Mike Laszkiewicz Mike.Laszkiewicz@rub.de Faculty of Computer Science, Ruhr University Bochum, North Rhine-Westphalia Germany Imant Daunhawer Department of Computer Science, ETH Zurich Zurich, Zurich, Switzerland Julia E. Vogt* Department of Computer Science, ETH Zurich Zurich, Zurich, Switzerland

Asja Fischer* Faculty of Computer Science, Ruhr University Bochum, North Rhine-Westphalia Germany Johannes Lederer* Department of Mathematics, Computer Science, and Natural Sciences, University of Hamburg Hamburg, Hamburg, Germany



Figure 1: Upsampling algorithms fail to reconstruct samples from minority groups: Reconstructing low-resolution images using models trained on UnfairFace, a subset of the FairFace dataset replicating the racial distribution of common large-scale datasets.

*Joint supervision.



This work is licensed under a Creative Commons Attribution International 4.0 License.

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil

ABSTRACT

Recent years have witnessed a rapid development of deep generative models for creating synthetic media, such as images and videos. While the practical applications of these models in everyday tasks are enticing, it is crucial to assess the inherent risks regarding

© 2024 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0450-5/24/06 https://doi.org/10.1145/3630106.3658921 their fairness. In this work, we introduce a comprehensive framework for benchmarking the performance and fairness of conditional generative models. We develop a set of metrics—inspired by their supervised fairness counterparts—to evaluate the models on their fairness and diversity. Focusing on the specific application of image upsampling, we create a benchmark covering a wide variety of modern upsampling methods. As part of the benchmark, we introduce UnfairFace, a subset of FairFace that replicates the racial distribution of common large-scale face datasets. Our empirical study highlights the importance of using an unbiased training set and reveals variations in how the algorithms respond to dataset imbalances. Alarmingly, we find that none of the considered methods produces statistically fair and diverse results. All experiments can be reproduced using our provided repository.¹

CCS CONCEPTS

• Computing methodologies \rightarrow Reconstruction.

KEYWORDS

Conditional Generative Models, Computer Vision, Image Upsampling, Fairness

ACM Reference Format:

Mike Laszkiewicz, Imant Daunhawer, Julia E. Vogt, Asja Fischer, and Johannes Lederer. 2024. Benchmarking the Fairness of Image Upsampling Methods. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24), June 03–06, 2024, Rio de Janeiro, Brazil.* ACM, New York, NY, USA, 29 pages. https://doi.org/10.1145/3630106.3658921

1 INTRODUCTION

Remarkable advancements in deep generative models (e.g. [18, 28, 47, 51]) give rise to striking applications, including image inpainting (e.g. [46]), style transfer (e.g. [70]), image superresolution (e.g. [7]), and text-to-image translation (e.g. [45]). Despite their apparent success, our understanding of these models and the biases they are subject to lags behind, often revealing unexpected fairness violations. Earlier research on fairness in deep learning discovered biases in supervised prediction models, leading to racist [9, 71], sexist [32, 53], and homophobic [40] decisions, which can be critical in many applications, such as criminal justice, medicine, hiring, or personalized advertising. These observations sparked a general interest in a research field that concentrates on developing and quantifying fair machine learning models [5, 11, 44, 71]. However, compared to the supervised setting, there is a noticeable gap in addressing fairness in generative models [8, 20] and conditional generative models [24, 56].

In this work, we focus on the latter and investigate conditional generative models regarding their fairness and diversity. We introduce a novel set of metrics that serves as a practical tool for assessing fairness violations of conditional generative models. These metrics are associated with the fairness definitions introduced by Jalal et al. [24] in the sense that fairness is achieved if and only if our fairness violation metrics equal zero. These metrics form the foundation of a comprehensive framework that supplements standard performance criteria with fairness and diversity metrics. Based on these metrics, we design a versatile benchmark, facilitating a

unified evaluation of the performance, fairness, and diversity of a wide variety of conditional generative models. We showcase the capacities of the proposed framework by evaluating one particular type of conditional generative modeling, namely image upsampling, which has triggered a compelling discussion on racial biases in deep generative models [58]. We leverage the FairFace dataset [25] to introduce UnfairFace, a subset of FairFace mimicking the distribution of common large-scale face datasets. Because FairFace provides race² labels, it allows us to conduct an in-depth empirical evaluation of a range of upsampling methods and their influence on dataset bias. Our evaluation demonstrates that, while none of the methods are significantly fair, the degree of fairness varies significantly across methods. Unsurprisingly, we find that a racial bias in the training data reduces the fairness of most models. Notably, the effect is most apparent for Denoising Diffusion Restoration Models (DDRM [31]), which exhibit the highest discrepancy to fairness among all investigated models when trained on UnfairFace, but the least discrepancy to fairness when trained on FairFace. In contrast, the fairness discrepancy of Posterior Sampling [24] remains robust despite biased training data. These findings highlight the importance of careful data acquisition and emphasize how biases present in the dataset can impact model selection.

Notation. In this work, we define performance and fairness using probabilistic terminology. Specifically, we denote target samples by $x \in X$ and conditional arguments by $y \in \mathcal{Y}$. Target samples and conditions follow the distribution $\mathbb{P}_{x,y}$ with marginals \mathbb{P}_x and \mathbb{P}_y , respectively. $\mathbb{P}_{y|x}$ denotes the conditional distribution of y given a target x and $\mathbb{P}_{x|y}$ the conditional distribution of x given a condition y. The conditional generative model generates samples \hat{x} conditioned on y, which follows the distribution $\mathbb{P}_{\hat{x}|y}$. We denote the distribution of $\mathbb{P}_{\hat{x}|y}$ marginalized over all y by $\mathbb{P}_{\hat{x}}$. Moreover, we denote by $\mathbb{P}_{\hat{x}|x}$ the conditional distribution that one obtains by marginalizing the product $\mathbb{P}_{\hat{x}|y} \cdot \mathbb{P}_{y|x}$ over all y. For the sake of clarity, we use the notation $\mathbb{P}(\hat{x} \in A | x \in B) := \mathbb{P}_{\hat{x}|x}(A|B)$ for measurable sets A, B in the domain of \hat{x} and x, respectively. We use a similar notation for all other distributions introduced in this paragraph.

2 RELATED WORK

2.1 Fairness for Supervised Models and Unconditional Generative Models

The first investigations on fairness of machine learning models were conducted for supervised models[36]. Research in that field has led to several approaches, such as fair representation learning [4, 50, 63, 67], constrained optimization [1, 42], adversarial training [13], and post-processing methods [10, 17]. These works introduced numerous different definitions of fairness—for instance, demographic parity [12] and equalized odds [17]—that guarantee that the model predictions are not driven by sensitive features. These definitions do not apply to unconditional generative models, whose aim is to generate synthetic data. Instead, fairness for unconditional generative models is typically phrased in terms of feature

¹https://github.com/MikeLasz/Benchmarking-Fairness-ImageUpsampling

 $^{^{2}}$ We specify race in terms of phenotype and appearance, and ethnicity as a shared culture [3, 14]. We are aware that race and ethnicity are social constructs and that our definitions do not reflect the underlying complexities apparent in society. An elaborate ethics statement is provided in Section 8.

diversity, which is often violated in practice. For instance, deep generative models have been observed to generate less diverse samples [23], which is amplified by common practices such as the truncation trick [35]. Moreover, empirical results demonstrate that an increase in model performance correlates with a decrease in sample diversity [49, 66]. To resolve the issue, given a fair reference dataset, Choi et al. [8] and Zameshina et al. [66] enhance the diversity by training a model with importance weighting and stratification, Yu et al. [65] introduce another training loss based on implicit maximum likelihood, Teo et al. [57] apply modern transfer learning techniques to the reference dataset, whereas other works refine the diversity of the generative models by latent code manipulation [20, 55].

2.2 Fairness for Conditional Generative Models

Conditional generative models create a new sample \hat{x} given a condition y, which could, for instance, be a low-resolution image in image upsampling (e.g. Menon et al. [37]), a cropped image in image inpainting (e.g. Richardson et al. [46]), or a sentence in text-to-image synthesis (e.g. Rombach et al. [47]). Since the condition y specifies certain features that should be reflected in the generated samples, it might be too restrictive to assess the fairness of a generative model in terms of sample diversity. Doing so would ignore the coherence between the condition and the generated samples. To address this, Jalal et al. [24] extend traditional group-fairness [12] measures for conditional generative models to define *representation demographic parity* (RDP), *proportional representation* (PR), and *conditional proportional representation* (CPR). In particular, let $C = \{C_1, \ldots, C_k\}$ be a partition³ of the data space X, defining k classes. Then, RDP, PR, and CPR are satisfied if and only if

$$\mathbb{P}(\hat{x} \in C_i | x \in C_i) = \mathbb{P}(\hat{x} \in C_j | x \in C_j) \quad \forall i, j \in [k], \forall x \in \mathcal{X}, (1)$$

$$\mathbb{P}(\hat{x} \in C_i) = \mathbb{P}(x \in C_i) \quad \forall i \in [k] \quad , \tag{2}$$

$$\mathbb{P}(\hat{x} \in C_i | y) = \mathbb{P}(x \in C_i | y) \quad \forall i \in [k], \, \forall y \in \mathcal{Y} \quad , \tag{3}$$

respectively, where \mathcal{Y} denotes the space of possible conditions and $[k] := \{1, \ldots, k\}$. We provide further explanation and synthetic examples highlighting the difference between these definitions in Section 9.1.

For instance, in face image upsampling we could have the highresolution image manifold X, the low-resolution image manifold \mathcal{Y} , and a partition $C := \{C_{\text{glasses}}, C_{\text{no-glasses}}\}$ into a set of people that do or do not wear glasses. In this specific example, Tanjim et al. [56] approximates equation (1) by measuring the performance of a classifier that predicts whether or not a given image depicts a person wearing glasses.

2.3 Image Upsampling Methods

In image upsampling, the goal is to create a high-resolution image given a low-resolution input. To make the notation more illustrative in this case, we denote the condition by x_{LR} , the true highresolution image by x_{HR} , and the conditionally generated output by \hat{x}_{HR} . A myriad of image upsampling methods have been introduced in recent years; we focus on five carefully selected methods representative of most previously proposed methods. To cover a diverse set of approaches, we select PULSE [37], which performs gradient-based latent-code optimization of a pretrained GAN [22], Pixel2Style2Pixel [46], which utilizes an encoder to obtain latentcodes of a pretrained GAN, and denoising diffusion restoration models [31] based on diffusion models [64]. Furthermore, we include two models that were specially designed for fair image upsampling: Posterior Sampling [24] and a fairness-improved version of Pixel2Style2Pixel [56], which we refer to as fair-pSp. In the following, we provide a brief overview of these methods.

Upsampling in PULSE is conducted by fixing a pretrained generative model G and then optimizing⁴ for a latent code $z^* :=$ $\arg \min_{z} ||x_{\text{LR}} - DS(G(z))||$, where x_{LR} is the low-resolution input and DS is a downscaling operator. The resulting high-resolution input is given by $\hat{x}_{HR} := G(z^*)$. In contrast, Pixel2Style2Pixel (pSp) learns an encoder *E* that can directly map from x_{LR} to the latent code $z^* = E(x_{LR})$ and returns $\hat{x}_{HR} := G(z^*)$. To improve the fairness of pSp during its training, Tanjim et al. [56] propose to use a stratified sampling scheme, contrastive learning, and a cross-entropy regularization based on the predicted fairness labels. Posterior Sampling generates samples by sampling from the posterior $p(x_{\text{HR}}|x_{\text{LR}}) \propto p(x_{\text{LR}}|x_{\text{HR}})p(x_{\text{HR}})$. Jalal et al. [24] implement Posterior Sampling by leveraging a NCSNv2 score-based unconditional model [54], which gives access to $\nabla_{x_{\text{HR}}} \log p(x_{\text{HR}})$. They model $p(x_{LR}|x_{HR})$ by fixing a linear downscaling operation A and modeling $x_{LR} = Ax_{HR} + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma I)$ for some $\sigma > 0$. This allows to model the score function of $p(x_{HR}|x_{LR})$ by $\nabla_{x_{\text{HR}}} \log p(x_{\text{HR}}|x_{\text{LR}}) = \nabla_{x_{\text{HR}}} \log p(x_{\text{LR}}|x_{\text{HR}}) + \nabla_{x_{\text{HR}}} \log p(x_{\text{HR}})$ and to sample using Langevin dynamics [61]. Denoising diffusion restoration models (DDRM) also attempt to sample from $p(x_{\text{HR}}|x_{\text{LR}})$, but, instead of running the diffusion process in pixel space, they propose running it in the spectral space of the downscaling operator A. This makes DDRM remarkably efficient, requiring only around 20 diffusion steps in the reported experiments.

3 BENCHMARKING FAIRNESS OF CONDITIONAL GENERATIVE MODELS

In this section, we introduce a collection of metrics that allow us to quantify the performance and fairness of conditional generative models. Although evaluating a model's performance is standard practice, investigating its fairness and diversity is often neglected. Hence, our framework expands conventional evaluation methodologies by systematically examining the fairness and diversity of conditional generative models. The following metrics are specifically designed for image data but can readily be generalized to other data types.

3.1 Performance

We categorize the performance criteria into three types: expected reconstruction losses, referenceless quality losses, and expected attribute reconstruction losses. While the first two groups are common, the expected attribute reconstruction loss is rarely considered.

Formally, the expected reconstruction loss is defined as

$$\mathbb{E}_{x,y\sim\mathbb{P}_{x,y},\hat{x}\sim\mathbb{P}_{\hat{x}|y}}\left[L_{\text{rec}}(x,\hat{x})\right] \approx \frac{1}{n} \sum_{i=1}^{n} L_{\text{rec}}\left(x^{(i)},\hat{x}^{(i)}\right) ,$$

³Note, the metrics also generalize to non-disjoint covers of X.

⁴The optimization procedure is accompanied by several regularization terms.

where $x^{(i)}$, $y^{(i)}$ are samples from $\mathbb{P}_{x,y}$, $\hat{x}^{(i)}$ are samples from $\mathbb{P}_{\hat{x}|y^{(i)}}$, and L_{rec} is a reconstruction loss. For image data, we choose L_{rec} to be either the LPIPS distance [68], or the structural dissimilarity

$$DSSIM(x, \hat{x}) \coloneqq \frac{1 - SSIM(x, \hat{x})}{2} \in [0, 1]$$

where SSIM denotes the structural similarity index measure [60].

In certain conditional generation tasks, multiple outputs x align with a given condition y. For instance, in text-to-image synthesis multiple images x may depict images that correspond to an input text y, rendering a reconstruction loss inappropriate. Therefore, to account for such settings, we measure the expected referenceless quality loss

$$\mathbb{E}_{y \sim \mathbb{P}_{y}, \, \hat{x} \sim \mathbb{P}_{\hat{x}|y}} \left[L_{\text{qual}}(\hat{x}) \right] \approx \frac{1}{n} \sum_{i=1}^{n} L_{\text{qual}}\left(\hat{x}^{(i)} \right) \;,$$

where $y^{(i)}$ are samples from \mathbb{P}_y , $\hat{x}^{(i)}$ are samples from $\mathbb{P}_{\hat{x}|y^{(i)}}$, and L_{qual} is a quality assessment metric. For image data, we set L_{qual} to be either the NIQE score [39] or the negative of the blurriness index introduced by Pech-Pacheco et al. [43], for which in both cases a lower score indicates a better image quality.

The last performance metric focuses on the reconstruction quality of certain categorical attributes. Given a classifier f_{att} of these attributes, we define the expected attribute reconstruction loss

$$\mathbb{E}_{x,y\sim\mathbb{P}_{x,y},\hat{x}\sim\mathbb{P}_{\hat{x}|y}}\left[L_{\text{att}}(x,\hat{x};f_{\text{att}})\right] \approx \frac{1}{n}\sum_{i=1}^{n}L_{\text{att}}\left(x^{(i)},\hat{x}^{(i)};f_{\text{att}}\right) ,$$

where L_{att} is an attribute-related loss. One choice for L_{att} is the binary attribute prediction loss, i.e.,

$$L_{\text{att}}^{0-1}(x, \hat{x}; f_{\text{att}}) := \begin{cases} 1 & , & \text{if } f_{\text{att}}(x) \neq f_{\text{att}}(\hat{x}) \\ 0 & , & \text{if } f_{\text{att}}(x) = f_{\text{att}}(\hat{x}) \end{cases}$$
(4)

To relax the hard 0–1 penalty in the binary prediction loss, we also compute the cosine similarity of the latent representations of x and \hat{x} , that is,

$$L_{\text{att}}^{\cos}(x,\hat{x};f_{\text{att}}) := \frac{\langle l_x, l_{\hat{x}} \rangle}{\|l_x\| \cdot \|l_{\hat{x}}\|} , \qquad (5)$$

where l_x and $l_{\hat{x}}$ denote the last activations⁵ of $f_{\text{att}}(x)$ and $f_{\text{att}}(\hat{x})$, respectively. While attribute reconstruction can be related to LPIPS, using an attribute-specific classifier allows for a more precise evaluation with respect to the labeled attributes.

3.2 Fairness and Diversity

There are no unified metrics for assessing the amount of fairness in unconditional models. For instance, Tanjim et al. [56] use the expected attribute reconstruction loss as a heuristic for fairness, while Jalal et al. [24] evaluate fairness by visual inspection of confusion matrices. Therefore, to formalize the amount of fairness in a principled way, we extend the definitions of RDP (1) and PR (2) through divergence measures between probability distributions. Let \mathbb{P}_{RDP} be the distribution quantifying the probability that the conditional sample \hat{x} and the original sample x belong to the class Laszkiewicz et al.

c. That is, for a classifier f_{att} that predicts the class $c \in \{c_1, \ldots, c_k\}$ of *x*, we define

$$\mathbb{P}_{\text{RDP}}(c = c_j) := \frac{\mathbb{P}(\hat{x} \in C_j | x \in C_j)}{\sum_{l=1}^{k} \mathbb{P}(\hat{x} \in C_l | x \in C_l)} \\
\approx \frac{\frac{1}{n_j} \sum_{i \in I_j} L_{\text{att}}^{0-1}(x^{(i)}, \hat{x}^{(i)}; f_{\text{att}})}{\sum_{l=1}^{k} \frac{1}{n_l} \sum_{i \in I_l} L_{\text{att}}^{0-1}(x^{(i)}, \hat{x}^{(i)}; f_{\text{att}})} \quad \forall j \in [k] ,$$
(6)

where C_l is the set of samples corresponding to class c_l , $I_l := \{i \in [n] : x^{(i)} \in C_l\}$, and $n_l = |I_l|$. According to (1), an algorithm satisfies RDP if and only if $\mathbb{P}_{\text{RDP}}(c = c_j) = \mathbb{P}_{\text{RDP}}(c = c_i)$ for all $i, j \in [k]$, or equivalently, if and only if $\mathbb{P}_{\text{RDP}} = \mathcal{U}([k])$, i.e., if \mathbb{P}_{RDP} is uniformly distributed over [k]. Hence, to quantify the amount of violation of fairness, we introduce

$$\Delta_{\text{RDP}} \coloneqq D(\mathbb{P}_{\text{RDP}} \parallel \mathcal{U}([k])) \quad , \tag{7}$$

where *D* is a divergence measure. We set *D* to the Pearson χ^2 -divergence or the Chebyshev-distance to obtain

$$\Delta_{\text{RDP-}\chi^2} := \sum_{j=1}^k \frac{(\mathbb{P}_{\text{RDP}}(c=c_j) - 1/k)^2}{1/k}$$
$$= k \sum_{j=1}^k \left(\mathbb{P}_{\text{RDP}}(c=c_j) - \frac{1}{k} \right)^2 , \qquad (8)$$

$$\Delta_{\text{RDP-Cheb}} := \max_{j \in [k]} \left| \mathbb{P}_{\text{RDP}}(c = c_j) - \frac{1}{k} \right| . \tag{9}$$

By empirically estimating (6), we obtain plug-in estimates of $\Delta_{\text{RDP-}\chi^2}$ and $\Delta_{\text{RDP-Cheb}}$. One could think of other divergences, such as KLdivergence or Total-Variation-divergence but we decided on the χ^2 -divergence to relate the scores to the test statistics of a Pearson's χ^2 -test. This test allows us to statistically test the hypothesis $\mathbb{P}_{\text{RDP}} = \mathcal{U}([k])$. Furthermore, we propose to consider the Chebyshev distance, since it can be interpreted as the maximum violation of fairness.

However, note that considering RDP in isolation might be insufficient for fairness assessment as demonstrated by the following example. Let the conditional generative model be such that $\mathbb{P}(\hat{x} \in C_j | x \in C_j) = 0.5$ for all $j \in [k]$, and therefore $\mathbb{P}_{\text{RDP}}(c = c_j) = 1/k$, thus satisfying RDP. Let us further assume that all misclassified samples belong to C_1 , i.e., $\hat{x} \in \{C_j, C_1\}$ if $x \in C_j \neq C_1$. It follows that $\mathbb{P}(\hat{x} \in C_1 | x \in C_j) = 1 - \mathbb{P}(\hat{x} \in C_j | x \in C_j) = 0.5$ for all $j \neq 1$ and consequently, it holds that

$$\mathbb{P}(\hat{x} \in C_1) = \sum_{j=1}^k \mathbb{P}(\hat{x} \in C_1 | x \in C_j) \mathbb{P}(x \in C_j) = \sum_{j=1}^k 0.5 \cdot \frac{1}{k} = 0.5 ,$$

under the assumption that $\mathbb{P}(x \in C_j) = 1/k$ for all $j \in [k]$. Hence, even though the conditional generative model is fair with respect to RDP, it has an unreasonable bias for producing samples from C_1 more frequently than samples from $C_j \neq C_1$ if k > 2. We provide further simplified synthetic examples highlighting the differences between RDP and PR in Example 9.1 of the Appendix. This motivates us to complement RDP with quantitative measures of PR. As above in (6), we define a discrete distribution \mathbb{P}_{PR} over classes

 $^{^5\}mathrm{Note}$ that by incorporating activations, we implicitly assume f_{att} to be a neural network classifier.

 $\{c_1, ..., c_k\}$ by

$$\mathbb{P}_{\text{PR}}(c=c_j) \coloneqq \mathbb{P}(\hat{x} \in C_j) \approx \frac{1}{n} \sum_{i=1}^n \mathbb{1}\left[\hat{x}^{(i)} \in C_j\right] \ \forall j \in [k]$$

where 1 is the indicator function. In practice, we use the attribute classifier f_{att} to quantify whether $\hat{x}^{(i)} \in C_j$. Once again, PR is satisfied if and only if $\mathbb{P}_{\text{PR}} = \mathcal{U}([k])$, which motivates to introduce $\Delta_{\text{PR-}\chi^2}$ and $\Delta_{\text{PR-Cheb}}$ analogous to their RDP-versions in (8) and (9), respectively.

A different viewpoint of fairness is that of diversity in the presence of conditional arguments that are uninformative for certain attributes. For illustration, let us consider the low-resolution sample for image upsampling depicted in Figure 4. This sample does not contain any information regarding its race, and hence, a diverse upsampling algorithm should produce high-resolution samples without favoring any race over the other. We quantify that intuition by building upon CPR (3) and defining $\mathbb{P}_{CPR+\mu}$ by

$$\mathbb{P}_{\text{CPR}|y}(c=c_j) := \mathbb{P}(\hat{x} \in C_j|y) \approx \frac{1}{n_x} \sum_{i=1}^{n_x} \mathbf{1} \left[\hat{x}^{(i)} \in C_j \right] \ \forall j \in [k]$$

for n_x conditional samples $\hat{x}^{(i)}$ based on a fixed condition y. Certainly, we are not interested in the diversity of $\mathbb{P}_{\text{CPR}|y}$ for any y but only for non-informative conditions y. Hence, we define the *uninformative conditional proportional representation* (UCPR) distribution \mathbb{P}_{UCPR} by taking the expectation of $\mathbb{P}_{\text{CPR}|y}$ over all uninformative y, that is,

$$\mathbb{P}_{\text{UCPR}}(c = c_j) := \mathbb{E}_{y \sim U} \left[\mathbb{P}_{\text{CPR} \mid y}(c = c_j) \right] \\ \approx \frac{1}{n_y} \sum_{m=1}^{n_y} \frac{1}{n_x} \sum_{i=1}^{n_x} \mathbf{1} \left[\hat{x}^{(i)}(y^{(m)}) \in C_j \right] \; \forall j \in [k] \; ,$$
(10)

where *U* is a random variable that generates uninformative *y*, $y^{(1)}, \ldots, y^{(n_y)}$ are n_y samples from *U*, and $\hat{x}^{(i)}(y^{(m)})$ denotes the *i*th output of the conditional generative model given the condition $y^{(m)}$. We say that the conditional generative model satisfies UCPR if and only if $\mathbb{P}_{\text{UCPR}}(c = c_j) = \mathbb{P}_{\text{UCPR}}(c = c_i)$ for all $i, j \in [k]$. As above, we define the plug-in estimators $\Delta_{\text{UCPR}-\chi^2}$, $\Delta_{\text{UCPR-Cheb}}$ to measure the diversity of the conditional model in the presence of uninformative inputs. Generally, the generation of conditions that are uninformative regarding the classes *C* depends on the application. In the following section, we explain how we obtain uninformative samples in the specific case of image upsampling.

Conveniently, since the introduced fairness metrics are evaluations of a divergence (see e.g. (7)), they satisfy $D(\mathbb{P} || \mathbb{Q}) = 0$ if and only if $\mathbb{P} = \mathbb{Q}$. Consequently, we conclude that RDP, PR, and UCPR of a conditional generative model can be verified using our proposed metrics.

Corollary 3.1. A conditional generative model satisfies

(1) *RDP if and only if* $\Delta_{\text{RDP} - \gamma^2} = \Delta_{\text{RDP-Cheb}} = 0$;

(2) *PR if and only if* $\Delta_{PR-\chi^2} = \Delta_{PR-Cheb} = 0$;

(3) UCPR if and only if $\Delta_{\text{UCPR}-\chi^2} = \Delta_{\text{UCPR-Cheb}} = 0$.

4 INTRODUCING UNFAIRFACE

In the further course of this work, we apply the framework derived in Section 3 to evaluate the performance, fairness, and diversity of a particular type of conditional generative model, namely for the application of image upsampling on human face images. Specifically, we empirically analyze the upsampling performance, but also evaluate the fairness and diversity of the reconstructions using race labels. Unfortunately, standard face datasets used for generative models-such as CelebA [34] and FFHQ [29]-do not come with race labels. On the other hand, datasets that come with race labels like PPB [6] and UTKFace [69] are often too small to be used to train a (conditional) generative model. Suited for our experiments are the BUPT [59] and FairFace [25] datasets, which both provide sufficient examples to train a generative model and include race labels. However, the BUPT database is not consistent in the shape and size of the images and contains only celebrities. Therefore, we decided to base our experiments on the FairFace dataset, which consists of around 87k training samples of resolution 224×224 . The samples are labeled by race (7 categories), which are approximately uniformly distributed across the races, see Figure 8(a).

While it is often desirable to train on a balanced dataset, having access to such a dataset is rare in practice. Instead, existing public face datasets are usually constructed from online sources, which are biased towards faces that have a light skin tone [38]. In this study, we aim to investigate the effect of such a dataset bias on image upsampling algorithms. Hence, to mimic the bias apparent in large-scale datasets while having access to race labels, we subsample from the FairFace dataset such that it imitates the racial distribution of CelebA⁶ (see Figure 8(b)). The resulting dataset, which we refer to as *UnfairFace*, consists of around 20k samples of which more than 80% are labeled as "White".

For testing, we take a subset of the test set of FairFace consisting of 200 samples per race, totaling a test-set size of 1400 images. We base our expected attribute reconstruction loss on race reconstruction, i.e., f_{att} in (4) and (5) is a race classifier. Note, that we do not have access to an unambiguous, ground-truth race classifier; instead, we leverage the pretrained race classifier provided by Kärkkäinen and Joo [25]. To diminish the effect of the imperfect classifier on the evaluations, we select only the first 200 samples per race that are correctly labeled by the classifier. To emphasize the fact that in the following we focus on race reconstruction, we denote the losses in (4) and (5) by $L_{\text{race}}^{\text{cos}}$ and $L_{\text{race}}^{\text{cos}}$, respectively and henceforth use the notation x_{HR} , \hat{x}_{HR} , x_{LR} introduced in Section 2.3.

5 EXPERIMENTS

In this section, we apply our framework derived in Section 3 for the evaluation of image upsampling algorithms. To assess the bias of training on an unbalanced population with multiple underrepresented races, we compare upsampling methods trained on the UnfairFace dataset to the same methods trained on the original FairFace dataset with balanced subpopulations. All models are evaluated on the same holdout set with balanced races. The experiments can be reproduced using the official code repository.

 $^{^6{\}rm The}$ approximate racial distribution is taken from Figure 1 in Kärkkäinen and Joo [25].

Laszkiewicz et al.



(a) Models trained on UnfairFace.



(b) Models trained on FairFace.

Figure 2: Upsampling results for models trained on UnfairFace and FairFace using test samples categorized as "Black".

Experimental setup. Many architectures, such as StyleGAN [27–30] are designed for outputs of resolution $2^k \times 2^k$ for some $k \in \mathbb{N}$. Since FairFace contains images of resolution 224×224 , we trained all image upsampling models to generate outputs of the next-smallest power of two, which is 128×128 . For testing, we first downsample each test sample to resolution 128×128 using bilinear downscaling with antialiasing to obtain x_{HR} . To get x_{LR} , we downsample x_{HR}

using the same downscaling. For the experiments evaluating performance and fairness, we set x_{LR} to resolution 16 × 16. For evaluating the diversity, we employ a two-step procedure to produce uninformative samples. First, we average over all test samples of the same race. Second, we downsample to a resolution of 4 × 4 to obtain uninformative samples. In our experiments, we use 7 uninformative

Table 1: Performance metrics for each algorithm trained on UnfairFace (UFF) and FairFace (FF). Values are highlighted in bold if the Null hypothesis that the results on the UFF and FF datasets coincide is not rejected. Lower scores indicate a better performance.

	LPIPS		DS	DSSIM		$L_{\rm race}^{\cos}$		$L_{\rm race}^{0-1}$		NIQE			BLUR	
	UFF	FF	UFF	FF	UFF	FF		UFF	FF	 UFF	FF	_	UFF	FF
PULSE	0.24	0.24	0.21	0.21	0.25	0.22		0.65	0.59	3.87	3.80		-2.85	-3.07
pSp	0.18	0.16	0.18	0.16	0.16	0.11		0.56	0.37	4.76	4.73		-0.60	-0.62
fair-pSp	0.17	0.16	0.17	0.16	0.14	0.11		0.46	0.30	4.62	4.74		-0.62	-0.61
Post.Samp.	0.27	0.29	0.16	0.16	0.09	0.08		0.33	0.31	5.84	5.45		-0.89	-1.08
DDRM	0.23	0.20	0.16	0.12	0.23	0.12		0.68	0.36	6.38	6.78		-0.24	-0.24

conditions, for which we generate 100 reconstructions each, resulting in a total of 700 samples. Note, that we first average to wipe out sample-specific biases apparent in a single image. In Section 9.2, we describe the individual models and specify the hyperparameters used for training and evaluation, respectively.

5.1 Qualitative Results

We start with the qualitative evaluation of upsampling methods, which provides a first intuition and valuable insights into the strengths and weaknesses of the individual methods.

For underrepresented subpopulations, we observe that the upsampling results show faces of a considerably lighter skin tone compared to the original images if the models are trained on Unfair-Face; for example, in Figure 2(a), we show the upsampling results for samples categorized as "Black". The racial bias appears especially pronounced for models that are not designed for fair image upsampling, i.e., PULSE, pSp, and DDRM. While the diffusion-based models exhibit a less severe bias, the reconstructions appear more blurry. Notably, the blurriness appears to be reduced when reconstructing samples categorized as "White" (see Figure 2(b) in the Appendix). As a reference, in Figure 2(b), we present the results for the same methods trained on the original FairFace dataset with balanced subpopulations. Further examples provided in Figure 1 and in Section 9.4 in the Appendix demonstrate that the models fail to reconstruct particular phenotypes related to ethnicity, such as bindis and headscarves. When trained on FairFace, the reconstructions match the original samples better across all races. Unsurprisingly, all methods faithfully reconstruct faces categorized as "White" if trained on UnfairFace (see Figure 9 in the Appendix). In summary, the qualitative results highlight the potential biases in image upsampling and motivate a thorough quantitative analysis of the performance and fairness of different methods.

5.2 Upsampling Performance

Next, we present the quantitative evaluation of the upsampling performance based on the metrics described in Section 3.1. In Table 1, we compare the results across all considered methods trained on the UnfairFace and FairFace datasets, respectively. Based on these results, we perform a Wilcoxon signed-rank test [62] to test the Null hypothesis that the distributions of metrics derived from a model trained on UnfairFace and FairFace coincide.⁷ In bold, we indicate if





(a) Models trained on UnfairFace.



(b) Models trained on FairFace.

Figure 3: Comparing the race reconstruction loss $L_{\text{race}}^{0-1}(x_{\text{HR}}, \hat{x}_{\text{HR}})$ if $x_{\text{HR}} \in C_j$ for varying C_j . Lower scores indicate a better reconstruction.

the null hypothesis cannot be rejected based on a significance level of α = 0.05. In most cases, we see a significant difference between models trained on the UnfairFace and FairFace datasets. In the case of PULSE, we notice that it is less affected by the racial distribution, as it exhibits more similar results for both datasets. For DDRM, we observe a notable improvement in the race reconstruction loss when trained on FairFace. Additionally, in Section 9.5, we report the p-values of the Wilcoxon signed-rank test to assess how significant the difference in performance is for each method.

Overall, we find that StyleGAN-based models perform better in terms of image quality. We observe that almost all models are sensitive to the training data bias when it comes to race reconstruction loss. Posterior Sampling constitutes the only exception,

distribution, which prevents us from using a paired t-test to test whether the means coincide. Since $L_{\text{race}}^{0.1}$ is a binary value, we employ a Pearson's χ^2 -test in that case.

Table 2: Evaluating the fairness discrepancy $D(\mathbb{P}_{RDP} || \mathcal{U}([k]))$ and $D(\mathbb{P}_{PR} || \mathcal{U}([k]))$ for different divergences D for each algorithm trained on UnfairFace (UFF) and FairFace (FF). Lower scores indicate more fairness. The λ illustrates that the Null hypothesis $\mathbb{P}_{RDP} = \mathcal{U}([k])$ or $\mathbb{P}_{PR} = \mathcal{U}([k])$ is rejected.

	$\Delta_{ m RD}$	P- χ^2	$\Delta_{\rm RDF}$	P-Cheb	$\mathbb{P}_{\text{RDP}} = \mathcal{O}$	$\mathcal{U}([k])$	Δ_{PF}	$R - \chi^2$	$\Delta_{\rm PR}$ -	Cheb	$\mathbb{P}_{\mathrm{PR}} = \mathcal{I}$	$\mathcal{U}([k])$
	UFF	FF	UFF	FF	UFF	FF	UFF	FF	UFF	FF	UFF	FF
PULSE	0.34	0.12	0.16	0.10	X	X	0.86	0.21	0.29	0.11	×	X
pSp	0.24	0.04	0.15	0.07	×	X	0.71	0.13	0.24	0.10	×	X
fair-pSp	0.16	0.03	0.11	0.05	×	×	0.53	0.07	0.24	0.06	×	X
Post.Samp.	0.05	0.04	0.05	0.04	×	×	0.11	0.07	0.10	0.08	×	X
DDRM	0.70	0.05	0.27	0.04	×	×	1.49	0.04	0.41	0.06	×	X

which indicates that it has a smaller racial bias when trained on unbalanced data.

5.3 Fairness and Diversity

Before evaluating the fairness metrics introduced in Section 3.2, let us first break down parts of the evaluation from Table 1 and compare how the performance varies across the races.

In Figure 3, we visualize $L_{\text{race}}^{0-1}(x_{\text{HR}}, \hat{x}_{\text{HR}})$ conditioned on $x_{\text{HR}} \in$ C_i , i.e., it shows a rescaled version of \mathbb{P}_{RDP} , which we provide in Figure 19 of the Appendix. Remember, we quantified RDP (discrepancy) by the divergence between the scaled race-conditioned performance and the uniform distribution (see Equation (7)). The results in Figure 3(a) illustrate that all models trained on UnfairFace have a comparably low race reconstruction loss for images labeled as "White". For all other races, we observe high reconstruction losses, especially for PULSE and DDRM. Interestingly, we find that "Southeast Asian" and "Indian", which represent the least frequent races in UnfairFace (see Figure 8(b)), have the largest reconstruction loss. When trained on FairFace, we observe that the reconstruction losses approach uniformity. As before, we find that DDRM is most sensitive with respect to the training data. A similar plot that visualizes PPR is provided in Figure 20 in Section 9.5 of the Appendix. In Table 2, we summarize these results by evaluating the divergences proposed in Section 3.2. If the models are trained on UnfairFace, we find that Posterior Sampling achieves the highest degree of fairness. Generally, training on FairFace has a large influence on the fairness of the models; all fairness scores improve. In that setting, DDRM, Posterior Sampling, and fair-pSp obtain approximately similar scores in all metrics, while PULSE performs worst in all metrics. To test whether $\mathbb{P}_{\text{RDP}} = \mathcal{U}([k])$ and $\mathbb{P}_{\text{PR}} = \mathcal{U}([k])$, we run a Pearson's χ^2 -test with significance $\alpha = 0.05$. We find that the Null hypothesis is rejected for all models, even when they are trained on FairFace. This means that even though some methods are more fair than others, the statistical evidence suggests that $\mathbb{P}_{\text{RDP}} \neq \mathcal{U}([k])$ and $\mathbb{P}_{\text{PR}} \neq \mathcal{U}([k])$, i.e., no method can be considered fair. Once again, we want to highlight that the fairness of DDRM spikes from the worst score when trained on UnfairFace to the best when trained on FairFace.

To measure the diversity of the methods, we propose to upsample uninformative samples multiple times (see Figure 4). Qualitatively, we observe a clear bias towards "White" reconstructions in PULSE

Table 3: Evaluating the diversity discrepancy $D(\mathbb{P}_{UCPR} \parallel \mathcal{U}([k]))$ for different divergences D for each algorithm trained on UnfairFace (UFF) and FairFace (FF). Lower scores indicate more diversity. The λ illustrates that the Null hypothesis $\mathbb{P}_{UCPR} = \mathcal{U}([k])$ is rejected.

	$\Delta_{\text{UCPR-}\chi^2}$		$\Delta_{\rm UCP}$	R-Cheb	$\mathbb{P}_{\mathrm{UCPR}} = \mathcal{U}([k])$		
	UFF	FF	UFF	FF	UFF	FF	
PULSE	5.67	2.43	0.83	0.49	X	X	
Post.Samp.	2.85	0.24	0.57	0.11	×	X	
DDRM	5.20	0.08	0.80	0.09	×	X	

and DDRM.8 This bias becomes even more evident when examining PUCPR in Figure 5. When trained on UnfairFace, almost all reconstructions are labeled as "White". When trained on FairFace, Posterior Sampling, and DDRM generate samples such that the resulting racial distribution is close to uniformity for all races but "Black", which is still highly underrepresented. Surprisingly, PULSE generates no faces labeled as "White" anymore but almost exclusively generates samples labeled as "Southeast Asian" and "Latino Hispanic". Therefore, PULSE is subject to a strong racial bias even when trained on FairFace. We provide the quantitative evaluation based on $\Delta_{\text{UCPR}\,-\,\chi^2}$ and $\Delta_{\text{UCPR-Cheb}}$ in Table 3. Note, that since the generation process in pSp and fair-pSp is deterministic,⁹ we cannot generate different values given the same input. Therefore, these models do not apply to our analysis. Unsurprisingly, we see that training on FairFace benefits the diversity of the output for all methods. DDRM provides the most diverse outputs when trained on FairFace. We run a similar Pearson's χ^2 -test and find that no method provides significant diversity in the sense that $\mathbb{P}_{\text{UCPR}} = \mathcal{U}([k])$.

5.4 Summary

Our qualitative results reveal a discernible racial bias when models are trained on UnfairFace. This holds for 16×16 inputs and uninformative 4×4 inputs. The effect is most prominent in PULSE, while being least apparent in Posterior Sampling. Yet, this comes with a drawback, as Posterior Sampling tends to generate blurry images.

⁸Further reconstructions can be found in Section 9.4 in the Appendix.

 $^{^{9}}$ In Section 9.5, we provide an additional experiment on the diversity, which leverages randomly perturbed 4 × 4 inputs. This setup allows to evaluate pSp and fair-pSp with respect to their diversity.

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



(a) Models trained on UnfairFace.



(b) Models trained on FairFace.

Figure 4: Upsampling results for models trained on UnfairFace and FairFace using uninformative test samples. The real image is an average over images classified as "White".

The quantitative measurements of the performance metrics highlight that it is significantly influenced by the dataset choice, particularly pronounced in DDRM. Additionally, the fairness performance of DDRM is very volatile to the training dataset, which should be taken into consideration for practical applications. When trained on FairFace, fair-pSp, Posterior Sampling, and DDRM achieve comparable levels of fairness and diversity. But still, statistical tests reject the hypothesis that RDP, PR, and UCPR are satisfied. Finally, it is worth highlighting the advantages of fair-pSp over pSp. Stratification, contrastive learning, and regularizing with a race classifier could,





(b) Models trained on FairFace.

Figure 5: Comparing the uninformative conditional proportional representation distribution \mathbb{P}_{UCPR} of models trained on UnfairFace and FairFace. The horizontal dashed line indicates the bar height corresponding to a uniform distribution.

in principle, also be used to improve other techniques regarding their fairness.

6 LIMITATIONS AND FUTURE WORK

128 × 128-*Resolution.* While our provided framework (Section 3) is generally applicable, our evaluation of upsampling methods considers 128×128-resolution images. For comparison, PULSE [37] was evaluated on CelebA-HQ [26] images of resolution up to 1024×1024. However, we also want to emphasize that this choice is driven by the fact that large-scale datasets, such as CelebA-HQ and FFHQ, do not come with fairness-relevant labels. Therefore, we believe that the research community would highly benefit from a high-quality dataset with labels comparable to those in FairFace, which we consider as one valuable future research direction.

Fairness-relevant Labels. There is a general concern in choosing the fairness-relevant labels. For instance, the exact partitioning and its granularity can lead to different results [19]. Additionally, the boundaries between races can be intricate, given that phenotypic features might vary within races [48]. As an alternative to race labels, we experimented with ITA-based skin tone estimations [33]. However, we found skin tone estimates to be inconsistent and unreliable in FairFace (see e.g. Figure 3 in Kärkkäinen and Joo [25]). From a technical perspective, the labels should be easy and reliable to predict, since the empirical evaluation is based on predicted labels. For reference, the FairFace race classifier [25] used in our experiments has an average prediction accuracy of around 71% on FairFace. To diminish the effect of the imperfect classifier, we select only real test samples that are correctly classified. Although a favored approach would involve utilizing a more powerful label classifier.

Model Parameters. We found that upsampling models pretrained on FFHQ and CelebA do not yield good reconstructions on FairFace samples. This is due to differently shaped image crops and camera angles apparent in FairFace, which is in contrast to the clean images in CelebA. This is why we must retrain all models on FairFace and UnfairFace. Our hyperparameter choice (Section 9.2) was guided by the settings reported in the respective works, which might be suboptimal in this setting. All utilized models can be downloaded from our open-source repository.

7 CONCLUSION

This work advances toward a principled assessment of conditional generative models. Our contribution lies in the formulation of a comprehensive evaluation framework that extends beyond conventional performance criteria as it incorporates novel fairness and diversity metrics. To validate our framework, we introduce UnfairFace, a dataset derived from the original FairFace dataset, and perform an in-depth evaluation of the fairness of image upsampling methods. Our empirical analysis underscores the sensitivity of the fairness of all considered methods to the underlying training data. Furthermore, our findings reveal that, while some models exhibit a closer alignment with fairness ideals, no model achieves statistically significant fairness. This outcome emphasizes the pressing need for further research in this domain. In light of these insights, we encourage researchers and practitioners in this field to embrace our proposed framework in the evaluation of conditional generative models beyond image upsampling methods.

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil

8 ETHICAL CONSIDERATIONS

Various definitions of race and ethnicity have been developed over time and a unified definition remains elusive. In previous works, both terms are often used interchangeably [16, 21].

We adopted the terminology employed by Balestra and Fleischer [3] and the Canadian Institute for Health Information [14] to describe "ethnicity" as a community belonging to a common cultural group, and "race" as a group defined by similarities of physical phenotypes. It is important to note that contemporary scientific understanding supports the view that there is no biological basis for a categorization into distinct racial groups [15, 19, 48]. We are aware of the controversies surrounding these definitions, specifically that they are imprecise and do not capture the full heterogeneity of human societies and cultures. For instance, Schaefer [52] describes race as an interplay between physical similarities, social similarities, and self-identification as such. Emphasizing the socially constructed nature of the terms "race" and "ethnicity", we recognize their definitions and connotations as subject to variation across time and context.

In our specific example of image upsampling, the evaluation is restricted to visible phenotypes expressed in images with labels adopted from previous work [25]. Strictly speaking, the adopted categorization violates the nuanced meaning of race and ethnicity, which are multifaceted concepts that might vary over time, even for a single individual. For instance, an Asian immigrant can be of American ethnicity; however, a single image cannot express the underlying social and cultural characteristics in their full diversity, complexity, and variability. Consequently, assessing whether a reconstruction accurately represents a sample categorized as a specific race becomes a highly debatable endeavor. In fact, even the evaluation of reconstruction performance of specific phenotypes, such as a headscarf, may prove insufficient, as ethnicity can never be reduced to sole phenotypes. Thus, being aware of the limitations of the adopted categorization, in this work, we merely use it as a proxy to quantify the potential biases of upsampling methods in a relevant context to highlight the potential issues in real-world applications. We encourage future work to remain aware of these ethical considerations and advocate for developing a more nuanced evaluation methodology.

ACKNOWLEDGMENTS

We thank Ričards Marcinkevičs for encouraging us to enhance contemporary evaluation procedures, which eventually resulted in this work, and Mohamed Hebiri for his insightful comments. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2092 CASA – 390781972.

REFERENCES

- [1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna M. Wallach. 2018. A Reductions Approach to Fair Classification. In Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018 (Proceedings of Machine Learning Research, Vol. 80), Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 60–69. http://proceedings.mlr.press/v80/agarwal18a.html
- [2] T. W. Anderson and D. A. Darling, 1954. A Test of Goodness of Fit. J. Amer. Statist. Assoc. 49, 268 (1954), 765–769. http://www.jstor.org/stable/2281537
- [3] Carlotta Balestra and Lara Fleischer. 2018. Diversity statistics in the OECD. (2018). https://doi.org/https://doi.org/10.1787/89bae654-en

- [4] Mislav Balunovic, Anian Ruoss, and Martin T. Vechev. 2022. Fair Normalizing Flows. In *The Tenth International Conference on Learning Representations*, *ICLR 2022, Virtual Event, April 25-29, 2022.* https://openreview.net/forum?id= BrFIKuxrZE
- [5] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2023. Fairness and Machine Learning: Limitations and Opportunities. MIT Press.
- [6] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In Conference on Fairness, Accountability and Transparency, FAT 2018, 23-24 February 2018, New York, NY, USA (Proceedings of Machine Learning Research, Vol. 81), Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 77–91. http://proceedings.mlr.press/v81/buolamwini18a. html
- [7] Xiangyu Chen, Xintao Wang, Jiantao Zhou, Yu Qiao, and Chao Dong. 2023. Activating More Pixels in Image Super-Resolution Transformer. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. IEEE, 22367–22377. https://doi.org/10.1109/CVPR52729. 2023.02142
- [8] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. 2020. Fair Generative Modeling via Weak Supervision. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119). PMLR, 1887–1898. http://proceedings.mlr.press/v119/choi20a.html
- [9] Alexandra Chouldechova. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data* 5, 2 (2017), 153–163. https://doi.org/10.1089/BIG.2016.0047
- [10] Evgenii Chzhen, Christophe Denis, Mohamed Hebiri, Luca Oneto, and Massimiliano Pontil. 2020. Fair regression with Wasserstein barycenters. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). 7321–7331. https://proceedings.neurips.cc/paper/2020/ hash/51cdbd2611e844ecc5d80878eb770436-Abstract.html
- [11] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. 2021. Fairness in Deep Learning: A Computational Perspective. *IEEE Intell. Syst.* 36, 4 (2021), 25–34. https: //doi.org/10.1109/MIS.2020.3000681
- [12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, Shafi Goldwasser (Ed.). ACM, 214–226. https://doi.org/10.1145/2090236.2090255
- [13] Harrison Edwards and Amos J. Storkey. 2016. Censoring Representations with an Adversary. In 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings, Yoshua Bengio and Yann LeCun (Eds.). http://arxiv.org/abs/1511.05897
- [14] Canadian Institute for Health Information. 2022. Guidance on the use of standards for race-based and indigenous identity data collection and health reporting in Canada.
- [15] Joan H Fujimura, Troy Duster, and Ramya Rajagopalan. 2008. Introduction: Race, genetics, and disease: Questions of evidence, matters of consequence. *Social Studies of Science* 38, 5 (2008), 643–656. https://doi.org/10.1177/0306312708091926
- [16] Ramán Grosfoguel. 2004. Race and Ethnicity or Racialized Ethnicities?: Identities within Global Coloniality. *Ethnicities* 4, 3 (2004), 315–336. https://doi.org/10. 1177/1468796804045237
- [17] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain, Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). 3315–3323. https://proceedings. neurips.cc/paper/2016/hash/9d2682367c3935defcb1f9e247a97c0d-Abstract.html
- [18] Ali Hatamizadeh, Jiaming Song, Guilin Liu, Jan Kautz, and Arash Vahdat. 2023. DiffiT: Diffusion Vision Transformers for Image Generation. *CoRR* abs/2312.02139 (2023). https://doi.org/10.48550/arXiv.2312.02139
- [19] Junia Howell and Michael O. Emerson. 2017. So What "Should" We Use? Evaluating the Impact of Five Racial Measures on Markers of Social Inequality. *Sociology* of Race and Ethnicity 3, 1 (2017), 14–30. https://doi.org/10.1177/2332649216648465
- [20] Åhmed Imtiaz Humayun, Randall Balestriero, and Richard G. Baraniuk. 2022. MaGNET: Uniform Sampling from Deep Generative Network Manifolds Without Retraining. In The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022. https://openreview.net/forum?id= r5qumLiYwf9
- [21] Linda M. Hunt and Mary S. Megyesi. 2008. The ambiguous meanings of the racial/ethnic categories routinely used in human genetics research. *Social Science* & Medicine 66, 2 (2008), 349–361. https://doi.org/10.1016/j.socscimed.2007.08.034
- [22] Abdul Jabbar, Xi Li, and Bourahla Omar. 2022. A Survey on Generative Adversarial Networks: Variants, Applications, and Training. ACM Comput. Surv. 54, 8 (2022), 157:1–157:49. https://doi.org/10.1145/3463475
- [23] Niharika Jain, Alberto Olmo Hernandez, Sailik Sengupta, Lydia Manikonda, and Subbarao Kambhampati. 2022. Imperfect ImaGANation: Implications of GANs exacerbating biases on facial data augmentation and snapchat face lenses. Artif.

Intell. 304 (2022), 103652. https://doi.org/10.1016/J.ARTINT.2021.103652

- [24] Ajil Jalal, Sushrut Karmalkar, Jessica Hoffmann, Alex Dimakis, and Eric Price. 2021. Fairness for Image Generation with Uncertain Sensitive Attributes. In Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139), Marina Meila and Tong Zhang (Eds.). PMLR, 4721–4732. http://proceedings.mlr. press/v139/jalal21b.html
- [25] Kimmo Kärkkäinen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In IEEE Winter Conference on Applications of Computer Vision, WACV 2021, Waikoloa, HI, USA, January 3-8, 2021. IEEE, 1547–1557. https://doi.org/10.1109/WACV48630. 2021.00159
- [26] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. https://openreview.net/ forum?id=Hk992CeAb
- [27] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2020. Training Generative Adversarial Networks with Limited Data. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). 12104–12114. https://proceedings.neurips.cc/ paper/2020/hash/8d30aa96e72440759f74bd2306c1fa3d-Abstract.html
- [28] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-Free Generative Adversarial Networks. In Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 852–863. https://proceedings.neurips.cc/ paper/2021/hash/076ccd93ad68be51f23707988e934906-Abstract.html
- [29] Tero Karras, Samuli Laine, and Timo Aila. 2019. A Style-Based Generator Architecture for Generative Adversarial Networks. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019. Computer Vision Foundation / IEEE, 4401–4410. https://doi.org/10.1109/CVPR. 2019.00453
- [30] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020. Analyzing and Improving the Image Quality of StyleGAN. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. Computer Vision Foundation / IEEE, 8107-8116. https://doi.org/10.1109/CVPR42600.2020.00813
- [31] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. 2022. Denoising Diffusion Restoration Models. In Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022, Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (Eds.). 23593–23606. http://papers.nips.cc/paper_files/paper/2022/hash/ 95504595b6169131b6ed6cd72eb05616-Abstract-Conference.html
- [32] Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI 2015, Seoul, Republic of Korea, April 18-23, 2015, Bo Begole, Jinwoo Kim, Kori Inkpen, and Woontack Woo (Eds.). ACM, 3819–3828. https://doi.org/10.1145/2702123.2702520
- [33] Newton M. Kinyanjui, Timothy Odonga, Celia Cintas, Noel C. F. Codella, Rameswar Panda, Prasanna Sattigeri, and Kush R. Varshney. 2019. Estimating Skin Tone and Effects on Classification Performance in Dermatology Datasets. *CoRR* abs/1910.13268 (2019). http://arxiv.org/abs/1910.13268
- [34] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep Learning Face Attributes in the Wild. In 2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015. IEEE Computer Society, 3730–3738. https://doi.org/10.1109/ICCV.2015.425
- [35] Vongani H. Maluleke, Neerja Thakkar, Tim Brooks, Ethan Weber, Trevor Darrell, Alexei A. Efros, Angjoo Kanazawa, and Devin Guillory. 2022. Studying Bias in GANs Through the Lens of Race. In Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XIII (Lecture Notes in Computer Science, Vol. 13673), Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer, 344–360. https://doi.org/10.1007/978-3-031-19778-9_20
- [36] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. A Survey on Bias and Fairness in Machine Learning. ACM Comput. Surv. 54, 6 (2022), 115:1-115:35. https://doi.org/10.1145/3457607
- [37] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. 2020. PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. Computer Vision Foundation / IEEE, 2434-2442. https://doi.org/10.1109/CVPR42600.2020.00251

- [38] Michele Merler, Nalini K. Ratha, Rogério Schmidt Feris, and John R. Smith. 2019. Diversity in Faces. CoRR abs/1901.10436 (2019). http://arxiv.org/abs/1901.10436
- [39] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. 2013. Making a "Completely Blind" Image Quality Analyzer. IEEE Signal Process. Lett. 20, 3 (2013), 209–212. https://doi.org/10.1109/LSP.2012.2227726
- [40] Jack Morse. 2017. Google's AI has some seriously messed up opinions about homosexuality. *mashable* (2017). https://mashable.com/article/google-machinelearning-bias
- [41] Alexander Quinn Nichol and Prafulla Dhariwal. 2021. Improved Denoising Diffusion Probabilistic Models. In Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139), Marina Meila and Tong Zhang (Eds.). PMLR, 8162–8171. http://proceedings.mlr.press/v139/nichol21a.html
- [42] Luca Oneto, Michele Donini, and Massimiliano Pontil. 2020. General Fair Empirical Risk Minimization. In 2020 International Joint Conference on Neural Networks, IJCNN 2020, Glasgow, United Kingdom, July 19-24, 2020. IEEE, 1–8. https://doi.org/10.1109/IJCNN48605.2020.9206819
- [43] José Luis Pech-Pacheco, Gabriel Cristóbal, Jesús Chamorro-Martínez, and Joaquín Fernández-Valdivia. 2000. Diatom Autofocusing in Brightfield Microscopy: a Comparative Study. In 15th International Conference on Pattern Recognition, ICPR'00, Barcelona, Spain, September 3-8, 2000. IEEE Computer Society, 3318–3321. https://doi.org/10.1109/ICPR.2000.903548
- [44] Dana Pessach and Erez Shmueli. 2023. A Review on Fairness in Machine Learning. ACM Comput. Surv. 55, 3 (2023), 51:1–51:44. https://doi.org/10.1145/3494672
- [45] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *CoRR* abs/2204.06125 (2022). https://doi.org/10.48550/ARXIV.2204.06125
- [46] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. 2021. Encoding in Style: A StyleGAN Encoder for Image-to-Image Translation. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. Computer Vision Foundation / IEEE, 2287–2296. https://doi.org/10.1109/CVPR46437.2021.00232
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. IEEE, 10674–10685. https://doi.org/10. 1109/CVPR52688.2022.01042
- [48] Wendy D. Roth. 2016. The multiple dimensions of race. *Ethnic and Racial Studies* 39, 8 (2016), 1310–1338. https://doi.org/10.1080/01419870.2016.1140793
- [49] Joni Salminen, Soon-Gyo Jung, and Bernard J. Jansen. 2019. Detecting Demographic Bias in Automatically Generated Personas. In Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019, Regan L. Mandryk, Stephen A. Brewster, Mark Hancock, Geraldine Fitzpatrick, Anna L. Cox, Vassilis Kostakos, and Mark Perry (Eds.). ACM. https://doi.org/10.1145/3290607.3313034
- [50] Prasanna Sattigeri, Samuel C. Hoffman, Vijil Chenthamarakshan, and Kush R. Varshney. 2019. Fairness GAN: Generating datasets with fairness properties using a generative adversarial network. *IBM J. Res. Dev.* 63, 4/5 (2019), 3:1–3:9. https://doi.org/10.1147/JRD.2019.2945519
- [51] Axel Sauer, Katja Schwarz, and Andreas Geiger. 2022. StyleGAN-XL: Scaling StyleGAN to Large Diverse Datasets. In SIGGRAPH '22: Special Interest Group on Computer Graphics and Interactive Techniques Conference, Vancouver, BC, Canada, August 7 - 11, 2022, Munkhtsetseg Nandigjav, Niloy J. Mitra, and Aaron Hertzmann (Eds.). ACM, 49:1-49:10. https://doi.org/10.1145/3528233.3530738
- [52] Richard T Schaefer. 2008. Encyclopedia of race, ethnicity, and society. Vol. 1. Sage.
- [53] Tom Simonite. 2015. Porbing the Dark Side of Google's Ad-Targeting System. MIT Technology Review (2015). https://www.technologyreview.com/2015/07/06/ 110198/probing-the-dark-side-of-googles-ad-targeting-system/
- [54] Yang Song and Stefano Ermon. 2020. Improved Techniques for Training Score-Based Generative Models. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). 12438–12448. https://proceedings.neurips.cc/paper/2020/hash/ 92c3b916311a5517d9290576e3ea37ad-Abstract.html
- [55] Shuhan Tan, Yujun Shen, and Bolei Zhou. 2020. Improving the Fairness of Deep Generative Models without Retraining. CoRR abs/2012.04842 (2020). https: //arxiv.org/abs/2012.04842
- [56] Md. Mehrab Tanjim, Krishna Kumar Singh, Kushal Kafle, Ritwik Sinha, and Garrison W. Cottrell. 2022. Debiasing Image-to-Image Translation Models. In 33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022. BMVA Press, 182. https://bmvc2022.mpi-inf.mpg.de/182/
- [57] Christopher T. H. Teo, Milad Abdollahzadeh, and Ngai-Man Cheung. 2023. Fair Generative Models via Transfer Learning. In Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press,

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil

2429-2437. https://doi.org/10.1609/AAAI.V37I2.25339

- [58] Kevin Truong. 2020. This Image of a White Barack Obama Is AI's Racial Bias Problem In a Nutshell. Vice News (2020). https: //www.vice.com/en/article/7kpxyy/this-image-of-a-white-barack-obama-isais-racial-bias-problem-in-a-nutshell
- [59] Mei Wang, Yaobin Zhang, and Weihong Deng. 2022. Meta Balanced Network for Fair Face Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 11 (2022), 8433–8448. https://doi.org/10.1109/TPAMI.2021.3103191
- [60] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13, 4 (2004), 600–612. https://doi.org/10.1109/TIP.2003.819861
- [61] Max Welling and Yee Whye Teh. 2011. Bayesian Learning via Stochastic Gradient Langevin Dynamics. In Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011, Lise Getoor and Tobias Scheffer (Eds.). Omnipress, 681–688. https://icml.cc/2011/papers/ 398_icmlpaper.pdf
- [62] Frank Wilcoxon. 1945. Individual Comparisons by Ranking Methods. Biometrics Bulletin 1, 6 (1945), 80-83. http://www.jstor.org/stable/3001968
- [63] Depeng Xu, Shuhan Yuan, Lu Zhang, and Xintao Wu. 2018. FairGAN: Fairnessaware Generative Adversarial Networks. In *IEEE International Conference on Big Data (IEEE BigData 2018), Seattle, WA, USA, December 10-13, 2018*, Naoki Abe, Huan Liu, Calton Pu, Xiaohua Hu, Nesreen K. Ahmed, Mu Qiao, Yang Song, Donald Kossmann, Bing Liu, Kisung Lee, Jiliang Tang, Jingrui He, and Jeffrey S. Saltz (Eds.). IEEE, 570–575. https://doi.org/10.1109/BIGDATA.2018.8622525
- [64] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. 2024. Diffusion Models: A Comprehensive Survey of Methods and Applications. ACM Comput. Surv. 56, 4 (2024), 105:1–105:39. https://doi.org/10.1145/3626235
- [65] Ning Yu, Ke Li, Peng Zhou, Jitendra Malik, Larry Davis, and Mario Fritz. 2020. Inclusive GAN: Improving Data and Minority Coverage in Generative Models. In Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXII (Lecture Notes in Computer Science, Vol. 12367), Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 377–393. https://doi.org/10.1007/978-3-030-58542-6_23
- [66] Mariia Zameshina, Olivier Teytaud, Fabien Teytaud, Vlad Hosu, Nathanaël Carraz Rakotonirina, Laurent Najman, and Markus Wagner. 2022. Fairness in generative modeling: do it unsupervised!. In GECCO '22: Genetic and Evolutionary Computation Conference, Companion Volume, Boston, Massachusetts, USA, July 9 - 13, 2022, Jonathan E. Fieldsend and Markus Wagner (Eds.). ACM, 320–323. https://doi.org/10.1145/3520304.3528992
- [67] Richard S. Zemel, Yu Wu, Kevin Swersky, Toniann Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013, Vol. 28. PMLR, 325–333. http://proceedings.mlr.press/v28/zemel13.html
- [68] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. Computer Vision Foundation / IEEE Computer Society, 586–595. https://doi.org/10.1109/CVPR.2018.00068
- [69] Zhifei Zhang, Yang Song, and Hairong Qi. 2017. Age Progression/Regression by Conditional Adversarial Autoencoder. In 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. IEEE Computer Society, 4352–4360. https://doi.org/10.1109/CVPR.2017.463
- [70] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. IEEE Computer Society, 2242–2251. https://doi.org/10.1109/ICCV. 2017.244
- [71] Yongshuo Zong, Yongxin Yang, and Timothy M. Hospedales. 2023. MEDFAIR: Benchmarking Fairness for Medical Imaging. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.* https://openreview.net/pdf?id=6ve2CkeQe5S

9 APPENDIX

9.1 Further Intuitions and Recommendations for Practitioners

The proposed fairness metrics in Section 3.2 are based on the notion of RDP, which is satisfied if

 $\mathbb{P}(\hat{x} \in C_i | x \in C_i) = \mathbb{P}(\hat{x} \in C_j | x \in C_j) \quad \forall i, j \in [k], \, \forall x \in \mathcal{X} \quad (11)$

and PR, which is satisfied if

$$\mathbb{P}(\hat{x} \in C_i) = \mathbb{P}(x \in C_i) \quad \forall i \in [k]$$
(12)

for a partition $C = \{C_1, \ldots, C_k\}$ of the dataset *X*. In this section, we aim at sheding more light on these definitions and providing the reader some more intuition.

Note that RDP (11) focuses on equalizing the class error rates of the reconstructions \hat{x} irrespective the actual class of x. Intuitively speaking, the conditional generative model should not perform better at correctly reconstructing one class over another. In contrast, PR aims at retaining the true class distribution. Specifically, in the case $\mathbb{P}(x \in C_i) = \mathbb{P}(x \in C_j)$ for all $i, j \in [k]$, PR enforces a balance of the overall class distribution of the reconstructions, which has a direct relation to fairness. Therefore, PR does not relate to correct class predictions and hence captures different goals than RDP.

In certain applications, we may favor one fairness notion over another illustrated by the following examples. For simplicity, let us reconsider the specific task of image upsampling. For the sake of illustration, let us assume that we use image upsampling for lossy data decompression, that is, the low-resolution images constitutes the compressed image and the decompression is conducted my applying an upsampling algorithm.¹⁰ Since data compression typically targets retaining the semantics of the image—and therefore the predicted class—equalizing the error rates is crucial. This goal is well-captured by enforcing RDP.

In contrast, if image upsampling is used for generating a novel synthetic dataset under coarse guidance provided by low-resolution images, we may not necessarily prioritize reconstructing the exact classes. Instead, we may focus on generating an unbiased synthetic dataset in which each class is equally represented. This is captured by PR. Ultimately, we recommend practitioners to be aware of which facets of fairness are covered by each definition. It is rarely recommended to discard one definition completely. Also, while numerical scores can only tell us whether fairness is violated, figures akin to Figure 3 and Figure 20 allow tracing down the underrepresented class causing the violation. Hence, we recommend practicioners to also analyzing these figures in their fairness assessment.

Finally, we present two extreme cases, illustrating that RDP does not induce PR and vice versa.

Example 9.1. Consider the image upsampling scenario tackled in Section 5 with just three races, "White", "Black", and "Asian". Mathematically, we define

$$C_1 := \{x \in \mathcal{X} : x \text{ is "White"}\},$$

$$C_2 := \{x \in \mathcal{X} : x \text{ is "Black"}\},$$

$$C_3 := \{x \in \mathcal{X} : x \text{ is "Asian"}\}.$$

¹⁰This is just a hypothetical non-practical usecase that exchanges data storage with computing power.

Let us assume that the conditional class distributions $\mathbb{P}(\hat{x} \in C_j | x \in C_i)$ are given by Figure 6. This means that, for instance, if the low-resolution input corresponds to a real image showing a "Black" person, then the reconstruction is "White"/"Black" with probability 50%/50% (Figure 6(b)). RDP aims at balancing the striped bars in Figure 6, which is satisfied in this concrete example. However, under the assumption that $\mathbb{P}(x \in C_1) = \mathbb{P}(x \in C_2) = \mathbb{P}(x \in C_3)$, the resulting class distribution of the reconstructions \hat{x} is given by

$$\mathbb{P}(\hat{x} \in C_1) = \sum_{i=1}^{3} \mathbb{P}(\hat{x} \in C_1 | x \in C_i) \mathbb{P}(x \in C_i) = 0.5$$

whereas

$$\mathbb{P}(\hat{x} \in C_2) = \mathbb{P}(\hat{x} \in C_3) = 0.25$$

Especially, since $\mathbb{P}(\hat{x} \in C_1) > \mathbb{P}(\hat{x} \in C_j)$ for $j \in \{2, 3\}$, i.e., "White" is overrepresented, we observe that PR is violated.

Another extreme case arises when the conditional class distributions $\mathbb{P}(\hat{x} \in C_j | x \in C_i)$ are as depicted in Figure 7. "White" faces are 100% correctly classified (Figure 7(a)), whereas "Black" and "Asian" have 0% correct class reconstructions (Figure 7(b) and Figure 7(c)). Hence, RDP is violated but PR is satisfied under the assumption $\mathbb{P}(x \in C_1) = \mathbb{P}(x \in C_2) = \mathbb{P}(x \in C_3)$:

$$\mathbb{P}(\hat{x} \in C_j) = \sum_{i=1}^{3} \mathbb{P}(\hat{x} \in C_j | x \in C_i) \mathbb{P}(x \in C_i) = \mathbb{P}(x \in C_j)$$

These examples highlight that the utilized fairness definitions are no one-size-fits-all solutions. Instead, PR and RDP capture different aspects of fairness.

9.2 Image Upsampling Models

In the following, we describe the models and specify all hyperparameters used for training and evaluation, respectively.

PULSE. PULSE upsampling leverages the expressiveness of a pretrained StyleGAN. We trained a StyleGAN2 using the official repository¹¹ and set $\gamma = 0.0128$. The final FID-score is 5.89 and 5.61 if trained on UnfairFace and FairFace, respectively. To generate the reconstructions, we used the default values provided in the official repository¹² but set the number of *bad noise layers* to 11 (instead of 17) because the StyleGAN2 we use has only 12 layers (instead of 18).

pSp. We utilize the same StyleGAN2 backbone as in PULSE. The pSp encoder is trained for 300 000 steps using the default values as provided in the official repository¹³. Note that its training procedure is based on downsampling a training sample and measuring its reconstruction to the original. We observe that if the training procedure does not contain downsampling to 4×4 resolution, reconstructions of 4×4 inputs—which is the setup in the diversity study—are not meaningful. To compensate for that lack of generalization, we added 4×4 -downscaling to the training procedure.



(a) $\mathbb{P}(\hat{x}\in C_j\,|\,x\in C_1),$ that is, the low-resolution input corresponds to "White".



(b) $\mathbb{P}(\hat{x}\in C_j\,|\,x\in C_2),$ that is, the low-resolution input corresponds to "Black".



(c) $\mathbb{P}(\hat{x}\in C_j\,|\,x\in C_3),$ that is, the low-resolution input corresponds to "Asian".

Figure 6: Given a low-resolution input ("White", "Black", or "Asian"), we assume that the class distributions of the reconstructions are given by the following probability mass functions.

¹¹https://github.com/NVlabs/stylegan3

¹² https://github.com/adamian98/pulse

¹³ https://github.com/eladrich/pixel2style2pixel



(a) $\mathbb{P}(\hat{x} \in C_j \,|\, x \in C_1),$ that is, the low-resolution input corresponds to "White".



(b) $\mathbb{P}(\hat{x} \in C_j \,|\, x \in C_1),$ that is, the low-resolution input corresponds to "Black".



(c) $\mathbb{P}(\hat{x}\in C_j|\,x\in C_1),$ that is, the low-resolution input corresponds to "Asian".

Figure 7: Given a low-resolution input ("White", "Black", or "Asian"), we assume that the class distributions of the reconstructions are given by the following probability mass functions. *fair-pSp.* We use the same setting as in pSp but additionally apply the resampling and the curriculum learning scheme as described in Tanjim et al. [56]. We found that using an MLP after standardizing the latents leads to bad results. Hence, we standardize the latents to compute the contrastive loss but proceed with the unstandardized latents, i.e., we ignore the MLP originally proposed by Tanjim et al. [56]. The different behavior could be attributed to the fact that FairFace has 7 classes, which is significantly more complex than the binary-class setup considered by Tanjim et al. [56]. Other than that, we adopt the hyperparameters from the original paper. Since there is no official implementation, we reimplemented fair-pSp¹⁴.

PosteriorSampling. For training the NCSNv2¹⁵ backbone model, we select the hyperparameters according to the techniques recommended in Song and Ermon [54]. For UnfairFace and FairFace, this results in L = 1022, $\sigma_1 = 170$, T = 3, $\varepsilon = 1.86e - 6$. We trained the models for 150 000 Iterations. We generate samples from the posterior leveraging the official repository¹⁶.

DDRM. As a backbone model, we used a DDPM as suggested by Nichol and Dhariwal [41] trained using the official repository¹⁷. We set the number of diffusion steps to 1000 and the channel multiplier of the UNet stages to 1, 1, 2, 3, 4, respectively. Other than that, we stick with the suggested baseline default hyperparameters and trained for 500 000 iterations. For computing the reconstructions, we used the default values provided by the official repository¹⁸.

9.3 Racial Distribution of UnfairFace

In Figure 8, we compare the racial distribution of FairFace and UnfairFace.

9.4 Additional Qualitative Results

Figure 9 to Figure 14 show additional upsampling results of test samples categorized as the remaining six races, "White", "Indian", "Southeast Asian", "East Asian", "Middle Eastern", "Latino Hispanic", respectively. Figure 15 reconstructs the samples provided in Figure 1 when the models are trained on FairFace. We observe that even if the models are trained on FairFace, reconstructing headscarves remains a difficult task, as shown by the blurry reconstructions. Additional reconstructions of people having bindis, headscarves, and monolid eyes are provided in Figure 16, Figure fig:scarves, and Figure fig:monolid, respectively.

9.5 Additional Quantitative Results

To indicate the significance of the difference of a performance score obtained by a model trained on UnfairFace and FairFace, we provide the corresponding P-values in Table 4. For the reconstruction-based losses (i.e., LPIPS and DSSIM), we find the biggest significance for pSp and DDRM. The training data set has the most significant influence on race reconstruction performance (i.e., L_{race}^{\cos} and L_{race}^{0-1}) for DDRM as indicated by the smallest p-value. The image quality differs most significantly for DDRM in terms of the NIQE score and for Posterior Sampling in terms of the blurriness index.

¹⁴https://github.com/MikeLasz/Fair-pSp

¹⁵https://github.com/ermongroup/ncsnv2

¹⁶https://github.com/ajiljalal/code-cs-fairness

¹⁷https://github.com/openai/improved-diffusion

¹⁸ https://github.com/bahjat-kawar/ddrm



(a) Racial distribution in FairFace.



(b) Racial distribution in UnfairFace. We leave out the value for Southeast Asian (0.05%) to avoid value overcrowding.

Figure 8: Racial distribution of FairFace and UnfairFace in comparison.

In Figure 19, we visualize \mathbb{P}_{RDP} for each model. Let us highlight two points to clarify the proposed metrics based on RDP. First, this figure presents a rescaled version of Figure 3 in which the bars corresponding to a model are normalized such that they add up to one (compare with equation 6). Second, the proposed metrics are divergences that measure the discrepancy of \mathbb{P}_{RDP} to the horizontal dashed line representing the uniform distribution in Figure 19. Specifically, $\Delta_{\text{RDP}-\chi^2}$ (see equation 8) measures the scaled meansquare distance while $\Delta_{\text{RDP-Cheb}}$ (see equation 9) measures the maximum distance of the bars to the dashed line. The results align with the numbers presented in Table 2.

Figure 20 visualizes \mathbb{P}_{PR} , which—once again—highlights the influence that a lack of diversity in the training data has on the resulting racial distribution.

Lastly, we modify the experiment on diversity in a way that allows us to measure the diversity of pSp and fair-pSp. Recall that since the upsampling procedure in pSp and fair-pSp is deterministic, these algorithms cannot generate diverse reconstructions given a fixed low-resolution input. To circumvent this issue, we randomly perturbed the input by adding a Gaussian noise¹⁹ $\varepsilon \sim \mathcal{N}(0, 10)$ as illustrated in Figure 21 and Figure 22. The resulting diversity discrepancies are presented in Table 5. Compared to Table 3, the resulting numbers tend to be slightly lower, which is not surprising given the added source of randomness. More randomness in the inputs is, intuitively, promoting more randomness and diversity in the outputs. While fair-pSp leads to the most diverse reconstructions if trained on UnfairFace, DDRM is still superior if trained on the balanced FairFace dataset. The qualitative results from Figure 21 and Figure 22 align with the superior performance of DDRM if trained on FairFace. In both cases, fair-pSp creates more diverse reconstructions than pSp. Again, in all cases the Null hypothesis $\mathbb{P}_{UCPR} = \mathcal{U}([k])$ is rejected. Nevertheless, it must be emphasized that we cannot recommend the evaluation of diversity using noisy versions, yet. While the results of this experiment align with the original experiments in Section 5, it is generally unclear how to pick the noise variable ε . Specifically, we do not conduct experiments on how the performance is influenced by the particular choice of ε . Relatedly, [46] propose to replace certain components of the inferred latent code *z* by random noise to produce multiple reconstructions given a single low-resolution input. But again, we think that this procedure introduces ambiguity in the specific choice of components and the noise pattern, leading to a non-trivial evaluation. Especially since a suitable ε may vary from generative model to generative model.

 $^{^{19}}$ In practice, the color channels of the pixels are encoded by 8-bit integers, i.e., each color value is in {0, 1, ..., 255}. To obtain a perturbed image, we first add the continuous noise, then clip the noisy pixel values to [0, 255], and then map each color channel to an 8-bit representation by rounding off.



(a) Trained on UnfairFace.



Figure 9: Upsampling results for models using test samples categorized as "White".



(a) Trained on UnfairFace.



Figure 10: Upsampling results for models using test samples categorized as "Indian".

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



(a) Trained on UnfairFace.



Figure 11: Upsampling results for models using test samples categorized as "Southeast Asian".

Laszkiewicz et al.



(a) Trained on UnfairFace.



Figure 12: Upsampling results for models using test samples categorized as "East Asian".

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



(a) Trained on UnfairFace.



Figure 13: Upsampling results for models using test samples categorized as "Middle Eastern".



(a) Trained on UnfairFace.



Figure 14: Upsampling results for models using test samples categorized as "Latino Hispanic".

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



Figure 15: Upsampling results for the samples provided in Figure 1 using models trained on the FairFace dataset.

Table 4: P-values for the tests utilized in Table 1. Bold values indicate values that are above the α -level.

	LPIPS	DSSIM	$L_{\rm race}^{\cos}$	$L_{\rm race}^{0-1}$	NIQE	BLUR
PULSE	9.37e-01	3.80e-01	6.06e-10	7.22e-3	7.37e-05	7.03e-05
pSp	3.71e-190	2.76e-192	7.39e-59	2.31e-17	4.20e-02	2.51e-03
fair-pSp	1.42e-119	5.73e-145	1.74e-49	2.31e-17	7.78e-18	4.14e-07
Post.Samp.	4.60e-08	9.51e-31	5.81e-06	4.18e-1	1.35e-58	2.25e-82
DDRM	1.35e-64	1.66e-217	1.92e-159	4.19e-61	3.30e-69	5.54e-04

Table 5: Evaluating the diversity discrepancy $D(\mathbb{P}_{UCPR} || \mathcal{U}([k]))$ for different divergences D for each algorithm trained on UnfairFace (UFF) and FairFace (FF) based on noisy inputs. Lower scores indicate more diversity. The λ illustrates that the Null hypothesis $\mathbb{P}_{UCPR} = \mathcal{U}([k])$ is rejected.

	$\frac{\Delta_{\text{UCPR-}\chi^2}}{\text{UFF} \text{FF}}$		$\frac{\Delta_{\text{UCPR-Cheb}}}{\text{UFF}}$		$\frac{\mathbb{P}_{\text{UCPR}}}{\text{UFF}} = \frac{1}{2}$	$\frac{\mathcal{U}([k])}{\mathrm{FF}}$
PULSE	5.11	1.63	0.79	0.35	X	X
pSp	2.20	2.57	0.43	0.55	×	X
fair-pSp	1.21	0.73	0.33	0.18	×	X
Post.Samp.	2.55	0.28	0.54	0.14	×	X
DDRM	4.88	0.13	0.77	0.08	×	×

Laszkiewicz et al.



(a) Trained on UnfairFace.



(b) Trained on FairFace.

Figure 16: Upsampling results for models using test samples showing a bindi.

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



(a) Trained on UnfairFace.



Figure 17: Upsampling results for models using test samples showing a headscarf.



(a) Trained on UnfairFace.



(b) Trained on FairFace.

Figure 18: Upsampling results for models using test samples showing monolid eyes.

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



Figure 19: Comparing the representation demographic parity distribution \mathbb{P}_{RDP} for each model trained on UnfairFace and FairFace. The horizontal dashed line indicates the bar height corresponding to a uniform distribution.



Figure 20: Comparing the proportional representation distribution \mathbb{P}_{PR} for each model trained on UnfairFace and FairFace. The horizontal dashed line indicates the bar height corresponding to a uniform distribution.

Laszkiewicz et al.



(a) Models trained on UnfairFace.



(b) Models trained on FairFace.

Figure 21: Upsampling results for models trained on UnfairFace and FairFace using uninformative perturbed test samples. The real image is an average over images classified as "White".

FAccT '24, June 03-06, 2024, Rio de Janeiro, Brazil



(a) Models trained on UnfairFace.



(b) Models trained on FairFace.

Figure 22: Upsampling results for models trained on UnfairFace and FairFace using uninformative perturbed test samples. The real image is an average over images classified as "Black".