

# Investigation of accuracy and bias in face recognition trained with synthetic data

Pavel Korshunov, Ketan Kotwal, Christophe Ecabert, Vidit  
Amir Mohammadi, and Sébastien Marcel  
Idiap Research Institute, Martigny, Switzerland  
{firstname.lastname}@idiap.ch

## Abstract

*Synthetic data has emerged as a promising alternative for training face recognition (FR) models, offering advantages in scalability, privacy compliance, and potential for bias mitigation. However, critical questions remain on whether both high accuracy and fairness can be achieved with synthetic data. In this work, we evaluate the impact of synthetic data on bias and performance of FR systems. We generate balanced face dataset, FairFaceGen, using two state of the art text-to-image generators, Flux.1-dev and Stable Diffusion v3.5 (SD35), and combine them with several identity augmentation methods, including Arc2Face and four IP-Adapters. By maintaining equal identity count across synthetic and real datasets, we ensure fair comparisons when evaluating FR performance on standard (LFW, AgeDB-30, etc.) and challenging IJB-B/C benchmarks and FR bias on Racial Faces in-the-Wild (RFW) dataset. Our results demonstrate that although synthetic data still lags behind the real datasets in the generalization on IJB-B/C, demographically balanced synthetic datasets, especially those generated with SD35, show potential for bias mitigation. We also observe that the number and quality of intra-class augmentations significantly affect FR accuracy and fairness. These findings provide practical guidelines for constructing fairer FR systems using synthetic data.*

## 1. Introduction

The use of synthetic data to train face recognition (FR) models has gained increasing attention in recent years, primarily, due to its potential to avoid ethical, legal, and licensing challenges associated with using real facial images, especially in the context of privacy regulations such as the GDPR [1]. Synthetic data offers the potential to generate large-scale datasets to train models for commercial use without infringing on individual privacy, hence, facilitating the development of safer FR systems. Moreover, it allows a more fine-grained control over the data generation, which potentially can help mitigating biases in FR systems.




Figure 1. Examples of images cropped to  $112 \times 112$  from WebFace42M [2] and synthetic datasets Digi2Real [3], HyperFace [4], DCFace [5], and generated by us with SD35 and Flux.

The main focus of the recent work is on generating synthetic face datasets that could be used to train FR models with performance approaching that of models trained on real data [1, 3–8]. However, several issues are still missing from the current research discourse pertaining to training FR models with synthetic data:

- **Dual-generator framework:** Synthetic face data generation often employs a two-stage process: a *seed generator* for creation of distinct identities and an *augmentation generator* for producing intra-class variations such as different poses, lighting conditions, and expressions. Although considerable effort is directed toward enhancing the diversity of seed identities [3, 4], the role of augmentation generators in influencing FR performance remains underexplored.

- **Unfair dataset comparisons:** Comparative studies between synthetic and real datasets frequently suffer from inconsistencies in dataset sizes and compositions. A synthetic dataset with 10 K identities and 64 images per identity gets compared with another dataset of 50 k identities and 20 images per identity or to a WebFace-12M’s 1.5 M identities and 12 M images [2]. In these comparisons, the number of identities are different and the total number of images are different, which are the parameters affecting the performance of the model. Such mismatches conceal the true effect of generation parameters [9] and a fair evaluation needs to acknowledge the issue and equalize either (i) the number of identities and total images or (ii) the images-per-identity ratio.

- **Dependence on real data:** Most synthetic face generation methods, including generative adversarial networks (GANs) and diffusion-based models, are trained on datasets of real facial images [10]. Many generative methods also rely on an existing FR models pre-trained on large datasets of real faces. Therefore, the generated synthetic data inherently reflects the biases present in the original datasets [11]. This dependency questions the purity of synthetic datasets and their efficacy in mitigating bias.

- **Lack of bias evaluation:** Synthetic data generation allows, at least in principle, for controlled manipulation of demographic attributes (e.g., race, gender, age), enabling balanced training data that may improve fairness [12, 13]. However, comprehensive evaluations of bias in FR models trained solely on synthetic data are scarce. Existing studies suggest that synthetic data can mirror the biases of their training datasets, emphasizing the need for deliberate balancing strategies [11].

In this work, we aim to systematically investigate the impact of synthetic data generation on the performance and bias of FR models. Specifically, we:

1. Evaluate two state-of-the-art text-to-image models, Flux.1-dev [14], referred to as Flux in this paper, and Stable Diffusion v3.5 [15], referred to as SD35, for generation of different identities to investigate the impact of the *seed generators* on the resulting FR model.

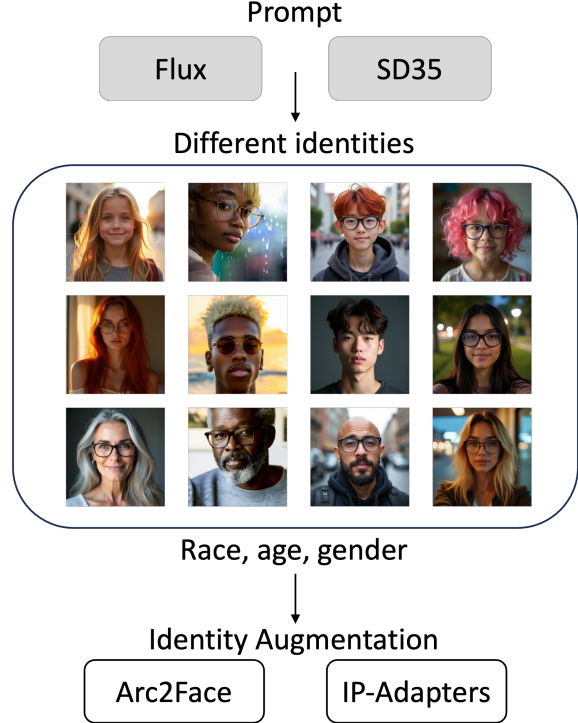


Figure 2. Generation process of *FairFaceGen* dataset.

2. Use prompt-based techniques with Flux and SD35 to generate balanced synthetic dataset, *FairFaceGen*<sup>1</sup> (see Figure 2 for the illustration), across gender, age (categorized into seven groups from children to seniors), and race (White, Black, Asian, and Hispanic). This approach to generating balanced data does not require special model training or tuning, which, if effective, makes it very practical and simple.
3. Compare several augmentation generators, including Arc2Face [16] pre-trained on WebFace-42M and four face IP-Adapter variants [17], with Stable Diffusion v1.5 (SD15) or Stable Diffusion XL (SDXL) models relying on the embeddings from either InsightFace [18] or CLIP [19], pretrained on LAION-2B [20] and COYO-700M [21] datasets.
4. Measure inter-identity and intra-identity similarity with an independent EdgeFace [22] face recognition model and compare different variations of the synthetic data with HyperFace [4], Digi2Real [3] and a size-matched 10K identities subset from the WebFace-42M, which is a cleaned version of WebFace-260M [2].
5. Besides using standard benchmarking datasets, including LFW [23] with its variants, AgeDB-30 [24]

<sup>1</sup><https://www.idiap.ch/paper/fairfacegen>

and others, we report verification performance on the challenging IJB-B/C [25, 26] and bias on Racial Faces in-the-Wild (RFW) [27] datasets, comparing our balanced sets to synthetic datasets, DCFace [5], Digi2Real [3], HyperFace [4], and to real datasets, CASIA-WebFace [28] and a 10K identities subset of WebFace 42M [2].

In this paper, our objectives are i) to provide quantitative guidelines on choosing *seed* and *augmentation generators*, ii) to investigate the impact of the number of images per identity on the FR model performance, and iii) to evaluate how the demographically balanced synthetic datasets mitigate bias.

## 2. Related work

**Synthetic identities for face recognition** Early work showed that high-fidelity GANs such as StyleGAN2 [29] can support FR pre-training provided sufficient inter- and intra-class diversity. Diffusion-based methods (e.g. DCFace [5]) improved identity preservation, but these and subsequent works [3, 4, 8] still combine *seed* identity creation and *augmentation* of each identity. The impact of the latter stage on FR accuracy remains under-explored.

**Augmentation generators and identity fidelity** Arc2Face [16] and face IP-Adapter variants [17] demonstrated realistic pose and expression changes without label flips. A systematic benchmark that fixes the backbone while varying only the augmentation generator is however missing. In this paper, we provide such an analysis and compare different augmentation generators.

**Balanced synthetic data for fairness** Demographic bias in FR is well documented [13, 27, 30]. Several recent studies leverage synthetic data for bias mitigation. For instance, Qraitem *et al.* [12] proposed a two-step strategy “From Fake to Real” (FTR) that first pre-trains on a *balanced* synthetic set and then fine-tunes on real data to reduce spurious correlations in general image recognition tasks. However, FTR is agnostic to facial identities and does not study FR-specific synthetic generation or the trade-off between intra- and inter-class variations. Focusing specifically on face recognition, Melzi *et al.* fine-tuned existing FR models with demographically balanced synthetic data produced by GANDiffFace and reported reductions in racial bias in RFW [31]. However, they did not equalize dataset scale when comparing to real-data baselines and did not evaluate on the challenging IJB-B/C benchmark. The authors of VariFace [8] claimed that they generated balanced synthetic dataset which is close in performance to the real data and they equalized all dataset sizes when comparing the performance, however, they do not show the results on IJB-B/C,

and they compared with CASIA WebFace, which, despite being a real dataset, leads to poor generalization capabilities [3].

Initiatives such as the FRSyn competition series have helped progress in the development and use of synthetic datasets [32, 33]. Despite these advancements, FR systems trained only on synthetic data continue to exhibit inferior performance compared to those trained on real-world datasets of similar size [3, 34]. This performance disparity is evident not only in recognition metrics—such as False Match Rate (FMR) and False Non-Match Rate (FNMR)—but also in fairness across demographic groups, as indicated by variability in group-specific outcomes. In [11], Huber *et al.* noted that FR models trained on synthetic dataset are likely to amplify demographic disparities, compared to the models trained on real datasets.

**Scale-controlled comparisons** Performance gains reported for synthetic data are often confounded by unmatched sample counts. Datasets such as WebFace-42M or even WebFace260M [2] are much larger than most synthetic datasets and therefore always demonstrate higher performance when used to train FR models. We therefore enforce a similar number of identities *and* number of images-per-identity across all training data, making our experimental results comparable in a fairer way, an approach largely missing from the existing works. Only a recent work by Yeung *et al.* [8] compared FR models trained on the same number of IDs and images per ID. However, the paper only used standard small benchmarking datasets for evaluation (LFW [23], AgeDB-30 [24], and others), missing the results on IJB-B/C [25, 26], which shows an incomplete picture with respect to the generalizability of the FR models trained on synthetic data.

## 3. Datasets

In this section, we provide details on how we have generated synthetic faces balanced by gender, race, and age that we used to train AdaFace [35] face recognition model. We then describe the existing datasets that we use as baselines to compare our generated data for the recognition accuracy. We also describe the test datasets used to benchmark the trained FR models.

### 3.1. Balanced synthetic data generation

To investigate the effect of *seed generator* on face recognition training, we use the prompt based technique to generate initial identities with Flux.1-Dev [14] (denoted as Flux) and Stable Diffusion v3.5 [15] (denoted as SD35). Figure 3a shows an image example generated by Flux in the top row and one generated by SD35 in the bottom row. We focused on generating balanced facial data by varying






Figure 3. Seed generators SD35 (top of (a)) and Flux (bottom of (a)) and their augmentations generated with Arc2Face and IP-Adapters.

each of these prompt attributes: seven age groups (below 8 years old, between ages 8 and 13, between ages 13 and 18, between ages 18 and 25, between ages 25 and 35, between ages 35 and 50 ranging, and older than 50 years), four races (Caucasian White, Black, East-Asian, and Latino Hispanic), and two genders (male and female). Given that there are 56 variations of these attributes, with 200 for each variation, we have generated a total of 11200 different faces for each generator. When generating the prompts, we also randomly varied other soft-attributes such as hairstyle, hair-color, glasses (no glasses, reading glasses or sunglasses), presence of earrings, lighting conditions, i.e., daylight, sunset, indoor lighting, low light, and raining, backgrounds, e.g., beach, street, or studio, photos styles, e.g., portrait or artistic, quality of the generated images, such as compressed, professional, or amateur; and camera types, including professional DSLR, selfie webcam, or smartphone. Since these soft environmental attributes were chosen at random, they are equally present in all images regardless of race or age. Note that exactly the same prompts were used for both Flux and SD35, so that the differences in the resulting images are dictated solely by the differences in these text-to-image models.

To understand the impact of *augmentation generator* on face recognition training, we used these publicly available models trained to generate several variations of the same input identity: Arc2Face [16] which uses face embeddings from ArcFace model [36] pretrained on WebFace42M to enforce ID consistency when generating the variations with Stable Diffusion v1.5 (SD15) and IP-Adapter variants [17], which use either a face embedding from InsightFace [18] or from CLIP [19] model for ID consistency when using SD15 and Stable Diffusion XL (SDXL) backbones. Figure 3 shows SD35-based augmentations at the top and for Flux at the bottom rows. For each image generated from a

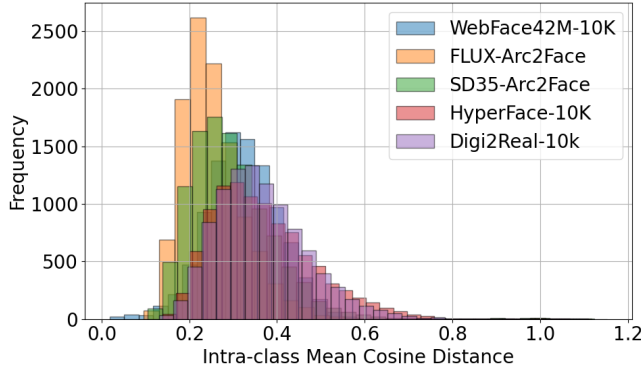
*seed generator*, which represents an identity, we generated the number of augmentations with these *augmentation generators* ranging from 8 images per identity, to 32 images per identity with a step of 8. We vary the number of images per identity to understand whether and how this number affects the performance of face recognition.

### 3.2. State of the Art datasets

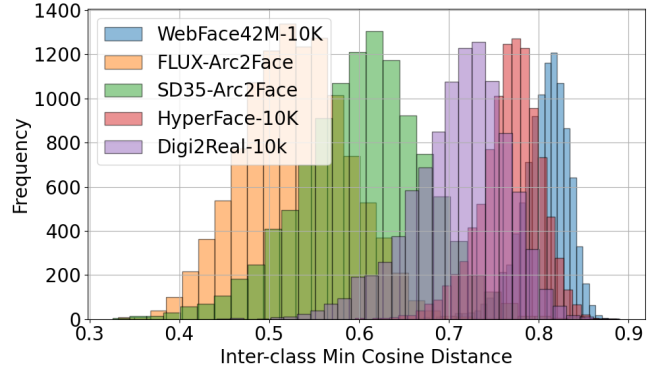
We utilized publicly available state-of-the-art synthetic datasets such as DCFace [5], Digi2Real [3], and HyperFace [4] as baselines to compare the performance of face recognition models trained on them. Since our synthetic datasets (see Section 3.1) have about 11K identities (but this number can reduce after face detection), we have limited all baseline datasets to about 10K identities, so that the accuracy of the face recognition model trained on all the data can be fairly compared. Similarly, we have used real datasets with about 10K identities, such as widely adopted CASIA-WebFace [28] dataset with about 10.5K identities and we took 10K identities subset from WebFace42M [2].

### 3.3. Benchmarking datasets

We evaluate face recognition models trained on the above synthetic and real data using the *standard* benchmark datasets, including Labeled Faces in the Wild (LFW) [23], Cross-Age LFW (CA-LFW) [37], Cross-Pose LFW (CP-LFW) [38], Celebrities in Frontal-Profile in the Wild (CFP-FP) [39], and AgeDB-30 [24], as well as, the large public test benchmarks such as IARPA Janus Benchmark-B (IJB-B) [25] and IARPA Janus Benchmark-C (IJB-C) [26]. Evaluating the models on the more challenging IJB-B/C datasets is very important because the models can show very high accuracies on the *standard* datasets but fail to generalize to IJB-B/C; and that is why the results on IJB-B/C are often missing from the state of the art work.



(a) Mean cosine distance within the same identity



(b) Min cosine between different identities

Figure 4. Mean intra-class and min inter-class cosine distances for synthetic datasets and WebFace-10K.

For bias assessment of trained face recognition models, we use the Racial Faces in the Wild (RFW) dataset [27], which has been widely adopted to benchmark racial bias or fairness in FR systems. It consists of four demographic subsets— African, Asian, Caucasian, and Indian— each containing approximately 3,000 unique individuals and 6,000 predefined image pairs curated by the dataset creators. Each subset includes an equal number of genuine (mated) and impostor (non-mated) pairs, offering a balanced and structured framework for evaluating verification performance across racial groups.

### 3.4. Analysis of the datasets

To analyze how well *augmentation generators* preserve the identity of its different variations, i.e., their intra-class ID similarity, we compute, for the same identity, the cosine distances between the EdgeFace [22] embeddings of its augmentations. The mean of these cosine distances for the identity indicates how similar the augmentations between each other. If the mean is close to zero, all augmentations/variations of this ID are similar.

Figure 4a shows the histograms of mean cosine distances computed for Arc2Face augmentation generator and for all identities of Flux or SD35 seed generators and compares them against HyperFace and Digi2Real synthetic datasets and real faces from the 10K IDs subset of WebFace42M. The figure shows that all datasets have similar distributions of mean cosine distances and they are as close to zero as the WebFace distribution, so the augmentation generators used in the synthetic datasets match closely the real data.

To understand how separated the identities generated by both *seed generators* and *augmentation generators* are from each other, i.e., their inter-class ID difference, we have computed the minimum cosine distance between the average EdgeFace embedding of a given identity with the average embedding of every other identity. Essentially, for each identity, we find the most similar looking other identity (ac-

cording to EdgeFace). If all identities are very distinct from each other, their min cosine distance will be large.

Figure 4b shows histograms of the min cosine distances computed for all identities of the same Flux or SD35-based seed generators, HyperFace and Digi2Real synthetic datasets, and WebFace42M real dataset. Figure shows that histograms of the synthetic datasets are to the left of the WebFace42M. It means the inter-class variability is lower for synthetic datasets, indicating that face recognition trained on synthetic data, potentially, should have lower discriminative power and may perform worse than the one trained on real data. The distribution of the HyperFace is the closest to that of WebFace, while the Flux-based data is the furthest away, which means that Flux seed generator may not be the best for training face recognition.

## 4. Face recognition training and evaluation

To have a fair comparison, we have trained the same model ResNet-50 with AdaFace loss<sup>2</sup> on different variations of our synthetic datasets and compared them with the baseline synthetic and real data. We evaluated all the models using the standard benchmark datasets (see Section 3.3) and IJC-B/C datasets for general face recognition assessment and RFW dataset for bias assessments.

### 4.1. Face recognition setup

For all experiments we followed a consistent FR pipeline. Each input image underwent standardized pre-processing that included face detection and facial landmark localization using MTCNN [40], followed by geometric alignment and resizing of the cropped face region to a resolution of  $112 \times 112$  pixels in RGB format that conforms to the input requirements of the ResNet-50 architecture trained using AdaFace loss [35]. On all datasets, the training was done for 30 epochs with batch size 256 and an initial learning rate of 0.1, which was reduced by a factor of 10 at

<sup>2</sup><https://github.com/mk-minchul/AdaFace>

Table 1. Accuracy results of ResNet-50 AdaFace model trained on different real (highlighted green) and synthetic datasets. All accuracy values are reported as %.

Training dataset	# Images per ID	Standard benchmark datasets					IJB-B		IJB-C	
		AgeDB30	CFP-FP	LFW	CP-LFW	CA-LFW	1E-04	1E-03	1E-04	1E-03
WebFace42M-10K [2]	22	91.7	92.8	99.3	89.6	93.3	86.5	92.8	<b>89.5</b>	<b>94.5</b>
CASIA WebFace [28]	46	94.7	95.3	99.5	90.5	93.8	13.6	75.3	15.9	73.8
DCFace-10K-Small	20	88.4	83.7	98.0	79.6	90.8	0.1	0.8	0.1	0.6
DCFace-10K [5]	54	90.3	88.3	98.3	83.5	91.7	77.0	86.3	<b>81.0</b>	<b>89.2</b>
Digi2Real-10K [3]	20	82.6	86.3	98.5	81.5	87.4	67.5	80.0	73.0	83.5
HyperFace-10K [4]	64	87.1	89.1	98.7	84.7	89.8	67.3	84.7	75.7	87.3
SD35-All-IPA	33	79.4	80.3	94.8	75.7	83.2	54.8	68.9	58.6	72.2
SD35-Arc2Face-All-IPA	41	79.4	81.0	95.2	75.7	83.4	55.7	69.3	59.0	72.6
SD35-Arc2Face-8	8	77.0	75.3	94.1	72.6	81.5	56.7	68.3	60.6	71.8
SD35-Arc2Face-16	16	82.2	83.2	96.6	77.4	85.4	64.9	75.8	<b>68.9</b>	<b>78.8</b>
SD35-Arc2Face-24	24	83.1	83.7	97.0	78.9	86.2	19.6	68.7	27.0	72.5
SD35-Arc2Face-32	32	83.9	85.3	97.2	78.9	85.6	6.6	51.4	9.0	57.8
Flux-All-IPA	33	78.4	79.9	95.1	74.4	83.6	55.4	66.5	59.0	69.9
Flux-Arc2Face-All-IPA	41	80.8	77.1	96.0	73.9	85.6	61.3	73.0	64.2	76.1
Flux-Arc2Face-8	8	73.0	73.1	93.5	71.7	81.3	53.8	66.4	57.8	70.0
Flux-Arc2Face-16	16	80.1	78.9	96.2	75.4	85.3	65.0	75.7	69.2	78.7
Flux-Arc2Face-24	24	80.6	80.4	97.0	76.7	86.5	65.0	76.0	<b>68.8</b>	<b>79.2</b>
Flux-Arc2Face-32	32	80.9	82.1	96.9	77.2	86.5	18.4	58.3	19.9	61.3

Table 2. Bias assessment on RFW of ResNet-50 AdaFace model trained on different real (highlighted green) and synthetic datasets. Accuracy values are reported as %.

Training dataset	Group-Level Accuracy $\uparrow$				Avg Acc $\uparrow$	Std $\downarrow$	TO $\uparrow$
	African	Asian	Caucasian	Indian			
WebFace42M-10K [2]	81.80	84.20	90.12	85.12	85.31	3.03	82.28
CASIA WebFace [28]	87.40	85.07	92.13	87.60	88.05	2.56	<b>85.49</b>
DCFace-10K-Small [5]	69.10	77.28	83.33	79.05	77.19	5.16	72.03
DCFace-10K [5]	73.33	79.38	86.22	82.00	80.23	4.67	75.56
Digi2Real [3]	77.52	76.82	82.97	78.63	78.98	2.39	76.59
HyperFace [4]	74.88	81.50	86.27	83.32	81.49	4.18	77.31
SD35-All-IPA	74.60	74.58	77.18	75.43	75.45	<b>1.06</b>	74.39
SD35-Arc2Face-All-IPA	73.55	74.85	76.82	74.72	74.98	1.17	73.81
SD35-Arc2Face-16	76.45	76.55	79.43	76.73	77.29	1.24	76.05
SD35-Arc2Face-24	78.17	78.45	81.42	77.73	78.94	1.45	<b>77.49</b>
Flux-All-IPA	68.25	71.45	74.12	69.25	70.77	2.25	68.51
Flux-Arc2Face-All-IPA	73.15	75.50	77.77	71.28	74.42	2.44	71.98
Flux-Arc2Face-16	74.78	76.32	77.90	73.90	75.72	1.52	74.20
Flux-Arc2Face-24	76.27	76.78	79.20	75.12	76.84	1.49	75.35

epochs 12, 20, and 26. We also applied with 0.2 probability the random data augmentations provided in the AdaFace code<sup>2</sup> by default: central cropping, resizing, horizontal flip, and photometric augmentation. At inference, the resulted model generates a 512- $d$  feature representation (*embedding*) for each input presentation. These embeddings were subsequently compared using cosine similarity to perform identity matching.

## 4.2. Evaluation metrics

We evaluated the overall recognition performance of the FR models using accuracy metrics on the benchmarking datasets (see Section 3.3), as it is typically done in the cur-

rent literature on face recognition [30, 33, 41]. For IJB-B/C datasets [25, 26], we report the True Accept Rate (TAR) at a False Accept Rate (FAR) of 1E-3 and 1E-4. We use RFW dataset [27] for the analysis of bias across racial demographics, measuring bias as the standard deviation of accuracy across its demographic subgroups. For all experiments, we first determined the optimal decision threshold—yielding the highest possible accuracy – on the complete RFW dataset, and then applied this same threshold across each racial subgroup to compute group-specific accuracy values.

Given the necessity to assess both accuracy and fairness in face recognition models, we also used the trade-off (TO)




Figure 5. ROC computed on IJB-B and IJB-C databases for AdaFace models trained on faces generated with 16, 24, or 32 images per ID by Arc2Face augmentation generator from Flux (blue color) or SD35 (golden color) seed identities.

metric, which has been used in recent studies on synthetic data in FR [32, 33]. Simply put, TO metric is the difference between an average accuracy and the standard deviation of accuracy across demographic groups.

## 5. Experimental results

In this section, ResNet-50 AdaFace model<sup>2</sup>, trained on different real and synthetic datasets, is evaluated on standard benchmark datasets for face recognition and on a balanced RFW dataset for bias assessment.

### 5.1. Benchmarking on standard datasets

Table 1 shows the performance of ResNet-50 AdaFace model<sup>2</sup> trained on different synthetic and real datasets and evaluated on standard benchmarking and IJB-B/C datasets (see Section 3.3 for details). All datasets in the table have about 10-11K identities but the number of images per ID can vary significantly and is indicated in the ‘#Images per ID’ column. In the table, real datasets are colored green, baseline synthetic dataset are colored red, the data generated from SD35 is orange, and from Flux is blue. For each category/color, the best performing on IJB-B/C dataset is set in **bold** and the worst is in *italic*.

Table 1 highlights several important issues:

- High accuracy on the five standard benchmark datasets (AgeDB30, CFP-FP, etc.) does not always translate to the high true accept rate (TAR) on IJB-B/C datasets. A notable example is CASIA WebFace showing the highest accuracy on the five benchmark datasets with an average at 94.8% but very low performance on both IJB-B/C datasets. The overall size of CASIA WebFace is third largest with 46 images per ID (It is an average number for real datasets), only full DCFace 0.5M

(marked as ‘DCFace-10K’ in the table) and HyperFace have more total images, however, it did not help it to generalize to IJB-B/C datasets.

- A 10K subset of real faces from WebFace42M leads to the best overall performance. The closest synthetic dataset in performance is DCFace-10K with an average accuracy of 90.4% on five standard benchmarks vs 93.4% for WebFace42M and with about 10% smaller in TAR on IJB-B/C.
- In synthetic datasets, the number of images per ID has a significant impact on the performance. If one looks at five benchmark datasets only, for the same generator, the more images per ID lead to better accuracy. However, when looking at IJB-B/C results, the performance can be very different. Notably, DCFace-10K-Small, which is obtained from the actual DCFace 0.5 (marked as ‘DCFace-10K’ in the table) by randomly sub-sampling the number of images per ID from 54 to 20, has almost zero true positive rates on IJB-B/C. The same identities, just about half of images per ID kept, and the performance drops from the highest for synthetic datasets (DCFace-10K) to zero (DCFace-10K-Small) on IJB-B/C. This surprising result illustrates that a dedicated evaluation of the DCFace would be required to better understand this behavior, but such investigation is out of the scope of this paper.
- The performance for Flux-based generated images, for the same parameters, is slightly higher on IJB-B/C than for SD35-based images. In terms of the augmentation generator, when mixing all augmentation generators, Arc2Face and all four IP-Adapters (denoted by ‘All-IPA’) and original seed images, it leads to the drop in performance on IJB-B/C. This is probably due



to too much intra-class variation inside each ID (see examples in Figure 3), with face augmentations of one ID possibly leaking into other IDs.

- We evaluate the direct impact of the number of augmentations per ID by fixing all parameters (for generation and training) except for the number of images generated for each ID by Arc2Face augmentation generator. From the results for both Flux and SD35, it is clear that when we increase the number of images per ID from 8 to 16, the performance show a clear relative improvement on all benchmarks. When we increase the the number of images per ID to 24, and then to 32, the accuracy on five benchmarks improves on average. However, it is different for IJB-B/C evaluations. For SD35-based data, the performance on IJB-B/C drops, once the number of images per ID become 24 and drops even more when the number is 32. For Flux-based images, for both 16 and 24 images per ID, the performance on IJB-B/C is similar and drops only when the number of images is 32. To demonstrate this phenomena in more details, we have plotted the receiver operating characteristic (ROC) curves for 16, 24, and 32 images per ID for SD35 and Flux-based data in Figures 5. The figures show ROCs moving to the right with each increase of number of images from 16, 24, and 32, for either SD35 or Flux based data.

## 5.2. Bias assessment

Table 2 shows the bias assessment results (on RFW [27] race-balanced dataset) of ResNet-50 AdaFace model<sup>2</sup> trained on different synthetic and real data. The main observation from ‘Std’ column, which essentially shows the bias of the model (the lower value, the lesser the bias) to the faces of different races, is that SD35-based synthetic data have the lowest overall standard deviation with the lowest value of 1.06 for SD35-All-IPA where all IP-Adapters are equally mixed together as augmentations of each identity. Even real datasets have higher standard deviation values. Clearly, the average accuracy of SD35-based synthetic data is lower than for real datasets or the synthetic HyperFace dataset (though its Std measuring bias is very high). High accuracy of HyperFace on RFW and IJB-B/C can be explained by its higher inter-class variability, which closely matches that of real data of WebFace42M, as illustrated by Figure 4b. Hence, if we make inter-class variability of SD35 higher, we may improve its overall face recognition performance, while still mitigating the bias better than others thanks to its race-balanced nature.

Overall, synthetic datasets generated without balancing the data exhibit high bias, with an exception of Dig2Real, which is on par with some of the Flux-based data. For real datasets, CASIA WebFace outperforms the 10K subset from WebFace42M, however, we should not forget that

CASIA WebFace has a very poor performance on a more challenging IJB-B/C benchmark (see Table 1).

If we consider a one-stop ‘TO’ metric [32,33], computed as a difference between average accuracy and standard deviation, SD35-Arc2Face-24 has the highest TO value among the synthetic datasets. However, it is important to remember that this dataset led to very poor performance on IJB-B/C in Table 1. It means that in practical applications, one needs to consider and conduct the evaluations of both bias and challenging face recognition benchmarks to find the tradeoff between an overall accuracy and bias mitigation properties of the synthetic data used for training.

From Table 2, we can note that SD35-based data is more suitable for bias mitigation. A possible reason is that the identities distribution of SD35 is more similar to the real faces compared to Flux as shown in Figure 4b. Another reason is the high variations in visual appearance of images generated by SD35 compared to Flux. Visually, the Flux generates more professionally looking portraits that are *visually appealing* but probably not suitable for training a fairer face recognition. SD35 images look more like images collected in the wild. Overall, the SD35-based images are suited more for bias mitigation, especially, if they can be combined with real images, thus, improving the overall accuracy while keeping the bias low.

## 6. Conclusion

This paper presents a comprehensive study of how balanced synthetic datasets affect both the accuracy and bias of face recognition systems. Through controlled experiments using text-to-image models and identity-consistent augmentation methods, we demonstrate that demographic balancing during synthetic data generation can significantly reduce racial bias in FR models, as evidenced by lower standard deviation in group accuracies on RFW. However, synthetic data alone still underperforms compared to real datasets on challenging benchmarks like IJB-B/C, particularly when seed identities are less separated from each other. In bias evaluations, synthetic data generated using SD35 consistently achieved better fairness metrics than other methods, including the real datasets, suggesting that using synthetic balanced datasets could help for training fairer practical face recognition systems. Our findings highlight the importance of thoughtful design choices in seed and augmentation generators, and indicate that hybrid training approaches, combining synthetic and real data, may be a promising path forward to achieve both high performance and fairness in face recognition systems.

## Acknowledgement

This work was funded by InnoSuisse 106.729 IP-ICT.



## References

- [1] P. Melzi, C. Rathgeb, R. Tolosana, R. Vera-Rodriguez, D. Lawatsch, F. Domin, and M. Schaubert, “GANDiffFace: Controllable generation of synthetic datasets for face recognition with gans and diffusion models,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2023. 1
- [2] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. 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*, pp. 10492–10502, 2021. 1, 2, 3, 4, 6
- [3] A. George and S. Marcel, “Digi2Real: Bridging the realism gap in synthetic-data face recognition via foundation models,” in *Proc. WACV Workshops*, 2025. 1, 2, 3, 4, 6
- [4] H. O. Shahreza and S. Marcel, “HyperFace: Generating synthetic face-recognition datasets by exploring the face-embedding hypersphere,” in *Accepted in International Conference on Learning Representations (ICLR)*, 2025. 1, 2, 3, 4, 6
- [5] M. Kim, F. Liu, A. K. Jain, and X. Liu, “DCFace: Synthetic face generation with dual condition diffusion model,” in *Proc. CVPR*, 2023. 1, 3, 4, 6
- [6] F. Boutros, J. H. Grebe, A. Kuijper, and N. Damer, “Idiff-face: Synthetic-based face recognition through fizzy identity-conditioned diffusion models,” in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 19593–19604, 2023. 1
- [7] D. Geissbühler, H. O. Shahreza, and S. Marcel, “Synthetic face datasets generation via latent space exploration from brownian identity diffusion,” in *Accepted in International Conference on Machine Learning (ICML)*, 2025. 1
- [8] M. Yeung, T. Teramoto, S. Wu, T. Fujiwara, K. Suzuki, and T. Kojima, “Variface: Fair and diverse synthetic dataset generation for face recognition,” 2025. 1, 3
- [9] A. Atzori, P. Cosseddu, G. Fenu, and M. Marras, “The impact of balancing real and synthetic data on accuracy and fairness in face recognition,” *arXiv preprint arXiv:2409.02867*, 2024. 2
- [10] P. Korshunov, H. Chen, P. N. Garner, and S. Marcel, “Vulnerability of automatic identity recognition to audio-visual deepfakes,” in *2023 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–10, 2023. 2
- [11] M. Huber, A. T. Luu, F. Boutros, A. Kuijper, and N. Damer, “Bias and diversity in synthetic-based face recognition,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024. 2, 3
- [12] M. Qraitem, K. Saenko, and B. A. Plummer, “From fake to real: Pretraining on balanced synthetic images to prevent spurious correlations in image recognition,” in *Proc. ECCV*, 2024. Extended version: *arXiv:2308.04553*. 2, 3
- [13] K. Kotwal and S. Marcel, “Review of demographic fairness in face recognition,” *arXiv preprint arXiv:2502.02309*, 2025. 2, 3
- [14] C. Yang, C. Liu, X. Deng, D. Kim, X. Mei, X. Shen, and L.-C. Chen, “1.58-bit flux,” *arXiv preprint arXiv:2412.18653*, 2024. 2, 3
- [15] S. AI, “Stable diffusion 3: Scaling rectified flow transformers for high-resolution image synthesis,” *arXiv preprint arXiv:2403.03206*, 2024. 2, 3
- [16] F. P. Papantoniou, A. Lattas, S. Moschoglou, J. Deng, B. Kainz, and S. Zafeiriou, “Arc2Face: A foundation model for id-consistent human faces,” *arXiv preprint arXiv:2403.11641*, 2024. 2, 3, 4
- [17] L. Zhang, H. Qiu, and J. Liao, “IP-Adapter: Text-compatible image prompt adapter for diffusion models,” *arXiv preprint arXiv:2308.06721*, 2023. 2, 3, 4
- [18] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, “Retinaface: Single-shot multi-level face localisation in the wild,” in *CVPR*, 2020. 2, 4
- [19] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*, pp. 8748–8763, PmLR, 2021. 2, 4
- [20] C. Schuhmann, R. Beaumont, R. Vencu, C. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman, *et al.*, “Laion-5b: An open large-scale dataset for training next generation image-text models,” *Advances in neural information processing systems*, vol. 35, pp. 25278–25294, 2022. 2
- [21] M. Byeon, B. Park, H. Kim, S. Lee, W. Baek, and S. Kim, “Coyo-700m: Image-text pair dataset.” <https://github.com/kakaobrain/coyo-dataset>, 2022. 2
- [22] A. George, C. Ecabert, H. O. Shahreza, K. Kotwal, and S. Marcel, “Edgeface: Efficient face recognition model for edge devices,” *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 6, no. 2, pp. 158–168, 2024. 2, 5
- [23] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” in *Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition*, 2008. 2, 3, 4
- [24] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou, “Agedb: The first manually collected, in-the-wild age database,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 51–59, 2017. 2, 3, 4
- [25] C. Whitelam, E. Taborsky, A. Beveridge, *et al.*, “IARPA Janus Benchmark — B (IJB-B): Face recognition benchmarking,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, p. 90–98, 2017. 3, 4, 6
- [26] B. Maze, J. Adams, J. Duncan, *et al.*, “IARPA Janus Benchmark — C: Face recognition in video,” in *IEEE International Conference on Biometrics (ICB)*, pp. 158–165, 2018. 3, 4, 6

- [27] M. Wang, W. Deng, J. Hu, X. Tao, and Y. Huang, "Racial faces in the wild: Reducing racial bias by information-maximization adaptation network," in *Proc. ICCV*, 2019. 3, 5, 6, 8
- [28] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning face representation from scratch," *arXiv preprint arXiv:1411.7923*, 2014. 3, 4, 6
- [29] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of style-GAN," in *Proc. CVPR*, pp. 8107–8116, 2020. 3
- [30] K. Kotwal and S. Marcel, "Demographic fairness transformer for bias mitigation in face recognition," in *2024 IEEE International Joint Conference on Biometrics (IJCB)*, pp. 1–10, IEEE, 2024. 3, 6
- [31] P. Melzi, C. Rathgeb, R. Tolosana, R. Vera-Rodríguez, A. Morales, D. Lawatsch, F. Domin, and M. Schaubert, "Synthetic data for the mitigation of demographic biases in face recognition," in *Proc. IJCB (Special Session: Synthetic Data in Biometrics)*, 2024. arXiv:2402.01472. 3
- [32] I. DeAndres-Tame, R. Tolosana, P. Melzi, R. Vera-Rodríguez, M. Kim, C. Rathgeb, X. Liu, A. Morales, J. Fierrez, J. Ortega-Garcia, *et al.*, "Frcsyn challenge at cvpr 2024: Face recognition challenge in the era of synthetic data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3173–3183, 2024. 3, 7, 8
- [33] P. Melzi, R. Tolosana, R. Vera-Rodríguez, M. Kim, C. Rathgeb, X. Liu, I. DeAndres-Tame, A. Morales, J. Fierrez, J. Ortega-Garcia, *et al.*, "FRCSyn challenge at WACV 2024: Face recognition challenge in the era of synthetic data," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 892–901, 2024. 3, 6, 7, 8
- [34] P. Melzi, R. Tolosana, R. Vera-Rodríguez, M. Kim, C. Rathgeb, X. Liu, I. DeAndres-Tame, A. Morales, J. Fierrez, J. Ortega-Garcia, *et al.*, "FRCSyn-onGoing: Benchmarking and comprehensive evaluation of real and synthetic data to improve face recognition systems," *Information Fusion*, vol. 107, p. 102322, 2024. 3
- [35] M. Kim, A. K. Jain, and X. Liu, "AdaFace: Quality adaptive margin for face recognition," in *Proc. CVPR*, 2022. 3, 5
- [36] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4690–4699, 2019. 4
- [37] T. Zheng, W. Deng, and J. Hu, "Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments," *arXiv preprint arXiv:1708.08197*, 2017. 4
- [38] T. Zheng and W. Deng, "Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments," Tech. Rep. 18-01, Beijing University of Posts and Telecommunications, 2018. 4
- [39] S. Sengupta, J.-C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs, "Frontal to profile face verification in the wild," in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1–9, IEEE, 2016. 4
- [40] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE signal processing letters*, vol. 23, no. 10, pp. 1499–1503, 2016. 5
- [41] K. Kotwal and S. Marcel, "Mitigating demographic bias in face recognition via regularized score calibration," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1150–1159, 2024. 6