2
- [9] Ayman Abaza, Mary Ann Harrison, Thirimachos Bourlai, and Arun Ross. Design and evaluation of photometric image quality measures for effective face recognition. *IET Biometrics*, 3(4):314–324, 2014. 3
- [10] Mohamed Abdel-Mottaleb and Mohammad H. Mahoor. Application notes - algorithms for assessing the quality of facial images. *IEEE Computational Intelligence Magazine*, 2(2):10–17, 2007. 7
- [11] Salem Hamed Abdurrahim, Salina Abdul Samad, and Aqilah Baseri Huddin. Review on the effects of age, gender, and race demographics on automatic face recognition. *Visual Computer*, 34(11):1617–1630, 2018. 3
- [12] Vítor Albiero and Kevin W Bowyer. Is face recognition sexist? no, gendered hairstyles and biology are. *British Machine Vision Conference (BMVC)*, 2020. 1
- [13] Vítor Albiero, Xingyu Chen, Xi Yin, Guan Pang, and Tal Hassner. img2pose: Face alignment and detection via 6dof, face pose estimation. In *Proc. Conf. Comput. Vision Pattern Recognition*, 2021. 3
- [14] Vítor Albiero, Kai Zhang, Michael C. King, and Kevin W. Bowyer. Gendered differences in face recognition accuracy explained by hairstyles, makeup, and facial morphology. *IEEE Transactions on Information Forensics and Security*, 17:127–137, 2022. 3
- [15] J. Ross Beveridge, David S. Bolme, Bruce A. Draper, Geof H. Givens, Yui Man Lui, and P. Jonathon Phillips. Quantifying how lighting and focus affect face recognition performance. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, pages 74–81, 2010. 2
- [16] Aman Bhatta, Vítor Albiero, Kevin W Bowyer, and Michael C King. The gender gap in face recognition accuracy is a hairy problem. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 303–312, 2023. 1
- [17] 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*, 3(1):101–111, 2021. 1
- [18] 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*, 1(1):32–41, 2019. 1, 3
- [19] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 1, 4
- [20] Pawel Drozdowski, Christian Rathgeb, and Christoph Busch. The watchlist imbalance effect in biometric face identification: Comparing theoretical estimates and empiric measurements. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3757–3765, October 2021. 3
- [21] Pawel Drozdowski, Christian Rathgeb, Antitza Dantcheva, Naser Damer, and Christoph Busch. Demographic bias in biometrics: A survey on an emerging challenge. *IEEE Transactions on Technology and Society*, 1(2), 2020. 1
- [22] Klemen Grm, Vitomir Štruc, Anais Artiges, Matthieu Caron, and Hazım K Ekenel. Strengths and weaknesses of deep learning models for face recognition against image degradations. *IET Biometrics*, 7(1):81–89, 2018. 2
- [23] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010. 2
- [24] Patrick Grother. NISTIR 8429: Face recognition vendor test (frvt) part 8: Summarizing demographic differentials. Technical report. 2
- [25] Patrick Grother, Mei Ngan, and Kayee Hanaoka. NISTIR 8280: Ongoing face recognition vendor test (frvt) part 3: Demographic effects. Technical report. 1, 4, 6
- [26] Jia Guo. Insightface: 2D and 3D face analysis project. <https://github.com/deepinsight/insightface>, last accessed on February 2021. 4
- [27] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 87–102, 2016. 3
- [28] Hu Han, Shiguang Shan, Xilin Chen, and Wen Gao. A comparative study on illumination preprocessing in face recognition. *Pattern Recognition*, 46(6):1691–1699, 2013. 2
- [29] T. Hoggins. ‘racist and sexist’ facial recognition cameras could lead to false arrests, Dec. 20 2019. <https://www.telegraph.co.uk/technology/2019/12/20/racist-sexist-facial-recognition-cameras-could-lead-false-arrests/>. 1
- [30] 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. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(4):550–560, 2021. 1

- [31] 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 07-49, University of Massachusetts, Amherst, October 2007. 2
- [32] Yu-Hsuan Huang and Homer H Chen. Deep face recognition for dim images. *Pattern Recognition*, page 108580, 2022. 2
- [33] KS Krishnapriya, Vitor 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*, 1(1):8–20, 2020. 1, 2, 3, 6
- [34] S. Lohr. Facial recognition is accurate, if you’re a white guy. *The New York Times*, Feb. 9 2018. 1
- [35] Jianfeng Long and Shutao Li. Near infrared face image quality assessment system of video sequences. In *International Conference on Image and Graphics*, pages 275–279. IEEE, 2011. 2, 7
- [36] Boyu Lu, Jun-Cheng Chen, Carlos D. Castillo, and Rama Chellappa. An experimental evaluation of covariates effects on unconstrained face verification. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 1(1):42–55, 2019. 1, 2
- [37] D. Maltoni, A. Franco, M. Ferrara, D. Maio, and A. Nardelli. Biolab-icao: A new benchmark to evaluate applications assessing face image compliance to iso/iec 19794-5 standard. In *IEEE Int. Conf. on Image Processing*, 2009. 3
- [38] Maria De Marsico, Michele Nappi, and Daniel Riccio. Measuring measures for face sample quality. In *ACM Workshop on Multimedia in Forensics and Intelligence*, 2011. 2
- [39] K Nasrollahi and Thomas B Moeslund. Complete face logs for video sequences using face quality measures. *IET signal processing*, 3(4):289–300, 2009. 2, 7
- [40] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. 2015. 3
- [41] 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 (FGR06)*, pages 341–345. IEEE, 2006. 2, 3
- [42] R. Rothe, R. Timofte, and L. Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *Int. J. Comput. Vision*, 126:144–157, 2018. 3
- [43] E. Santow. Emerging from AI utopia. *Science*, April 2020. 1
- [44] Torsten Schlett, Christian Rathgeb, Olaf Henniger, Javier Galbally, Julian Fierrez, and Christoph Busch. Face image quality assessment: A literature survey. *ACM Computing Surveys*, 2021. 2
- [45] Terence Sim, Simon Baker, and Maan Bsat. The cmu pose, illumination, and expression database. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1615–1618, 2003. 2
- [46] Philipp Terhörst, Marco Huber, Naser Damer, Florian Kirchbuchner, Kiran Raja, and Arjan Kuijper. Pixel-level face image quality assessment for explainable face recognition. *arXiv preprint arXiv:2110.11001*, 2021. 2
- [47] 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. *IEEE Transactions on Technology and Society*, pages 1–1, 2021. 1
- [48] J. Vincent. Gender and racial bias found in amazon’s facial recognition technology (again). *The Verge*, Jan. 25 2019. <https://www.theverge.com/2019/1/25/18197137/amazon-rekognition-facial-recognition-bias-race-gender>. 1
- [49] ISO/IEC JTC1 SC17 WG3. Portrait quality (reference facial images for mrtD). Technical report, International Civil Aviation Organization, 2018. 3
- [50] Cameron Whitelam, Emma Taborsky, Austin Blanton, Brianna Maze, Jocelyn Adams, Tim Miller, Nathan Kalka, Anil K. Jain, James A. Duncan, Kristen Allen, Jordan Cheney, and Patrick Grother. IARPA Janus Benchmark-B Face Dataset. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2017-July:592–600, 2017. 2
- [51] Haiyu Wu, Grace Bezold, Aman Bhatta, and Kevin W Bowyer. Logical consistency and greater descriptive power for facial hair attribute learning. *arXiv preprint arXiv:2302.11102*, 2023. 1
- [52] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 325–341, 2018. 3
- [53] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Jiwen Lu, Dalong Du, et al. Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10492–10502, 2021. 3