

Studying Bias in GANs through the Lens of Race

Vongani H. Maluleke*, Neerja Thakkar*, Tim Brooks, Ethan Weber, Trevor Darrell, Alexei A. Efros, Angjoo Kanazawa, and Devin Guillory

UC Berkeley

Abstract. In this work, we study how the performance and evaluation of generative image models are impacted by the racial composition of their training datasets. By examining and controlling the racial distributions in various training datasets, we are able to observe the impacts of different training distributions on generated image quality and the racial distributions of the generated images. Our results show that the racial compositions of generated images successfully preserve that of the training data. However, we observe that truncation, a technique used to generate higher quality images during inference, exacerbates racial imbalances in the data. Lastly, when examining the relationship between image quality and race, we find that the highest perceived visual quality images of a given race come from a distribution where that race is well-represented, and that annotators consistently prefer generated images of white people over those of Black people.

Keywords: GANs, Racial Bias, Truncation, Data Imbalance

1 Introduction

The computer vision community has wrestled with problems of bias for decades [45,50]. As vision algorithms are starting to become practically useful in the real world, this issue of bias has manifested as a serious problem in society [11,43,6,47]. In particular, GANs [17] have significantly increased in quality and popularity over the past few years [8,31], and these models have been shown to contain racial biases [28,14,27]. As GANs are increasingly used for synthetic data generation and creative applications, there is the potential for racial bias to propagate to downstream applications, and the need for an understanding of the cause of biased outputs. In generative image models, the question of whether the source of biased outputs comes from the data with which models are trained (data distribution bias) or the algorithms themselves (algorithmic bias) is unanswered.

In this work, we aim to understand the source of bias in GANs in the context of perceived race[†], i.e., can dataset imbalance alone sufficiently describe

*Equal contribution in alphabetical order.

[†]We do not not objectively evaluate the underlying actual race, but rather measure the perceived race of the image. This is because race is a complex social construct and it is not sufficient to evaluate race with only visual features. See section 3.1

issues of racial representation in generative image models? Or do algorithmic choices also contribute [25]? We consider the following types of bias in generative image models as they pertain to class distributions and image quality: 1) *Data distribution bias*: imbalances in training data that are replicated in the generated data, 2) *Symmetric algorithmic bias*: imbalances in training data that are exacerbated in the generated data, irrespective of which race labels are over or under-represented in the data, and 3) *Asymmetric algorithmic bias*: unequal effects on different classes, dependent on or independent of class representation in the training data.

We conduct a systematic study, exploring the following research questions:

1. Will a racially imbalanced training dataset lead to an even more imbalanced generated dataset?
2. Will improving sample quality using the commonly employed “truncation trick” exacerbate an underlying racial imbalance?
3. If a generator is trained on an imbalanced dataset, will perceived visual quality of the generated images change depending on class representation?

We explore these research questions in the context of StyleGAN2-ADA [33] trained to generate human faces. To measure the impact of dataset imbalance, we first label a subset of FFHQ [31], to understand the racial composition of this popular and representative dataset for training generative models of face images. We also train StyleGAN2-ADA on three datasets with varying controlled ratios of images of persons perceived as Black or white. We then measure the perceived racial distribution of training and generated data, with and without truncation, and study the relationship between quality and race class label distribution. To obtain the perceived racial distribution, we use Amazon Mechanical Turk Annotation (AMT) annotations, as well as a classifier that is calibrated against human performance. The AMT annotations are also used to measure the perceived visual quality of real and generated images.

Our findings show that 1) GANs appear to preserve the racial composition of training data, even for imbalanced datasets, exhibiting data distribution bias 2) however, truncation exacerbates discrepancies in the racial distribution equally amongst race class labels, exhibiting symmetric algorithmic bias, and 3) when ranking images by quality, we find that generated images of a given perceived race are of higher perceived quality when they come from a generator that is over-represented for images labeled as white, while images labeled as Black retain constant quality regardless of the training data’s racial composition. We also find that both real and generated white labeled facial images are consistently annotated as higher quality than real and generated images of Black people. It is unclear whether this observed asymmetric algorithmic bias is caused by StyleGAN2-ADA, our human-centric system of evaluation, underlying qualitative discrepancies in training data, or a combination thereof.

2 Related Work

Racial Bias in Computer Vision Machine learning models and their applications have a well-documented history of racial bias, spanning vision [11,35], language [10,9,46], and predictive algorithms that have a heavy impact on real peoples’ lives [43,42,29]. Numerous research efforts have aimed to evaluate, understand, and mitigate bias, particularly in computer vision. Buolamwini *et al.* [11] analyzed three automated facial analysis algorithms, and found that all classifiers performed worse on images of individuals with a darker skin type compared to counterparts with a lighter skin type [11]. A similar conclusion was made in earlier research by Klare *et al.* [35], who found that face recognition algorithms consistently performed poorly on young Black females. Phillip *et al.* [44] showed that machine learning algorithms suffer from the “other race effect” (humans recognize faces of people from their race more accurately compared to faces of other races) [44].

Racial Bias in Generative Models Image generation models have been shown to contain racial biases [28,14,27]. AI Gahaku [14], an AI art generator that turns user-submitted photos into Renaissance-style paintings, often turns photos of people of color into paintings that depict white people. The Face Depixelizer, a tool based on PULSEGAN [40], which super-resolves a low-resolution face image, also tends to generate an image of a white person, regardless of input race. Jain *et al.* [28] demonstrated that popular GAN models exacerbate biases along the axes of gender and skin tone when given a skewed distribution of faces; for example, Snapchat’s beautification face filter lightens skin tones of people of color and morphs their faces to have euro-centric features. GANs can inherit, reinforce and even exacerbate biases when generating synthetic data [27].

Racial Bias Mitigation Several works have proposed ways to mitigate racial bias in facial recognition systems by modifying models directly or with data sampling strategies [19,51,52]. To reduce racial bias via model modification, Wang *et al.* [51] proposed a deep information maximization adaptation network (IMAN), with white faces as the source domain and other races as target domains. Gwilliam *et al.* [19] performed facial recognition experiments by manipulating the race data distribution to understand and mitigate racial bias. Their work demonstrated that skewing the training distribution with a majority of African labeled images tends to mitigate racial bias better than balanced training data set.

Generative Adversarial Networks GANs [16], a class of implicit generative models, learn to generate data samples by optimizing a minimax objective between discriminator and generator networks. The discriminator is tasked with differentiating training images from generated images, and the generator aims to fool the discriminator. Modern GANs [8,31,33] are capable of producing high quality images and are increasingly leveraged for image manipulation tasks [26,1].

GAN Truncation The “truncation trick” introduced by Brock *et al.* [8] is a sampling technique that allows deliberate control of the trade-off between variety and fidelity in GAN models. At the loss of some diversity, the fidelity of generated images can be improved by sampling from a shrunk or truncated distribution [2,39,8]. StyleGAN implements truncation by interpolating towards the mean intermediate latent vector in W space [31]. In this work we evaluate the impact of truncation on racial diversity in images generated with StyleGAN2-ADA [30].

GAN Mode Collapse GANs are known to exhibit mode collapse or mode dropping, where certain features present in the training dataset distribution are missing from the distribution of generated images [15]. Many works propose solutions to address mode collapse, such as Wasserstein GANs [3,18], Prescribed GANs [13], and Mode Seeking GANs [38]. In spite of these works, mode dropping is not fully understood. Arora *et al.* [4] show that the generated distribution of a GAN has a relatively low support size (diversity of images) compared to the training distribution. In the work “Seeing What a GAN Cannot Generate,” Bau *et al.* [5] visualized regions of images that GANs are unable to reproduce. They found that higher-quality GANs better match the dataset in terms of the distribution of pixel area belonging to each segmentation class, and that certain classes, such as people, were particularly challenging for the GANs. These works indicate that GANs may exacerbate bias in the training data distribution by dropping certain features or classes in generated images — in this work, we also analyze whether these effects occur regarding racial bias.

3 Methodology

3.1 Racial Categorizations

Race is a dynamic and complex social construct. People can be of multiple races and perception of race is heavily dependent on culture and geography with different meanings and interpretations. As such, all discussion in this work pertains to perceived race by annotators. Despite the complexities and subjectivity involved in analyzing racial perceptions, we choose to study the bias in GANs through the lens of race as this is a topic of societal consequence [28,14,27].

The decision to use perceived racial classifications over skin color (*tone/shade*) estimates and categorizations (such as Individual Typology Angle (ITA) and Fitzpatrick skin phototype [53]) was driven by the notion that perceived racial categorization is an informative lens through which to study racial bias [20], and the availability of the FairFace dataset, a large-scale face dataset with seven race categorizations. Furthermore, Karkkainen *et al.* [36] found that solely relying on skin color/ITA in the FairFace dataset is not sufficient to distinguish race categories. We condense the seven race categories into three: Black, white, and Non-Black or Non-white, to further reduce perceptual ambiguity (see sec. 4.1). This study does not aim to minimize the importance of understanding these questions in Non-Black or Non-white races, but as a first study we simplify the categorization.




Fig. 1: **Training Data and Generated Data.** The top row shows real data from the FairFace dataset used for training generative models, sampled in ratios used for training. The bottom row shows generated data from models trained on the 80B-20W, 50B-50W, and 20B-80W datasets (left to right).

3.2 Datasets

This section describes the FFHQ and FairFace datasets. It also explains how we use the datasets: we quantify the racial distribution of FFHQ, train generative models on both datasets to answer our three research questions, and also use FairFace to train a classifier on perceived race.

Flickr Faces HQ (FFHQ) is one of the most commonly-used datasets for the development of high fidelity GANs that generate faces, such as StyleGAN [31] and StyleGANV2 [33]. It contains 70,000 high-quality 1024x1024 resolution images scraped from Flickr, which were automatically aligned and cropped [32]. We quantify the racial distribution of face images on a subset of FFHQ, as well as that of a StyleGANV2 model trained on FFHQ with and without truncation.

FairFace consists of 108,501 images annotated by age, gender and race, which are cropped, aligned and 224×224 [36]. The representative faces in this dataset come from public images without public figures, minimizing selection bias. As described in section 3.1, we use the FairFace race categorization choices as a starting point.

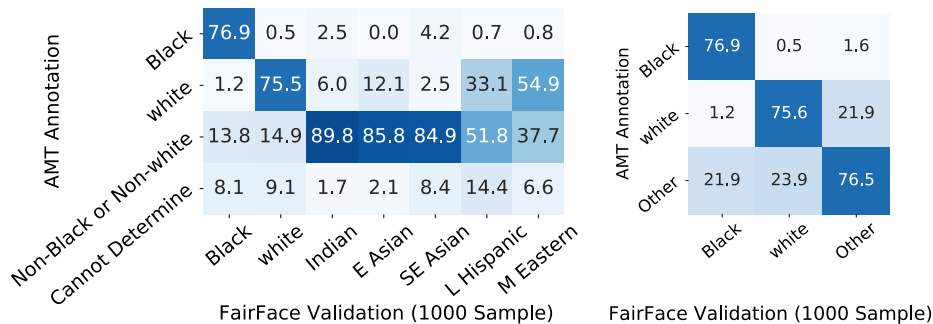


Fig. 2: **Consistency of AMT annotation labels vs. Full FairFace labels.** (Left) A confusion matrix of AMT annotations and the 7 FairFace labels on 1000 random samples of the FairFace validation set. (Right) A condensed version of the confusion matrix. All numbers are shown in percentages. Overall, there is 76% agreement between the FairFace labels and our collected annotations.

We use this dataset to train a perceived race classifier and multiple StyleGAN2-ADA models. From the FairFace dataset, we create three overlapping subsets of 12K images, each with images randomly sampled in the following ratios: 80% Black/20% white (80B-20W), 50% Black/50% white (50B-50W), and 20% Black/80% white (20B-80W). Six different 128×128 resolution StyleGAN2-ADA models [30] were trained using an “other races” dataset, an all Black dataset, an all white dataset, and each of the three FairFace datasets described above (i.e., 80B-20W, 50B-50W, and 20B-80W). All StyleGAN2-ADA models were trained with 2 GPUs, batch size 64, and an R1-gamma parameter value 0.0128, ensuring high fidelity images for 25000 kimg. Example generated images trained with the FairFace datasets can be seen in the bottom row of Fig. 1.

3.3 Amazon Mechanical Turk Annotation

Amazon Mechanical Turk (AMT) workers labeled tasks associated with the following three questions:

1. **Race Classification:** What is the race of the person in the image? [Choices: Black, white, Cannot Determine, Non-Black or Non-white.]
2. **Real/Fake Classification:** Is this image real or fake?
3. **Image Quality Ranking:** Which image is more likely to be a fake image?

In our tasks, 1000 randomly sampled real images from FairFace and FFHQ, respectively, and 1000 images from each dataset of generated images were labeled. More details are in the supplementary material.

4 Experiments and Results

In this section, we first establish the reliability of our AMT annotation process and how we condense and use race labels, and then use the AMT annotations

to compute the racial composition of a subset of FFHQ. We then assess the relationship between the racial distributions of the training and generated data, and evaluate the impact of truncation on the racially imbalance in the data. Finally, we assess the relationship between the training data racial distribution and the perceived image quality of the GAN-generated images.

4.1 The Racial Distribution of Training and GAN-Generated Data

Annotation Consistency Analysis We use annotations to measure the racial distribution of the training and generated images. We first assess the performance and reliability of our procedure by collecting annotations from a random sample of 1000 images from the FairFace validation set. The confusion matrix in Fig. 2 shows the difference in labels on the FairFace validation set and the annotations collected using our AMT protocol, demonstrating the inherent limitations in attempting to establish racial categorizations on visual perception alone. However, we find that the ambiguity in visual discernment is lowest between images perceived as Black and white, making these two classes suitable for analysis of racial bias. Limiting the observed racial categories leads to more consistent labeling, allowing for a thorough examination of the impacts between two race class labels from the data.

The Racial Distribution of FFHQ We analyze the racial distribution of FFHQ by selecting a random subset of 1000 images and collecting images on AMT using our procedure for task one described in Section 3.3. We find that FFHQ is composed of 69% white, 4% Black, and 27% non-Black or non-white facial images. Compared to the global Black population, FFHQ is under-representative.

Relationship between Training and GAN-Generated Data Distributions On our first research question, regarding if an imbalanced dataset further exacerbates the generated dataset distribution, our experiments indicate that StyleGAN2-ADA’s generated data distribution preserves the training data distribution. We compute the perceived racial distributions of FFHQ and FairFace (20B-80W, 50B-50W, and 80B-20W) training data and generated data, based on our AMT annotations. To explicitly showcase the ratio of Black and white race class labels in the training and generated data, we excluded the “Non-Black or Non-white” and “Cannot Determine” class labels. It can be seen in Fig. 3 that the training and generated data distributions are statistically close - the red bars represent the 95% Wald’s confidence interval (CI) of each generated data distribution. The training data distributions all fall within the 95% confidence interval of the observed sample means, and as such we conclude that the generators successfully preserve the training distributions and exhibits data distribution bias. See the supplementary material for more information on CI calculations.

Impact of Truncation on FFHQ Generated Data Distribution This section studies our second research question on the effect of truncation on racial

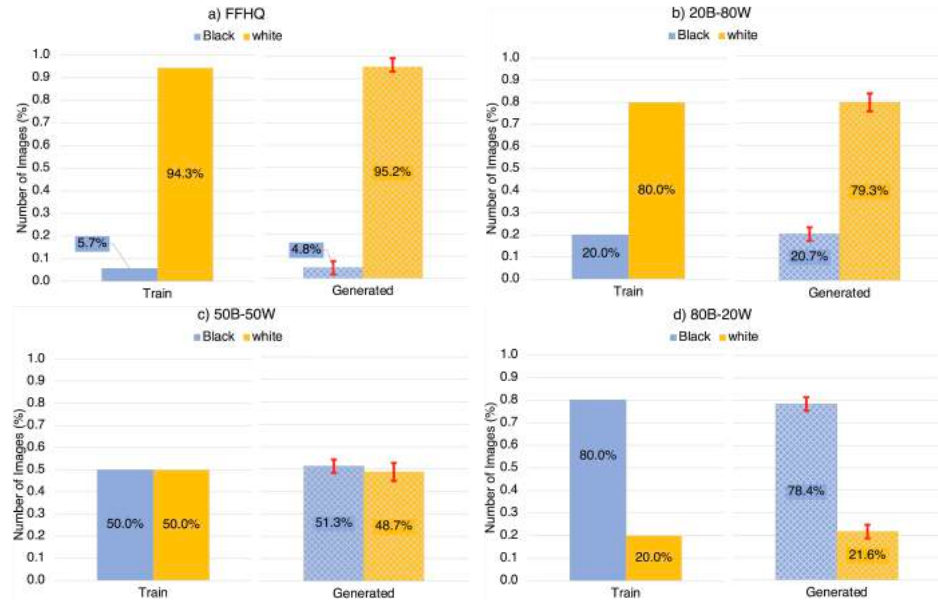


Fig. 3: **Racial distribution of training and GAN-generated data.** Distributions for (a) FFHQ, (b) 20B-80W, (c) 50B-50W and (d) 80B-20W. The red bars represent the 95% confidence interval for the expected distribution of the generated data. All class labels aside from Black and white are excluded. All of the generative models preserve the distribution of the training data.

imbalance. We follow the same protocol as above on truncation levels 1 (no truncation), 0.75, and 0.5, and find that that applying truncation when generating data exacerbates the racial imbalance in StyleGAN2-ADA. Fig. 4(a) shows the AMT annotation distribution of the FFHQ training data, and 4(b) shows the distribution of StyleGAN-2 ADA trained on FFHQ without truncation. As greater truncation levels are applied, the generated data becomes increasingly racially imbalanced. The percentage of images of Black people in the generated data distribution in Fig. 4 drops from 4% to 0% at a truncation level of 0.5. We observe an inverse effect for images of white people, where more truncation increases the percentage of the white class labeled images in the generated data distribution.

Automatic Race Classifier In order to conduct a more fine-grained study on the effect of truncation level, we scale the AMT annotation process by using an automatic race classifier to classify perceived race. A ResNet-18 model [21] was used to carry out three-way classification on face images. The model was trained on the FairFace training split augmented with equal quantities of generated images from StyleGAN2-ADA models trained on all-Black, all-white,




Fig. 4: **The racial distribution of FFHQ with truncation.** The top row shows the FFHQ training data distribution (left) and the generated data distribution without truncation (right). The bottom row shows the generated data distribution with a truncation level of 0.75 (left) and an increased truncation level of 0.5 (right). While the model without truncation closely preserves the original training data distribution, as the level of truncation increases, the ratio of white to Non-white class labels increases.

and all-“other”-races datasets. Confusion matrices showing the performance of the classifier on the FairFace validation set and our collected annotations are shown in Fig. 5. While the automatic classifier performance is not perfect, at 84% accuracy, it suffices as a reasonable proxy for AMT annotations, given that the confusion between the classifier labels and our collected annotations on images labeled as Black and white is similar to the confusion between our collected annotations and the FairFace labels seen in Fig. 2.

Evaluation of Truncation We evaluate levels of truncation and observe the following trend across all models: as the level of truncation increases, racial diversity decreases, converging to the average face of the dataset. Images were generated from StyleGAN2 trained on FFHQ and the 80B-20W, 50B-50W, and 80B-20W FairFace-trained generators at truncation levels ranging from $\gamma = 0$ to 1 at intervals of 0.1. The perceived race labels were automatically classified for 10K generated images at each truncation level, for a total of 110K images, with results in Fig. 6. We observe that truncation in a dataset with predominantly images of white people, such as FFHQ, increases the frequency of generating images classified as white. Similarly, when the majority of images in a dataset




Fig. 5: **Automatic Classifier Performance.** Confusion matrices for the automatic classifier on the entire FairFace validation set (left) and the automatic classifier on 1000 images with collected annotations (right). On our AMT annotations, the classifier confuses images labeled as Black and white at a level comparable to that of our annotations and the FairFace labels (Fig. 2).

are of Black people, as in the 80B-20W dataset, the truncated generated data distribution has predominantly images classified as Black. Examples of generated face images at different levels of truncation can be seen in the supplementary material.

4.2 GAN Quality

From the AMT Image Quality Ranking task, we perform pairwise comparisons of the generators’ respective data distributions against each other. We determine that when trained with FairFace splits, images from the generator trained on more images of white people are always preferred at a greater proportion, and on FFHQ with truncation, images with more truncation are seen as higher quality.

This finding comes from counting the number of times images from one generator are preferred over images from another generator. Results of this comparison can be seen in Table 1 for FairFace and Table 2 for FFHQ. Error bars are computed using the Wald’s method for a 95% confidence interval. For FairFace, we find that the generator trained on a higher percentage of images of white people tends to be preferred in more comparisons by a narrow margin that often surpasses error bounds. For FFHQ with truncation, an increased truncation level always leads to a generator being more preferred, indicating that truncation increases perceptual sample quality.

Correlation with FID. Fréchet inception distance (FID) [22] is a common GAN metric that measures image quality by comparing distributions of features extracted with the Inception network [49] between training and generated images. Higher quality models receive a lower FID score. The FIDs of our FairFace-trained generators are 5.60, 5.63, and 5.68 for the 20B-80W, 50B-50W, and 80B-20W models respectively, not revealing a clear difference in perceived visual quality based on this automatic metric.



Fig. 6: **Automatic evaluation of truncation.** Automatically evaluated results of StyleGAN models trained on various datasets, with 110,000 images total generated at levels of truncation from $\gamma = 1.0$ (no truncation) to $\gamma = 0.0$ (full truncation). The y-axis represents the racial breakdown of the dataset, which becomes more polarized as truncation increases.

4.3 Perceived Visual Image Quality and Race

To address the third research question, i.e., to determine if there is a relationship between perceived race and generated image quality, we examine the results of our binary real/fake classification task and our pairwise image quality ranking task. Our findings on the real or fake classification task do not yield a clear relationship between the training data distribution and generated image quality; please see the supplementary material for details. However, pairwise image quality comparisons provide a more fine-grained analysis. From a perceptual quality ranking obtained from pairwise comparisons, we find that the average perceived visual quality of generated images of a particular race increases as the proportion of training images of that race increases. We also find that generated images of white people tend to be perceived as higher quality than images of Black people, regardless of the training distribution.

Using 3000 FairFace 80B-20W, 50B-50W, 20B-80W dataset images, 54000 pairwise comparisons were evaluated within and across the datasets. From these pairwise comparisons, we use the `choix` package’s [37] implementation of the Bradley-Terry model [12] to rank the 3000 images in descending order of image quality. From this global ranking, we obtain a ranked ordering of all images labeled as Black and white. Table 3 investigates the breakdown of the top K

Table 1: **Pairwise image quality comparison of FairFace generators.** 9000 comparisons were conducted between each pair of generators, resulting in a total of 27000 comparisons. We report the percentage of images that are preferred from the left generator over the right, with the accompanying 95% Wald’s CI. Generators trained on datasets with a greater number of images of white people tend to be perceived as having better image quality.

Generator A	Generator B	Percentage Gen. A Preferred
20B-80W	80B-20W	53.9 ± 1.02
20B-80W	50B-50W	52.0 ± 1.03
50B-50W	80B-20W	51.0 ± 1.03

Table 2: **Pairwise image quality comparison of FFHQ at different truncation levels.** The percentage of images that are preferred when generated with the left truncation over the right, with a Wald’s 95% CI. A truncation level of $\gamma = 1$ corresponds to no truncation, and $\gamma = 0.5$ corresponds to the most truncation. Images generated with more truncation are perceived as being of higher quality.

Truncation A	Truncation B	Percentage Trunc. A Preferred
0.50	1.00	58.7 ± 1.02
0.75	1.00	55.4 ± 1.02
0.50	0.75	52.5 ± 1.03

images. In order to obtain weighted percentage scores, the raw counts for the top K images (which can be seen in the supplementary material) of a particular race are normalized by the expected frequency of the images from the corresponding race found in each data split. Then, the weighted numbers are divided by the sum of all scores for that race and value of K . The results indicate that the highest quality images of a particular race are more likely to come from a data split where the race class is over-represented or represented in parity. From the global ranking, a precision-recall curve for each race from each data split over the top K images, and the area under the PR curves, are shown in Fig. 7. Images labeled as white are overall ranked as higher quality than images labeled as Black. Furthermore, for white labels, being in the majority (i.e., from the 20B-80W split) yields better quality than in the minority (i.e., from the 80B-20W split).

These results raise the question of whether a predisposition towards white generated faces is a by-product of our learned generative models, or is a result of other parts of our data collection and evaluation process. In order to gain insight on this question, we conducted 1700 pairwise comparisons between real face images from the FairFace data labeled as Black and white, using the same AMT protocol as for generated data. By removing generative models from this evaluation, we can determine whether external factors such as original real

Table 3: **Top K image composition per-race.** Given a ranking of images labeled as Black and white across all data splits, we break down the data split that each image came from. The highest quality images ($K = 10, 25, 50$) are more likely to come from a data split where they are over-represented or represented in parity

white				Black			
K	80B-20W	50B-50W	20B-80W	K	80B-20W	50B-50W	20B-80W
10	0.00	0.29	0.71	10	0.49	0.23	0.28
25	0.00	0.38	0.62	25	0.57	0.19	0.24
50	0.22	0.36	0.42	50	0.45	0.37	0.17
100	0.26	0.41	0.33	100	0.37	0.32	0.31
500	0.33	0.36	0.31	500	0.35	0.33	0.32

image quality or annotator bias may play a role in our observed results. An evaluation procedure invariant to the perceived races of images should produce results where real images perceived as white are preferred over real images perceived as Black 50% of the time. Instead, they were preferred 55.2% of the time with a 95% Wald’s confidence interval of $55.2\% \pm 2.3\%$. This indicates that even though our system of evaluation is based on pairwise comparisons, a standard and well-regarded GAN quality metric [7], it has a detectable bias towards selecting images labeled as white over those labeled as Black.

The source of this propensity towards selecting images perceived as white is unclear. Captured images of Black people in the dataset could be of lower quality than that of images of white people, potentially because of camera or sensor bias. Due to prevalence, collecting high quality images of white-appearing faces might be easier. Another possibility is the “other race effect” [44], where annotators are biased toward their own race, however, the demographics of the annotators in our study are unknown. A future in-depth study of these factors causing asymmetric algorithmic bias should be a subject of future investigations.

5 Discussion

Through a systematic investigation into the role of racial composition in generative models, we find that state-of-the-art GANs such as StyleGAN2-ADA closely mirror the racial composition of their training data, exhibiting data distribution bias. Our study reveals that in FFHQ, the most prominent dataset of generative models of facial images, Black people are substantially underrepresented at 4%, as compared to the global population. Practitioners should be aware of this bias when using this dataset. We recommend that generative modeling practitioners have more awareness of the racial composition of their training dataset, particularly when the downstream application requires a well-balanced model. When the training data has known biases or imbalances, we recommend transparency through mechanisms such as model cards [41].

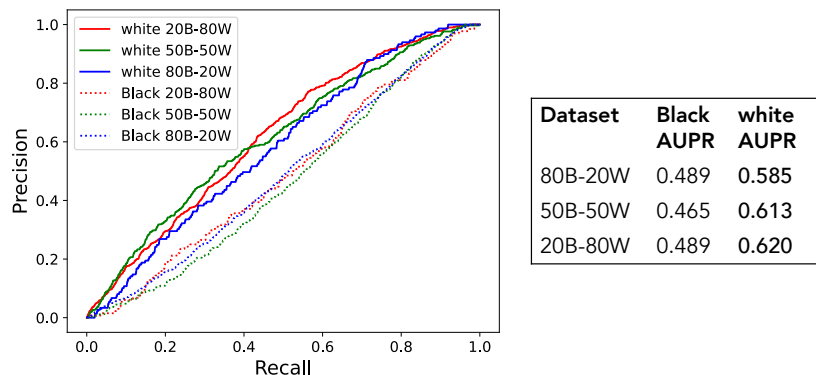


Fig. 7: **PR curves (left)**. PR curves for each race within each generated dataset. For a given K between 0 and 3000, precision, shown on the y-axis, is defined as the count of images of a particular perceived race and dataset in the top K images, normalized by the total number of images of each race and dataset. Recall, shown on the x-axis, is defined as the number of images seen out of the total images (K/N , where $N = 3000$). **Area under the PR curves (right)**. A larger number indicates high image quality; images labeled as white are consistently perceived as higher quality than those labeled as Black, regardless of the generated dataset they come from.

Downstream applications, even generative model demos, often employ the truncation trick to improve the visual fidelity of generated images. Our work shows that this qualitative improvement comes at the expense of exacerbating existing bias. Our studies show that using a well balanced dataset can mitigate this issue of symmetric algorithmic bias. We suggest researchers be transparent on their usage and level of truncation, and encourage research for alternative algorithms to truncation. Interesting future directions include correlating FID to other quality metrics, performing this study on different GAN architectures and other generative models such as VAEs [34,48] and diffusion models [23]. In particular, with diffusion models, it would be interesting to see if classifier-free guidance [24] exhibits the same symmetric algorithmic bias as the truncation trick. Another interesting direction is to perform an intersectional study probing similar questions, by considering other attributes such as gender in addition to race.

Acknowledgements We thank Hany Farid, Judy Hoffman, Aaron Hertzmann, Bryan Russell, and Deborah Raji for useful discussions and feedback. This work was supported by the BAIR/BDD sponsors, ONR MURI N00014-21-1-2801, and NSF Graduate Fellowships. The study of annotator bias was performed under IRB Protocol ID 2022-04-15236.

References

1. Abdal, R., Qin, Y., Wonka, P.: Image2stylegan++: How to edit the embedded images? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8296–8305 (2020)
2. Ackley, D.H., Hinton, G.E., Sejnowski, T.J.: A learning algorithm for boltzmann machines. *Cognitive science* **9**(1), 147–169 (1985)
3. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: Precup, D., Teh, Y.W. (eds.) Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 214–223. PMLR (06–11 Aug 2017), <https://proceedings.mlr.press/v70/arjovsky17a.html>
4. Arora, S., Zhang, Y.: Do gans actually learn the distribution? an empirical study (2017)
5. Bau, D., Zhu, J.Y., Wulff, J., Peebles, W., Strobel, H., Zhou, B., Torralba, A.: Seeing what a gan cannot generate (2019)
6. Benjamin, R.: Race after technology: Abolitionist tools for the new jim code. *Social forces* (2019)
7. Borji, A.: Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding* **179**, 41–65 (2019)
8. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096 (2018)
9. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
10. Brunet, M.E., Alkalay-Houlihan, C., Anderson, A., Zemel, R.: Understanding the origins of bias in word embeddings. In: International conference on machine learning. pp. 803–811. PMLR (2019)
11. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. In: Conference on fairness, accountability and transparency. pp. 77–91. PMLR (2018)
12. Caron, F., Doucet, A.: Efficient bayesian inference for generalized bradley–terry models. *Journal of Computational and Graphical Statistics* **21**(1), 174–196 (2012)
13. Dieng, A.B., Ruiz, F.J.R., Blei, D.M., Titsias, M.K.: Prescribed generative adversarial networks (2019)
14. Gahaku, A.: Ai gahaku (2019)
15. Goodfellow, I.: Nips 2016 tutorial: Generative adversarial networks (2017)
16. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014)
17. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks (2014)
18. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of wasserstein gans (2017)
19. Gwilliam, M., Hegde, S., Tinubu, L., Hanson, A.: Rethinking common assumptions to mitigate racial bias in face recognition datasets. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops. pp. 4123–4132 (October 2021)
20. Hanna, A., Denton, E., Smart, A., Smith-Loud, J.: Towards a critical race methodology in algorithmic fairness. *CoRR* **abs/1912.03593** (2019), <http://arxiv.org/abs/1912.03593>

21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
22. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 6629–6640. NIPS’17 (2017)
23. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
24. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
25. Hooker, S.: Moving beyond “algorithmic bias is a data problem”. *Patterns* **2**(4), 100241 (2021). <https://doi.org/https://doi.org/10.1016/j.patter.2021.100241>, <https://www.sciencedirect.com/science/article/pii/S2666389921000611>
26. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. In: Proc. NeurIPS (2020)
27. Jain, N., Manikonda, L., Hernandez, A.O., Sengupta, S., Kambhampati, S.: Imagining an engineer: On gan-based data augmentation perpetuating biases. arXiv preprint arXiv:1811.03751 (2018)
28. Jain, N., Olmo, A., Sengupta, S., Manikonda, L., Kambhampati, S.: Imperfect imaganation: Implications of gans exacerbating biases on facial data augmentation and snapchat selfie lenses. arXiv preprint arXiv:2001.09528 (2020)
29. Jeff, L., Surya, M., Lauren, K., Julia, A.: How we analyzed the compas recidivism algorithm. Propublica (2016)
30. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. arXiv preprint arXiv:2006.06676 (2020)
31. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
32. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks (2019)
33. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8110–8119 (2020)
34. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
35. Klare, B.F., Burge, M.J., Klontz, J.C., Vorder Bruegge, R.W., Jain, A.K.: Face recognition performance: Role of demographic information. *IEEE Transactions on Information Forensics and Security* **7**(6), 1789–1801 (2012). <https://doi.org/10.1109/TIFS.2012.2214212>
36. Kärkkäinen, K., Joo, J.: Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In: 2021 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1547–1557 (2021). <https://doi.org/10.1109/WACV48630.2021.00159>
37. Lucas, M., Brendan, H.: Github (2022), <https://github.com/lucasmaystre/choix.git>
38. Mao, Q., Lee, H.Y., Tseng, H.Y., Ma, S., Yang, M.H.: Mode seeking generative adversarial networks for diverse image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)

39. Marchesi, M.: Megapixel size image creation using generative adversarial networks (2017)
40. Menon, S., Damian, A., Hu, S., Ravi, N., Rudin, C.: PULSE: self-supervised photo upsampling via latent space exploration of generative models. CoRR **abs/2003.03808** (2020), <https://arxiv.org/abs/2003.03808>
41. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., Gebru, T.: Model cards for model reporting. CoRR **abs/1810.03993** (2018), <http://arxiv.org/abs/1810.03993>
42. Noble, S.U.: Algorithms of oppression: How search engines reinforce racism. New York University Press (2018)
43. O’Neil, C.: Weapons of math destruction how big data increases inequality and threatens democracy. Penguin Books (2018)
44. Phillips, P.J., Jiang, F., Narvekar, A., Ayyad, J., O’Toole, A.J.: An other-race effect for face recognition algorithms. ACM Trans. Appl. Percept. **8**(2) (feb 2011). <https://doi.org/10.1145/1870076.1870082>, <https://doi.org/10.1145/1870076.1870082>
45. Ponce, J., Berg, T.L., Everingham, M., Forsyth, D.A., Hebert, M., Lazebnik, S., Marszalek, M., Schmid, C., Russell, B.C., Torralba, A., Williams, C.K.I., Zhang, J., Zisserman, A.: Dataset issues in object recognition. In: Toward category-level object recognition, pp. 29–48. Springer (2006)
46. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
47. Raji, I.D., Buolamwini, J.: Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society. pp. 429–435 (2019)
48. Razavi, A., Van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. Advances in neural information processing systems **32** (2019)
49. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1–9 (2015)
50. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: CVPR 2011. pp. 1521–1528. IEEE (2011)
51. Wang, M., Deng, W., Hu, J., Tao, X., Huang, Y.: Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 692–702 (2019)
52. Wang, Z., Qinami, K., Karakozis, Y., Genova, K., Nair, P., Hata, K., Russakovsky, O.: Towards fairness in visual recognition: Effective strategies for bias mitigation. CoRR **abs/1911.11834** (2019), <http://arxiv.org/abs/1911.11834>
53. Wilkes, M., Wright, C.Y., du Plessis, J.L., Reeder, A.: Fitzpatrick skin type, individual typology angle, and melanin index in an african population: steps toward universally applicable skin photosensitivity assessments. JAMA dermatology **151**(8), 902–903 (2015)

Supplementary Material: Studying Bias in GANs through the Lens of Race

Vongani H. Maluleke*, Neerja Thakkar*, Tim Brooks, Ethan Weber, Trevor Darrell, Alexei A. Efros, Angjoo Kanazawa, and Devin Guillory

UC Berkeley

In this supplementary material document, we first discuss the selected perceived race label categorization. We then describe in detail the three Amazon Mechanical Turk tasks and implementation information. Followed by an analysis of the performance of both AMT annotations and the automatic classifier in evaluating perceived race. Visualizations of various truncation levels are shown next, and finally, we present more details and visualizations of quality ranking.

1 Race Labels Categorization

We start with the the FairFace dataset labels, and then collect annotations based on our own condensed categorization. The FairFace dataset started with the commonly accepted race categories from the U.S. Census Bureau—white, Black, Asian, Hawaiian and Pacific Islanders (HPI), Native Americans (NA), and Latino. They dropped the HPI and NA categories due to insufficient image examples, and expanded the Asian category into four distinct subgroups: Middle Eastern, East Asian, Southeast Asian, and Indian [?]. To reduce perceptual ambiguity (see main paper in section 4.1), we condense the race class labels from seven FairFace classes to three classes—Black, white, and Non-Black or Non-white—where Non-Black or Non-white comprises the Middle Eastern, East Asian, Southeast Asian, Latino Hispanic, and Indian as labeled by FairFace. We also relabel all the images we analyze using our own annotation protocol with three categories and a “Cannot Determine” category.

2 Amazon Mechanical Turk (AMT) Details

Amazon Mechanical Turk (AMT) was used to collect annotations for three label tasks, namely; (1) race classification, (2) quality classification, and (3) quality ranking. As mentioned in the paper, these tasks consist of the following questions:

1. **Race Classification:** What is the race of the person in the image?
2. **Real/Fake Classification:** Is this image real or fake?
3. **Image Quality Ranking:** Which image is more likely to be a fake image?

Implementation details. Our label tasks were deployed using a custom framework for deploying AMT tasks using our dynamically populated HTML/JavaScript template and the Python API Boto3¹. Our code enables creating human intelligence tasks (HITs) that show images with a corresponding question, and then the annotator completes a forced-choice answer among a set of specified choices. For quality control, we use both **accuracy** and **consistency** checks. As an accuracy test, workers must get eight of these hidden questions correct. As a consistency test, we duplicate these ten test cases and scatter them throughout the HIT, and the worker must be consistent for eight of the repeated examples. Our hidden test cases are chosen to be adequately obvious such that diligent workers will successfully pass them. If less than eight are answered correctly, the worker’s responses are discarded. Furthermore, we ensure that three unique workers answer each question. Each HIT starts with a consent form and a comprehensive description of the task with practice examples with accompanying answers and descriptions; this helps ensure annotators understand the tasks so they can pass our quality control checks.

Next, we go through each of the three tasks and provide qualitative examples for the questions being asked. In total, we asked a total of 50K questions in 1422 unique HITs, where each was labeled by 3 different workers. We did not collect the demographics or other information on AMT workers. Overall, 59 annotators participated in our tasks, and we paid on average \$8.71 per hour. Note that *we will release our code used to conduct our experiments for the benefit of the computer vision community.*

2.1 Race Classification

The Race Classification AMT task asked the workers to identify the race of the person(s) in the image, by selecting one of the options described below:

1. Black - This is an image of a Black person.
2. white - This is an image of a white person.
3. Non-Black or Non-white - I know the race of the person and the person is not Black or white.
4. Cannot Determine - I cannot tell the race of the person.

The AMT workers were given examples of images and corresponding race class labels- 1, as well as a demo of the interface before they could start the task. Fig. 2 is an example of the deployed Race Classification AMT interface.

2.2 Quality Classification

The Quality Classification AMT task asked workers to identify real photographs and fake images by selecting one of the options described below:

1. Real Photograph - This is a photograph of a real person taken using a camera.

¹ <https://boto3.amazonaws.com/v1/documentation/api/latest/index.html>




Fig. 1: Sample of the race classification examples with corresponding race class labels given to AMT workers before they started Task 1.

2. Fake/Manipulated Image - This is a computer-generated image of a person(s) who do not exist.

To further assist the workers in understanding the difference between the two options, the definition of each of the two options was provided to workers as “Real photographs are images of real person(s) captured using a camera” and “Fake/Manipulated images are computer generated images of a person(s) who do not exist”. Fig. 3 is an example of the Quality Classification AMT interface.

Results from Real/Fake Classification Using the Real/Fake indicator as a proxy for image quality, we are unable to determine any significant distinctions in generated image quality with respect to race. We use the label “fake” as a proxy for low quality and “real” to represent high quality. This image quality proxy measured was collected by using AMT annotators to determine if an image



Fig. 2: Race Classification AMT Interface Example.

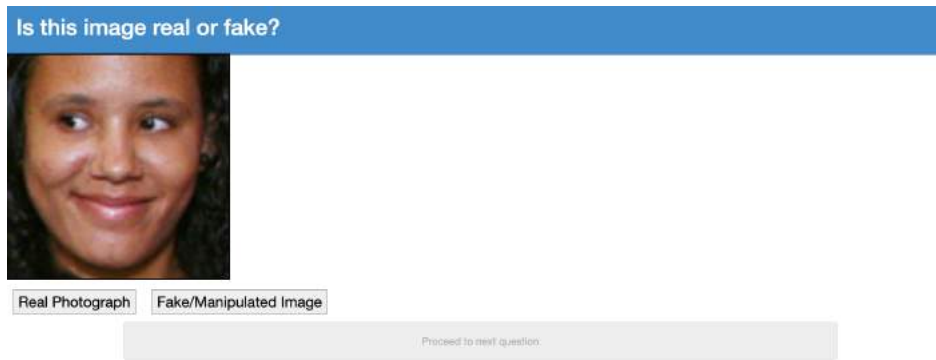


Fig. 3: Quality Classification AMT Interface Example.

was real or fake. Fig. 4 shows the racial distribution of the generated images that were classified as real/high quality for the three data splits.

In Fig. 4 we observe that the race class ratio of the training images has the same race class ratio as the training data. The race ratio of the 20B-80W, 50B-50W, and 80B-20W training data, respectively, has a race ratio of 0.25, 1, and 4, and the corresponding generated data race ratios of the high quality labeled images are 0.26, 1.1, and 3.8.

2.3 Quality Ranking

The Quality Ranking AMT task asked the workers to identify the Fake image between two images by selecting one of the options described below:

1. Image A - Image A is more likely to be a Fake image.




Fig. 4: Distribution of the Generated Images that were labeled as “real”, or high quality, images by annotators in the three FairFace data splits (20B-80W, 50B-50W and 80B-20W) with a corresponding table showing the race ratio (Black/white) of the different data splits. This shows that the race ratio of the training data is relatively the same as the race ratio of the generated images that were labeled Real.

2. Image B - Image B is more likely to be a Fake image.

To further assist the workers in understanding what Fake images are, the same definitions from above were given to the workers at the start of the task. Fig. 5 is an example of the Quality Ranking AMT interface.



Fig. 5: Quality Ranking AMT Interface Example.

3 Race Classification Performance Analysis

In this section we expand on the performance of the AMT annotation compared to the automatic perceived race classifier.

Our perceived race classifier obtained an accuracy of 84%, treating our FairFace labels as ground truth, on the same 1000 images used for human annotation. This gave us more confidence on the performance and validity of the race classifier and its role as a proxy for AMT annotation when conducting our experiments.

The classifier performs better on Black and Other class labels compared to the white class label. The classifier tends to classify white faces as Other which is also slightly observed in the AMT annotations.

Comparing the performance of the human annotation and the perceived race classifier we see that they are both aligned in terms of classifying the different race class labels, and therefore the automatic classifier can be used as a proxy for human annotation. Overall, the classifier outperforms the human annotations. We hypothesize that this could be due to subjective bias present in human annotation, or to the subjective nature of perceived race classification.

4 Experimental Results

4.1 Relationship between Training and Generated Data Distributions

In this section, we expand on the results that demonstrate that StyleGAN2-ADA’s generated data distribution preserves the training data distribution. In the paper, we excluded the “Non-Black or Non-white” and “Cannot Determine” class labels in the generated data to explicitly showcase the ratio of Black and white race class labels in the training and generated data. Fig. 6 (left) shows the generated data distribution with all the class labels where the actual number of the classes are in bracket in the pie charts and Fig. 6 (right) shows distribution for when the “Non-Black or Non-white” and “Cannot Determine” class labels were excluded. To get the generated data distribution with “Non-Black or Non-white” and “Cannot Determine” class labels, these two class labels were dropped and the distribution was recalculated with only white and Black class labels.

4.2 Truncation

In order to evaluate properties of truncation, images were generated from StyleGAN2-ADA trained on FFHQ and the 80B-20W, 50B-50W and 80B-20W FairFace-trained generators at truncation levels ranging from $\gamma = 0$ to 1 at intervals of 0.1. Random samples from FFHQ with various levels of truncation can be seen in Fig. 7. As the level of truncation increases, the ratio of faces of people of color to white faces decreases.




Fig. 6: Racial distribution of training and generated data. Distributions for 20B-80W, 50B-50W and 80B-20W data splits for (left) all the class labels and (right) Black and white class labels only. This figure shows that all of the generative models preserve interval distribution of the training data. The red bars represent the 95% Wald confidence interval (CI) of the generated data. See Table 1 for the corresponding CI.

4.3 Quality Ranking

The raw intra-split pairwise comparison results from our quality ranking experiments can be seen in Tables 2, 3 and 4. From these pairwise comparisons and the inter-split comparisons, we use the `choix` package [?] to produce a Bradley-Terry model that ranks the 3000 images in order of highest quality to lowest perceived image quality. From this global ranking, we visualize the top 25 and bottom 25 images, and also random images from each quartile, which can be seen in Fig. 8.

The raw counts for our top K image compositions per race class label can be seen in Table 6. In order to obtain weighted percentage scores seen in the main paper, the raw counts for the top K images of a particular race class label were first weighted by the expected frequency of the images in each split. This was done by multiplying the raw count by $\frac{1}{2}$, $\frac{1}{5}$, and $\frac{1}{8}$ if the race class label comprised 20%, 50% or 80% of the dataset, respectively. Then, the weighted numbers were divided by the sum of all scores for that particular race class label and given value of K .

Table 1: Wald’s 95% confidence interval (CI) of generated data from generators trained on FFHQ and 80B-20W, 50B-50W and 80B-20W FairFace .

Generated Data	Black CI @95%	white CI @95%
FFHQ	5.73 ± 1.7	94.34 ± 1.7
20B-80W	20.70 ± 3.1	79.3 ± 3.1
50B-50W	51.3 ± 3.7	48.7 ± 3.7
80B-20W	78.4 ± 3.1	21.6 ± 3.1

Table 2: **Intra-split pairwise perceived image quality comparison 20B-80W Dataset.**

More Preferred	Less Preferred	Count
white	white	2367
white	Black	708
white	Other	982
white	CD	1129
Black	Black	186
Black	white	465
Black	Other	205
Black	CD	270
Other	Other	411
Other	white	824
Other	Black	269
Other	CD	509
CD	CD	225
CD	Black	87
CD	white	230
CD	Other	133

A Bradley-Terry model [?] predicts the probability that a pairwise comparison $i > j$ is true. A ranking of all items can be derived by modeling the probability for pairs in a dataset. The `choix` package [?] produces a Bradley-Terry model by using the Iterative Luce Spectral Ranking algorithm [?]. This algorithm performs maximum-likelihood inference to rank items from a dataset of pairwise comparisons.

Table 3: **Intra-split pairwise perceived image quality comparison 50B-50W Dataset.**

More Preferred	Less Preferred	Count
white	white	990
white	Black	1305
white	Other	497
white	CD	756
Black	Black	1167
Black	white	888
Black	Other	448
Black	CD	578
Other	Other	252
Other	white	475
Other	Black	599
Other	CD	323
CD	CD	159
CD	Black	274
CD	white	201
CD	Other	88

Table 4: **Intra-split pairwise perceived image quality comparison of 80B-20W Dataset.**

More Preferred	Less Preferred	Count
white	white	195
white	Black	798
white	Other	216
white	CD	319
Black	Black	2595
Black	white	660
Black	Other	705
Black	CD	1064
Other	Other	216
Other	white	216
Other	Black	918
Other	CD	295
CD	CD	219
CD	Black	421
CD	white	83
CD	Other	80

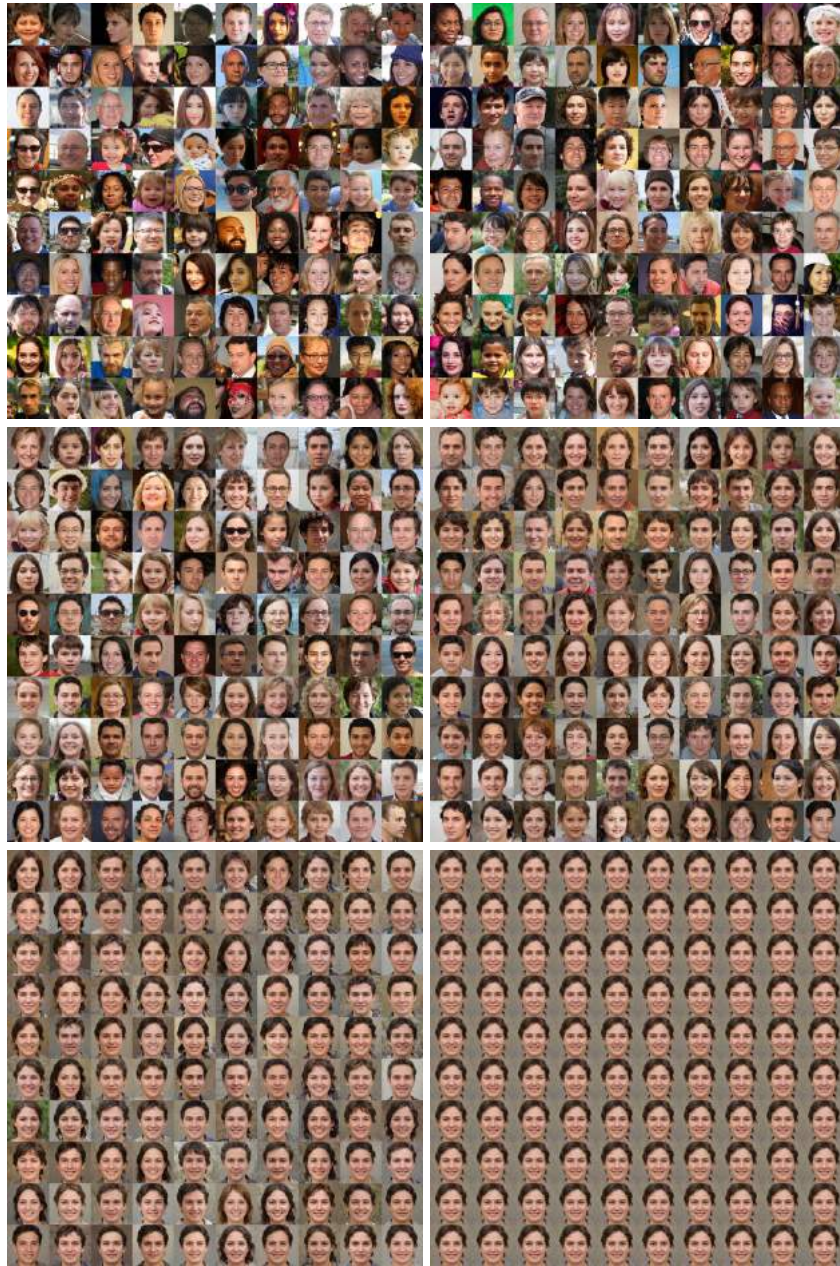


Fig. 7: **FFHQ samples with various levels of truncation.** (top left) truncation of $\gamma = 1$ (no truncation), (top right) truncation of $\gamma = 0.8$, (middle left) truncation of $\gamma = 0.6$, (middle right) $\gamma = 0.4$, (bottom left) truncation of $\gamma = 0.2$, and (bottom right) $\gamma = 0$ (full truncation). As the amount of truncation increases, racial diversity decreases, resulting in an increasingly larger proportion of white faces.

Table 5: **Race label breakdown of global ranking.** For each quartile of the global ranking of 3000 FairFace generated images compiled from different data splits, the percentage of faces annotated as white, Black, non Black/white, and Cannot Determine. white faces are over-represented in the top half of the quality ranking, and under-represented in the bottom half.

Quartile	white %	Black %	Non Black/white %	Cannot Determine %
Top	48.4	26.3	22.5	2.8
Second	40.0	37.7	19.3	4.0
Third	29.7	41.3	17.2	11.7
Bottom	15.2	32.1	11.6	41.1

Table 6: **Top K image composition per-race, raw counts.** Given a ranking of Black and white images across all data splits, we break down the data split that each image came from. The highest quality images ($K = 10, 25, 50$) are more likely to come from a data split where they are over-represented or represented in parity. These are the raw numbers, before normalization.

white				Black			
K	80B-20W	50B-50W	20B-80W	K	80B-20W	50B-50W	20B-80W
10	0	2	8	10	7	2	1
25	0	7	18	25	19	4	2
50	4	16	30	50	31	16	3
100	10	39	51	100	57	31	12
500	68	181	251	500	275	161	64




Fig. 8: Results of global quality ranking across the three FairFace data splits. Each row represents a particular image quality ranking: (top left) top 25 images in the quality ranking, (top right) samples from the top quartile of images, (middle left) samples from the second top quartile of images, (middle right) third quartile of ranked images, (bottom left) bottom quartile of ranked images, and (bottom right) the bottom 25 images in the quality ranking. As shown in Table 5, white faces are over-represented in the top half of the quality ranking and under-represented in the bottom half.