

Balancing Act: Distribution-Guided Debiasing in Diffusion Models

Rishubh Parihar^{*1} Abhijnya Bhat^{*1} Abhipsa Basu¹ Saswat Mallick¹
Jogendra Nath Kundu² R. Venkatesh Babu¹

¹Indian Institute of Science, Bangalore ²Meta Reality Labs

Abstract

Diffusion Models (DMs) have emerged as powerful generative models with unprecedented image generation capability. These models are widely used for data augmentation and creative applications. However, DMs reflect the biases present in the training datasets. This is especially concerning in the context of faces, where the DM prefers one demographic subgroup vs others (eg. female vs male). In this work, we present a method for debiasing DMs without relying on additional reference data or model retraining. Specifically, we propose **Distribution Guidance**, which enforces the generated images to follow the prescribed attribute distribution. To realize this, we build on the key insight that the latent features of denoising UNet hold rich demographic semantics, and the same can be leveraged to guide debiased generation. We train **Attribute Distribution Predictor (ADP)** - a small mlp that maps the latent features to the distribution of attributes. ADP is trained with pseudo labels generated from existing attribute classifiers. The proposed Distribution Guidance with ADP enables us to do fair generation. Our method reduces bias across single/multiple attributes and outperforms the baseline by a significant margin for unconditional and text-conditional diffusion models. Further, we present a downstream task of training a fair attribute classifier by augmenting the training set with our generated data.

Code is available at - [project page](#).

1. Introduction

Recent advancements in Diffusion Models (DM) [12, 37, 47] have garnered much interest among researchers in evaluating the quality of the generated content. These models are not only used to generate realistic content but also to augment real datasets [27, 50] for downstream tasks. However, existing DMs have been found to synthesize biased content with respect to multiple demographic factors like gender, race, *etc.*, which can be detrimental to society once these models are deployed in the real world [24, 32, 38].

^{*}Equal contribution

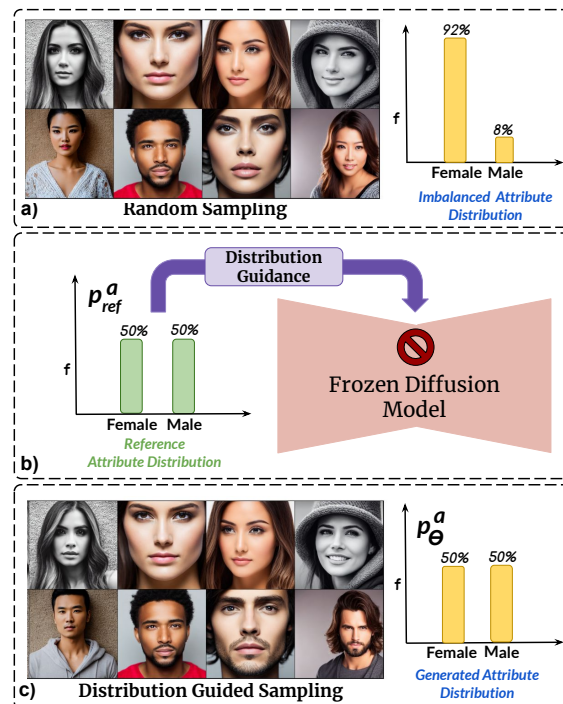


Figure 1. **a)** Random sampling from a pretrained DM [37] generates images with gender imbalance. **b)** Proposed method takes a user-defined reference attribute p_{ref}^a and performs *distribution guidance* on a pretrained DM. **c)** Sampling with distribution guidance results in *fair* generation that follow user define p_{ref}^a .

The problem is largely caused by the images used to train these models, as the outputs of these models are governed by these training datasets [25, 32]. Effects of such harmful biases have been shown by multiple recent works involving studies on DMs [24, 32, 38], GANs and other generative models [5, 13, 25]. In fact, Perera et al. [32] show that unconditional DMs—even when trained with balanced data—amplify racial biases, leading to the generation of more white-skinned faces than dark-skinned ones. Biases in the generated data are even more evident in large text-to-image DMs, e.g., models mostly tend to generate a specific gender with a given profession (like male and doctor) [26, 44, 57].

Existing works on debiasing, either require a reference dataset [5, 57] and/or allow retraining of the model [5, 55,

56]. On the contrary, our work aims to mitigate biases in both unconditional and conditional DMs, enabling *fair* generation without model retraining. We propose a practical setting where we provide a *reference attribute distribution* ($\mathbf{p}_{\text{ref}}^a$) for sensitive attribute a and query the DM to sample images following $\mathbf{p}_{\text{ref}}^a$. This framework allows for *adapting* any existing DM to a pre-defined distribution. *E.g.*, a user can provide $\mathbf{p}_{\text{ref}}^a$ as a uniform distribution for a given sensitive attribute to generate *balanced* attribute distribution. We believe defining $\mathbf{p}_{\text{ref}}^a$ provides just enough information to condition the DM for *fair* generation. This is an extremely practical setting for debiasing and is particularly important for large text-to-image DMs [35, 37, 41] where retraining or fine-tuning is computationally intensive.

One plausible approach for *fair* generation is to guide every generated sample with attribute classifiers following classifier guidance [8]. However, such a framework, though simple, is overly restrictive as it requires *presetting* and enforcing attributes for each sample individually (which we call sample guidance). Such constraints during denoising result in inferior generation quality, as discussed in Sec.4.1. Instead, we propose to *jointly* denoise a whole batch of samples and guide the process with $\mathbf{p}_{\text{ref}}^a$ (which we call distribution guidance). Specifically, we push the generated batch attribute distribution \mathbf{p}_θ^a and $\mathbf{p}_{\text{ref}}^a$ close to each other during the reverse process. Distribution guidance provides more flexibility to each sample during generation as it does not enforce a preset of attributes on a sample basis. Intuitively, distribution guidance prioritizes transforming easier samples close to the decision boundary. This results in *fair* generation without sacrificing the generation quality.

A major challenge for guidance-based conditioning is that it requires separate image classifiers for each noise scale of the diffusion process. To overcome this, we propose to perform guidance in a semantically rich feature space - *h-space* [21] of DMs. Specifically, we train an **Attribute Distribution Predictor (ADP)** that predicts attribute distribution directly from the *h-space* features. As ADP is trained on rich and discriminative *h-space* features, it - *a) is implemented as a linear layer, b) requires minimal training data, and c) is fast in training and inference.* Finally, during inference, we steer the *h-space* representation by matching the predictions from ADP to $\mathbf{p}_{\text{ref}}^a$.

We extensively evaluate our proposed method for the fair generation of single and multi-attribute cases for face generation. Additionally, we present the results of our method on Stable Diffusion [37], a large text-to-image DM. Further, as downstream application train *debiased* attribute classifiers by augmenting the training data for minority subgroups.

The major contributions of this work are the following:

1. A novel setting for debiasing existing DMs without retraining, given a reference attribute distribution.
2. Distribution guidance to condition the reverse diffusion

process on a reference attribute distribution.

3. Propose guidance in the intermediate features of diffusion network (*h-space*), which leads to data-efficient training and fast generation.

2. Related Works

Biases in Generative Models. While generative models like Generative Adversarial Networks and Diffusion Models have become the de-facto tools for image generation in recent times, studies show that they are not free of biases [14, 24, 25, 32, 38]. Perera et al. [32] show that unconditional diffusion models amplify the biases in the training data with respect to gender, race, and age. Luccioni et al. [24] identify and quantify social biases in images generated by popular text-to-image models like DALL-E 2, and Stable Diffusion v1.4 and 2. Maluleke et al. [25] study racial biases in GANs and find that GANs mimic the racial distribution in the training data.

Debiasing generative models by retraining. This line of work focuses on mitigating biases in generative models by retraining them [5, 30, 43, 49, 51, 52, 55, 56]. Some of these works [43, 52, 55] assume knowledge of the labels of the sensitive attribute and then debias the models such that there is no correlation between the decision attribute and the sensitive attribute. IMLE-GAN [56] ensures coverage of minority groups by combining GAN adversarial training with Implicit Maximum Likelihood Estimation (IMLE) [23]. Another body of works employs a balanced unlabelled reference dataset to ensure unbiased generations [5, 49, 51]. Choi et al. [5] use a density-ratio based technique to identify the bias in datasets via the reference dataset, and learn a fair model based on importance reweighting with the help of both the original biased and the reference dataset. To capture the distance between the small reference data and the generated data, Um et al. [51] use the LeCam Divergence [22]. On the other hand, Teo et al. [49] introduce a transfer learning approach to solve this problem by training the model on the biased dataset first and then adapting the model to the reference set.

Debiasing generative models without training. As training of GANs and DMs can be resource-consuming, many methods prefer fair generation of images without explicit training [13, 27, 34, 48]. MaGNET [13] aims to produce uniform sampling on the learned manifold of any generative model like GANs or VAEs, while Ramaswamy et al. [34] and Tan et al. [48] manipulate the latent space of GANs to generate balanced outputs. GANDiffFace [27], on the other hand, generates balanced synthetic for face recognition by first generating high-quality images from different demographics using GANs and then finetuning Stable Diffusion [37] using DreamBooth [39] to generate more images of such identities with different poses, expressions, etc. Multiple works attempt to mitigate biases in vision-language models and text-conditioned diffusion models as

well [3, 6, 46, 57, 58]. Chuang et al. [6] debias the text embedding using a calibrated projection matrix and shows unbiased generations without any additional training or data. However, debiasing the unconditional DMs has received less attention, which is the main focus of this work.

Guidance in Diffusion Models. One of the primary techniques to condition the diffusion model is to guide the reverse diffusion model with the gradients of additional network [8]. GLIDE [31] used CLIP [33] based guidance for open world caption to image generation. Guidance is used for image-to-image translation [54], counterfactual generation [42]. However, guidance in its original form requires retraining of guiding networks on the noisy data from scratch. Few works overcome this by learning a mapping function from the diffusion feature space to sketches for sketch-guidance [53] and a universal guidance that repurposes pretrained networks for guidance [1].

3. Method

We assume a setting where we are given a pretrained DM trained on *biased* data and a reference distribution of the sensitive attributes $\mathbf{p}_{\text{ref}}^a$. Our goal is to generate data from the DM, whose generated attribute distribution \mathbf{p}_θ^a best approximates the reference $\mathbf{p}_{\text{ref}}^a$ without retraining. The key idea is to *jointly* guide the denoising of a batch of samples such that $\mathbf{p}_{\text{ref}}^a \approx \mathbf{p}_\theta^a$. Directly computing \mathbf{p}_θ^a in a closed form is intractable. Instead, we train an **Attribute Distribution Predictor**, a linear projection that maps the intermediate batch features from *h-space* of a denoising network to an estimate of attribute distribution $\hat{\mathbf{p}}_\theta^a$.

3.1. Preliminary

Diffusion models have emerged as a powerful family of generative models trained to learn the data distribution by gradual denoising from a Gaussian distribution. Starting from a clean point \mathbf{x}_0 , and a set of scalar values $\{\alpha_t\}_{t=1}^T$, applying t steps of the forward diffusion process yields a noisy data point \mathbf{x}_t , where $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ and

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + (\sqrt{1 - \bar{\alpha}_t}) \epsilon, \epsilon \approx \mathcal{N}(0, I) \quad (1)$$

A diffusion model is learned as a neural network ϵ_θ that predicts the noise from given \mathbf{x}_t and t . The reverse process takes the form $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$, which is parameterized as a Gaussian distribution. In this work, we consider DDIM [47] sampling which first computes an estimate of the clean data point $\hat{\mathbf{x}}_0$ and then sample \mathbf{x}_{t-1} from $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \hat{\mathbf{x}}_0)$.

Classifier guidance is proposed to condition a diffusion model on class labels with the help of a pretrained classifier [8]. Specifically, a classifier $f_\phi(c | \mathbf{x}_t, t)$ is trained on noisy images to predict the class label c . The gradients of the classifier are used to guide the diffusion sampling process to generate an image of the prescribed class c . Concretely, the classifier guidance performs sampling by updating the noise prediction $\epsilon_\theta(\mathbf{x}_t, t)$ as follows:

$$\hat{\epsilon}_\theta(\mathbf{x}_t, t) = \epsilon_\theta(\mathbf{x}_t, t) - \sqrt{1 - \alpha_t} \nabla_{\mathbf{x}_t} \log f_\phi(c | \mathbf{x}_t, t) \quad (2)$$

3.2. Classifier guidance for debiasing

A promising approach is to leverage pretrained attribute classifiers to guide towards balanced generation. Assuming a reference attribute distribution $\mathbf{p}_{\text{ref}}^a$ for a binary attribute a (e.g., gender) and corresponding attribute classifier f_ϕ , we parameterize $\mathbf{p}_{\text{ref}}^a$ as a Bernoulli distribution with parameter \mathbf{r} denoting fraction of *males* samples. To generate samples following $\mathbf{p}_{\text{ref}}^a$, we can randomly select $N\mathbf{r}$ samples from the batch size of N and guide them towards *male* class using predictions from f_ϕ . Similarly, for the remaining $N(1 - \mathbf{r})$ samples, we guide them towards *female* class. In practice, such a *sample guidance* follows $\mathbf{p}_{\text{ref}}^a$ up to some extent, but results in inferior sample quality (see Fig. 3b).

Insight.1: Transforming samples close to the decision boundary is easier and results in higher quality generation.

Remark: We performed an insightful experiment for changing gender attribute (*female* to *male*) using classifier guidance. We group the samples in four quantiles based on their distance from the decision boundary of a pre-trained gender classifier in Fig. 2a). Next, we perform sample-based guidance over the samples from each quantile in Fig. 2b). The samples which are close to the decision boundary (quantiles Q1 and Q2) are easily transformed with guidance, whereas the samples away from the decision boundary (quantiles Q3 and Q4) are distorted during the guidance process. This is also quantified with the quantile-wise FID against a real set (of *male* images), where the samples from Q1/Q2 have better FID after conversion.

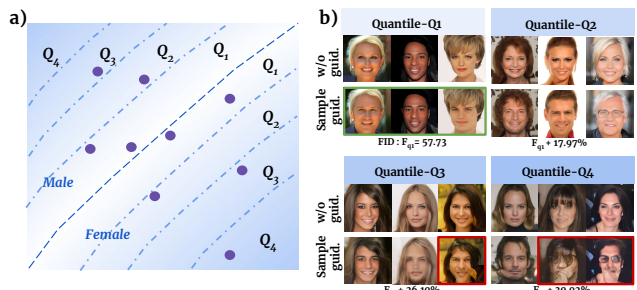


Figure 2. Distribution of samples w.r.t. decision boundary

Insight.2: Attempting to steer the generation of individual samples towards a pre-defined attribute class is overly restrictive and leads to inferior generation quality.

Remark: Sample guidance requires enforcing a preset attribute state (*male/female*) to each sample of a batch during the denoising process, which is too stringent and results in distorting the outputs. This is particularly bad in the case when samples from quantiles Q4 are selected for transformation, as shown in Fig. 2. Given an intermediate timestamp τ , the samples in the earlier stages ($t > \tau$) of reverse diffusion process are close to the noise space, resulting in the poor classifier of the corresponding stages. The guidance becomes effective only in the later stage of denoising ($t < \tau$). However, till timestep τ , some facial features are

already formed [28]. Hence, enforcing a preset attribute state at later stages ($t < \tau$) is restrictive and results in *sample collapse* (Fig. 3-b)). For e.g., for gender attribute, if a sample has formed dominant *female* features till τ it will be in Q4, and enforcing it to be *male* is overly restrictive, resulting in distorted generations.

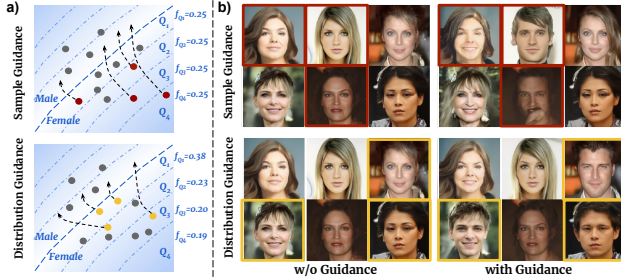


Figure 3. **Sample guidance vs Distribution guidance.** **a)** After a few steps of denoising ($t = \tau$), the generated samples have learned some discriminative features for gender. Sample guidance randomly selects **samples** from the batch uniformly from all the quantiles for conversion enforcing samples with dominant female features to also change. However, distribution guidance majorly converts the **samples** close to the decision boundary (Q1/Q2), which is *easy* to convert. **b)** Visualization of the generated samples. Samples transformed with distribution guidance are natural looking without any distortion, whereas images with sample guidance suffer from distortion or unnatural appearance.

3.3. Distribution Guidance

We propose an alternate guidance strategy termed *Distribution Guidance* for *fair* generation, which provides more flexibility to modify the attribute states during the denoising process. The key idea is to *jointly* denoise a batch of samples $\mathbf{x}_t^{[1:N]}$, with the guidance from reference attribute distribution $\mathbf{p}_{\text{ref}}^a$. To realize distribution guidance, we define a differentiable distribution prediction function g_ψ that maps the batch samples to an estimate of generated attribute distribution $\hat{\mathbf{p}}_\theta^a$. We learn the function g_ψ over an intermediate feature space - *h-space* [21] of denoising network instead of image space for efficiency (Sec.3.4). Hence, the batch estimate is given by, $\hat{\mathbf{p}}_\theta^a = g_\psi(\mathbf{h}_t^{[1:N]}, t)$, where $\mathbf{h}_t^{[1:N]}$ is the bottleneck U-Net feature representation of the batch samples $\mathbf{x}_t^{[1:N]}$. Further, we define a loss function $\mathcal{L}(\cdot)$ that measures the similarity of two distributions. During denoising we guide the batch of samples to bring $\hat{\mathbf{p}}_\theta^a$ closer to the reference, i.e., $\mathcal{L}(\hat{\mathbf{p}}_\theta^a, \mathbf{p}_{\text{ref}}^a) \approx 0$. This can be easily integrated into the reverse diffusion process as an extension of classifier guidance by modifying Eq.2.

Insight.3: *Distribution guidance provides flexibility to batch samples and transforms easier samples close to the decision boundary to match the required distribution.*

Remark: As distribution guidance does not require pre-setting for attribute states *for each sample*, it gives more flexibility during the generation as long as $\hat{\mathbf{p}}_\theta^a$ follows $\mathbf{p}_{\text{ref}}^a$.

E.g., in our running example of gender attribute, only those samples will change gender after $t = \tau$, which are close to the decision boundary (in Q1 & Q2). In contrast, in sample guidance, random samples are forced to change the attribute state (Fig. 3) and are spread equally in all four quantiles, resulting in inferior quality. This is quantified by the fraction of samples being transformed (f_{Q_i}) with each guidance in Fig. 3a), where distribution guidance majorly transforms samples from Q1 & Q2 and results in *fair* and high-quality samples (Fig. 3b)).

A major design decision for implementing distribution guidance is the selection of function g_ψ . The conventional approach is to train multiple attribute classifiers at each level of noise in the image space. Instead, we propose to leverage intermediate semantic features from the diffusion model itself and repurpose them for the task of attribute classification. This framework is very efficient as compared to image space classifiers.

3.4. Guidance in the H-space

Diffusion models, although trained for image generation, learn semantically rich representations in the internal features of denoising network ϵ_θ . These representations have been successfully used for segmentation [2] and classification [29]. Motivated by this, we ask *can we use these internal features for training guidance classifiers?* To answer this, we take features from the bottleneck layer of the diffusion U-Net, termed as the *h-space* [21] and attach a linear head for classification. The trained classifiers achieve good classification performance across multiple attributes, as shown in Fig. 5. Hence, we use the trained *h-space* classifiers to realize g_ψ . Specifically, we train a network Attribute Distribution Predictor that maps the batch *h-space* features to an estimate of attribute distribution $\hat{\mathbf{p}}_\theta^a$.



Figure 4. **Distribution guidance in the *h-space*.** For a given batch $\mathbf{x}_t^{[1:N]}$, we extract the intermediate *h-space* representation $\mathbf{h}_t^{[1:N]}$ and pass it through ADP to obtain attribute distribution $\hat{\mathbf{p}}_\theta^a$. Guidance updates $\mathbf{h}_t^{[1:N]}$ by backpropagating the derivative of loss.

Attribute Distribution Predictor (ADP) is realized via a linear attribute classifier conditioned on the diffusion time step t . Given a batch of generating samples $\mathbf{x}_t^{[1:N]}$, we extract the corresponding *h-space* features $\mathbf{h}_t^{[1:N]} = \epsilon_\theta^E(\mathbf{x}_t^{[1:N]}, t)$, where ϵ_θ^E is the encoder of U-Net network. Next, we pass the batch $\mathbf{h}_t^{[1:N]}$ to the attribute classifier and obtain a batch of softmax predictions. The softmax predictions are aggregated per class to obtain the attribute distribution estimate $\hat{\mathbf{p}}_\theta^a$. Finally, we update the intermediate *h-vectors* with the gradients of the distribution matching loss \mathcal{L} , where γ is the guidance strength parameter:

$$\begin{aligned}\tilde{\mathbf{h}}_t^{[1:N]} &= \mathbf{h}_t^{[1:N]} - \gamma * \nabla_{\mathbf{h}_t^{[1:N]}} \mathcal{L}(\hat{\mathbf{p}}_\theta^{\mathbf{a}}, \mathbf{p}_{\text{ref}}^{\mathbf{a}}) \\ \hat{\mathbf{p}}_\theta^{\mathbf{a}} &= g_\psi(\mathbf{h}_t^{[1:N]}, \mathbf{t})\end{aligned}$$

Finally, we obtain the distribution guided noise predictions $\tilde{\epsilon}$ by passing the $\mathbf{h}_t^{[1:N]}$ through the U-Net decoder ϵ_θ^D , i.e., $\tilde{\epsilon}(\mathbf{x}_t^{[1:N]}, \mathbf{t}) = \epsilon_\theta^D(\mathbf{h}_t^{[1:N]}, \mathbf{t})$. The predicted noise is then used to update batch $x_{t-1}^{[1:N]}$ using DDIM [47].

Insight.4: *H-space guidance is extremely effective and efficient as compared to image space guidance.*

Remark. As compared to conventional image space guidance, guidance in h-space has multiple advantages: **i)** it requires only a set of linear layers to implement the classifiers, **ii)** it is fast to backpropagate during guidance as compared to image models, **iii)** highly data-efficient and can be trained with only a few thousand examples due to semantically rich *h-space*. In the experiments section, we compare these properties of the *h-space* guidance.




Figure 5. Classification accuracy for linear *h-space* classifiers and ResNet-18 image space classifiers trained on 2K training examples. *h-space* classifiers are data efficient and achieve superior performance even with a linear layer

4. Experiments

In this section, we first discuss the evaluation metrics, followed by the implementation details. We take face datasets as the subject of study, as they are subject to very high demographic biases. We compare our method with other inference time debiasing approaches for both single-attribute and multi-attribute debiasing. Further, we present detailed ablations to validate design choices. Finally, we show the generalization of our debiasing approach on large text-to-image generation model Stable Diffusion [37].

4.1. Evaluation Metrics

A *fair* generative model is evaluated on two grounds: image quality and fairness. We discuss the metrics used to measure these two aspects of the generated images below.

Fairness. The primary goal of this paper is to generate images based on the reference distribution $\mathbf{p}_{\text{ref}}^{\mathbf{a}}$. We follow Choi et al. [5] to define the **Fairness Discrepancy** (FD)

metric. Given an attribute \mathbf{a} , we assume access to a high-accuracy classifier for \mathbf{a} (denoted as \mathcal{C}_a), and using the predictions from the latter, we compute the following [5, 49]:

$$\|\bar{\mathbf{p}} - \mathbb{E}_{\mathbf{x} \sim p_\theta(\mathbf{x})}(\mathbf{y})\|_2$$

where \mathbf{y} is the softmax output of the classifier $\mathcal{C}_a(\mathbf{x})$, $\bar{\mathbf{p}}$ is a uniform vector of the same dimension as \mathbf{y} , p_θ is the distribution of the generated images. The lower the FD score, the closer the distribution of the attribute values is to the uniform distribution – i.e., the generated images are fairer with respect to attribute \mathbf{a} .

Image Quality. To measure the quality of generation, we follow the standard Fréchet Inception Distance (FID) [11]. We compute this metric with the help of an attribute-wise balanced dataset sampled from the original training data.

4.2. Implementation Details

Training h-space classifiers. We start with creating a paired dataset $\mathcal{D}_{cl_f}^h$ of h-vector and attribute labels. Specifically, we take a subset \mathcal{D} of the CelebA-HQ [18] dataset and obtain attribute labels for the images using the pretrained attribute classifier \mathcal{C}_a . Next, we embed image $\mathbf{I}^i \in \mathbf{D}$ to obtain the corresponding h-space representation $\mathcal{H}^i = \{\mathbf{h}_t^i\}_{t=0}^T$ using DDIM [47] inversion. This yields a labelled dataset $\mathcal{D}_{cl_f}^h$ with pairs $(\mathcal{H}^i, \mathbf{y}^i)$, where $\mathbf{y}^i = \mathcal{C}_a(\mathbf{I}^i)$ is the predicted attribute label for image \mathbf{I}^i (e.g. *male / female*). Next, we train *h-space* attribute classifiers $\mathcal{C}_a^h(\mathbf{h}_t, t)$ as a *linear* head over \mathbf{h}_t and conditioned on time t (ignored sample index i for brevity). These obtained classifiers generate high-accuracy attribute predictions, as shown in Fig. 5. Further details about the dataset and classifiers are provided in Sec.D of supplement.

Distribution guidance. To perform distribution guidance, we realize the distribution prediction function, **ADP** with *h-space* classifiers \mathcal{C}_a^h . Specifically for a generating batch $\mathbf{x}_t^{[1:N]}$, we obtain the corresponding *h-space* representation $\mathbf{h}_t^{[1:N]}$ and obtain a set of attribute predictions $\hat{\mathbf{y}}_t^{[1:N]}$ from classifiers \mathcal{C}_a^h . Finally, we add all the softmax values for all the N images in the batch for each class to obtain the estimate of $\mathbf{p}_\theta^{\mathbf{a}}$. We use *Chi-square* distance as the loss \mathcal{L} with the reference distribution $\mathbf{p}_{\text{ref}}^{\mathbf{a}}$.

Diffusion Model Architecture. We evaluated our proposed debiasing method on two state-of-the-art pretrained DMs: a) an unconditional DM, P2 [4], trained on CelebA-HQ dataset, and b) a text conditional DM StableDiffusion v1.5 [37] trained on the LAION [45] dataset. Both these models have exceptional image generation quality; however, they have significant bias concerning the sensitive face attributes, as shown in the following subsections.

4.3. Baselines

We compare our proposed method against two techniques for guidance-based generations for DMs, one Latent-based editing method [21], and a state-of-the-art sampling-based technique for debiasing generative models, MagNet [13].



Figure 6. **Balancing generated data across *gender* and *eyeglass* attributes with various guidance strategies.** Minority group *male* and *eyeglasses* are marked. Image-space **Universal Guidance** and **Sample Guidance** results generate imbalanced and poor-quality images. For *gender*, **Sample Guidance** is able to generate only a few *males* while maintaining quality. **Latent Editing** produces collapsed images for *gender* and fails to generate *eyeglasses*. The proposed **Distribution Guidance** balances (close to the ratio of 0.50 : 0.50) both the attributes and generates high-quality images.

Sample guidance. We use classifier guidance as explained in Sec. 3.2 in the h -space using the trained h -space attribute classifiers \mathcal{C}_a^h . Such guidance requires presetting an attribute state for each generating image and pushing the trajectory to change the attribute state.

Universal guidance [1] performs guidance in the image space but uses pretrained classifiers trained on the clean image. This resolves the additional requirement of training image classifiers on noisy images. The key idea is to use the DDIM scheduler and predict the approximation of \hat{x}_0 from noise image x_t , and pass it through pretrained image classifier \mathcal{C}_A . However, this process has two shortcomings: it is slow as it backpropagates the gradients from the image classifier (Experimentally, we found it to be 7 times slower than h -space guidance) and it performs poorly at the early stage due to an inaccurate approximation of x_0 . We use two settings where we vary the number of images in the training set of the image space attribute classifier: (1) Using the same training set $|\mathcal{D}_{clf}| = 2K$ as the h -space classifier and (2) Using the entire CelebA-HQ dataset, i.e. 30k images.

Latent-based Editing [21] generates images with a specific set of attributes. Such a technique is popularly used in debiasing GANs [16, 34] because of the well-known disentangled latent spaces of GANs that allow for such edits. Recent works have shown [21] that similar semantic control is also present in the h -space of DMs and can be used for latent-based editing. We capitalize on this finding and perform latent editing in the h -space to generate images of desired attributes for *fair* generation.

MagNet [13] is an unsupervised method enabling *fair* sampling from a pretrained model. They propose a method for uniform sampling on the image manifold to generate under-represented groups equally. We generated results by MagNet sampling from a StyleGAN2 [20] model trained on the FFHQ [19] dataset from their official codebase. Notably,

Table 1. Evaluation of *balanced* generation for single attribute

Method	Gender		Race		Eyeglasses	
	FD ↓	FID ↓	FD ↓	FID ↓	FD ↓	FID ↓
Random Sampling	0.178	54.59	0.334	60.01	0.251	75.21
Universal Guidance (2k) [1]	0.193	52.10	0.377	93.42	0.189	64.55
Universal Guidance (30k) [1]	0.127	48.94	0.326	58.52	0.051	78.57
Latent Editing [21]	0.001	37.40	0.214	42.69	0.330	75.04
H-Sample Guidance (ours)	0.113	51.46	0.184	56.53	0.118	57.63
H-Distribution Guidance (ours)	0.049	50.27	0.113	52.38	0.014	51.78
StyleGAN2 - Random sampling	0.307	112.28	0.463	123.97	0.276	117.83
StyleGAN2 - Magnet [13]	0.267	91.15	0.454	97.05	0.281	106.55

as the base model for MagNet is StyleGAN2, we cannot directly compare FIDs with our DM debiased results and report random generations from StyleGAN2 as a reference.

4.4. Main Results

We first present the quantitative and qualitative results on debiasing single binary attributes. Second, we debias multiple attributes simultaneously. We finally present the case of debiasing attributes with multi-class labels.

Quantitative evaluation. We evaluate our debiasing method for the single attribute case by generating balanced generations of individual sensitive attributes - *gender*, *eyeglasses*, and *race*, in Tab. 1. As these attributes are binary, the synthesized images are expected to have a 1 : 1 ratio of the sensitive attributes (for e.g. 0.50 fraction of males and 0.50 fraction of females in case of gender). Specifically, we generate 10K images from each method per attribute and compute the metrics defined in Sec. 4.1. For most attributes, the proposed guidance method outperforms all the baselines in terms of visual quality as measured by FID and bias metric measured by FD. Although Latent editing has a better FID and FD for gender, on qualitative evaluation (as elaborated in the next section), artifacts are seen in the images (Fig. 6). Moreover, this methods fails to mitigate bias in case of multiple attributes (Tab. 3). Sample guidance achieves comparable FID for *gender*; however, higher FD indicates inferior debiasing. The tradeoff between FD and

Table 2. Evaluation of *balanced* generation for multiple attribute

Method	Gender + Race		Eyeglasses + Race		Gender + Eyeglasses	
	FD ↓	FID ↓	FD ↓	FID ↓	FD ↓	FID ↓
Random Sampling	0.256	60.68	0.292	89.14	0.214	70.97
Latent Editing [21]	0.124	64.84	0.219	90.63	0.230	74.93
Universal Guidance (2k) [11]	0.283	71.84	0.264	91.54	0.157	80.57
H-Sample Guidance (ours)	0.241	59.78	0.135	67.87	0.079	52.03
H-Distribution Guidance (ours)	0.075	49.91	0.101	57.46	0.057	47.45

FID is discussed in Sec.A.1 of supplement. This supports our thesis that the distribution guidance provides enough flexibility during generation, resulting in high-quality outputs even with a high guidance scale.

Qualitative evaluation. We present results for balancing *gender* and *eyeglasses* attributes in Fig. 6. We randomly sample 20 starting noise and use individual guidance methods for debiasing. Without guidance, the DM mostly generates *female* faces for the gender attribute. Although latent editing achieved better quantitative metrics, it produces images with artifacts and leads to collapsed results. Although Sample and Universal guidance increase the number of *males*, some images collapse. On the other hand, distribution guidance generates an almost equal number of males and females without affecting the generation quality. Moreover, all the baselines are not able to generate *eyeglasses*, whereas our method leads to highly balanced generations.

Multiple attributes. We apply our method for debiasing multiple attributes simultaneously in Tab. 3. Specifically, given two reference distributions $\mathbf{p}_{\text{ref}}^{\text{a1}}$ and $\mathbf{p}_{\text{ref}}^{\text{a2}}$, we add guidance from two pretrained attribute distribution predictors g_{ψ_1} and g_{ψ_2} . For this experiment, we define both the reference distribution as uniform for each attribute (50%-50% splits). The generated results follow the reference, resulting in a balanced generation across attributes. Further analysis is provided in Sec.B of supplement.

Multi Class attributes. We evaluate the efficacy of our approach in balancing multi-class attributes - *age* and *race* in Tab. 3. We use FFHQ [19] dataset for both attributes and obtained annotations using pretrained models as ground truth labels are unavailable. For *age* attribute, we use a pretrained VIT age classifier [36] to produce 3 classes: Young (< 20 yrs), Adult (20 – 60 yrs) and Old (> 60 yrs). For *race*, we use the Fairface race classifier [15] to obtain 4 classes: White, Black, Asian, and Indian. Our method successfully debiases multi-class attributes and beats the random and sample guidance in all cases.

4.5. Generating Imbalanced Distributions

We test our distribution guidance for generating imbalanced attribute distribution by providing skewed $\mathbf{p}_{\text{ref}}^{\text{a}}$ - **i)** 0.20 *female* and 0.80 *male* **ii)** 0.10 *white race* and 0.90 *black race*. These two settings are extremely challenging given that *male* and *black race* are minority groups in the dataset. We present qualitative results for both the settings in Fig.7, where our distribution guidance is able to generate the defined distribution with majority *males*. Note, we have binarized the race attribute as *black* and *whites* for sim-

Table 3. *Balanced* generation for multi-class attribute

Method	Age (3 classes)		Race (4 classes)	
	FD ↓	FID ↓	FD ↓	FID ↓
Random Sampling	0.256	60.68	0.292	89.14
H-Sample Guidance	0.124	64.84	0.219	90.63
H-Distribution Guidance	0.283	71.84	0.264	91.54

Table 4. Distribution guidance for imbalanced generation

Method	0.20F - 0.80M		0.10W - 0.90B	
	FD ↓	FID ↓	FD ↓	FID ↓
Random Sampling	0.478	72.26	0.734	77.63
H-Distribution Sampling	0.168	51.65	0.325	53.80

plicity, and hence in the generation, *brown race* is considered under *black* category. We report quantitative metrics in Tab. 4, where the proposed method can achieve good FD scores with the same FID. Additional experimental results are tabulated in Sec.B.1 of supplement.




Figure 7. Generating imbalanced data with two non-uniform reference distribution - **i)** 0.9B-0.1W, **ii)** 0.8M-0.2F. **Minority classes** are marked. Distribution guidance generates high-quality images and *closely follows* the skewed reference distribution.

4.6. Ablations

We provide ablation over the batch size here and that over the guidance scale, *h-space* classifier architectures and number of training examples in Sec.A of supplement. **Batch size.** As we approximate $\mathbf{p}_{\theta}^{\text{a}}$ with an estimate $\hat{\mathbf{p}}_{\theta}^{\text{a}}$ over a batch of size N , we ablate over different values of N in Tab. 5. Intuitively, using a larger batch size yields a better estimate of $\mathbf{p}_{\theta}^{\text{a}}$. We have found that the $N = 100$ works best for our experiments, balancing both FD and FID. Additionally, our model can also handle low data regime effectively.

Table 5. Ablation over batch size for gender balancing.

Batch size	2	4	8	10	25	50	75	100	125	150	200
FD ↓	0.108	0.088	0.073	0.062	0.059	0.049	0.059	0.046	0.052	0.053	0.058
FID ↓	60.12	51.86	49.80	50.54	51.51	50.86	50.98	51.64	49.91	49.81	49.74

4.7. Debiasing Text-to-Image Diffusion Models

We implemented our distribution guidance technique for debiasing a text-to-image generation model, Stable Diffusion (SD) v1.5 [37] concerning the *gender* attribute in this subsection. We provide results on other attributes and on mitigating spurious correlation on WaterBirds [40] generation in Sec.C of supplement. First, we generate a dataset \mathcal{D} , of 10K images from SD, with prompts ‘a photo of a male’ and ‘a photo of a female’ to generate a labeled dataset for training *h-space* classifiers. If the *h-space* classifier is



Figure 8. Debiasing results on stable diffusion for gender. Distribution guidance can balance the *gender* attributes in prompts involving other professions e.g. *firefighter* and *doctor*.

Table 6. Balancing *gender* on Stable Diffusion [37] model.

Method	Gender		Doctor		Firefighter	
	FD ↓	FID ↓	FD ↓	FID ↓	FD ↓	FID ↓
Random Sampling	0.317	72.37	0.355	70.11	0.235	71.86
ITI-Gen [57]	0.049	64.79	0.072	67.81	0.184	70.12
Fair Diffusion [9]	0.227	71.22	0.035	74.37	0.036	68.33
H-Sample Guidance (ours)	0.026	70.96	0.021	68.43	0.097	70.42
H-Distribution Guidance (ours)	0.024	70.69	0.015	67.36	0.093	69.41

trained on CelebA-HQ dataset, guidance is ineffective due to a significant domain shift from the SD generations. Next, we obtain the corresponding labeled dataset \mathcal{D}_{clf}^h in the *h-space* and train a *gender* classifier in the *h-space-h-space*. We used the trained classifier for distribution guidance following Sec. 3.4 to generate images with balanced *gender*.

It is observed that SD increases the gender bias when queried to generate certain professions (e.g. male and doctor) [57]. To this end, we implement our distribution guidance along the prompts ‘a photo of a doctor’ and ‘a photo of a firefighter’ to evaluate the effectiveness in this challenging setting. The qualitative results are shown in Fig. 8, and the quantitative results are reported in Tab. 6.

Baselines. We compare our method with the following methods. (1) Random sampling from SD [37]. (2) ITI-Gen [57] that learns prompt embeddings for each category of the attribute given image reference sets of each category. It then appends these prompts during generation to produce balanced images. (3) Fair Diffusion [9] that uses a lookup table to recognize the biased concept from the text input and adds scaled attribute expressions to the prompt. Note, these baselines are explicitly designed to debias text-conditioned diffusion model; however, our method can debias both conditional and unconditional diffusion models.

4.8. Class-imbalance in attribute classification

We explore an important downstream application of our proposed approach in balancing minority classes by augmenting the under-represented classes with generated data. Specifically, we train a race classifier (labels obtained from \mathcal{C}_A) on the CelebA-HQ [17] dataset. The race classifier is an ImageNet [7]-pretrained ResNet-18 [10] encoder, followed by 2 MLP layers and a classifier. We manually oversample the *Whites* and undersample the *Blacks* in the training

Table 7. Group-wise accuracies of identical classifiers trained on existing and generated (balanced) data

	Black female	White female	Black male	White male
Vanilla classifier	75.76	98.96	70.29	99.93
Balanced classifier	91.04	97.33	90.72	97.68

dataset such that the imbalanced dataset consists of 10k samples of white people and 1k black people (we keep the genders balanced within a race class). Consequently, the model performs poorly on the minority class (i.e., *Black*) due to under-representation. Next, we augment (class-balanced) the training data by generating samples whose distribution is inversely proportional to the class counts in the training set to increase images of minority classes using our distribution guidance approach. This adds 9k images of only *Black* race, and the classifier trained on this balanced data performs significantly better (Tab. 7). Even when gender is balanced in both the classes, we observe a significant disparity in the accuracies for Male and Female *Black* samples in the vanilla classifier. However, our proposed method helps reduce the accuracy gap between Black males and Black females. This shows a potential application in generating ‘class’ balanced datasets to train models for other downstream tasks, which can also mitigate bias.

5. Discussion.

Limitations. Although our method performs guidance in the *h-space*, which is efficient compared to the image space guidance, it still requires additional training of *h-space* classifiers. Another limitation is reliance on accurate attribute classifiers to obtain labels for training *h-space* classifiers.

Future works. An important future work is extending distribution guidance beyond de-biasing for controlled generation and data augmentation. In the context of debiasing DMs, extending the proposed approach without needing an attribute classifier or labeled data.

6. Conclusion

In this work, we aim to mitigate biases from pretrained diffusion models without retraining - given only a desired reference attribute distribution. We propose a novel approach leveraging *distribution guidance* that jointly guides a batch of images to follow the reference attribute distribution. The proposed method is effective and results in both high-quality and *fair* generations across multiple attributes and outperforms sample guidance strategies based on conditioning each sample individually. Extensive experiments demonstrate the effectiveness of our method in balancing both single and multiple attributes on unconditional DMs and conditional text-to-image diffusion models. We believe such a setting of debiasing without retraining is practical, especially in today’s era of large-scale generative models.

Acknowledgements. This work was supported by the Kotak IISc AIML Centre (KIAC) and Meesho. Rishubh Parihar and Abhipsa Basu are supported by PMRF fellowship.

References

- [1] Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 843–852, 2023. 3, 6, 7
- [2] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khruikov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021. 4
- [3] Hugo Berg, Siobhan Hall, Yash Bhargat, Hannah Kirk, Aleksandar Shtedritski, and Max Bain. A prompt array keeps the bias away: Debiasing vision-language models with adversarial learning. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 806–822, Online only, 2022. Association for Computational Linguistics. 3
- [4] Jooyoung Choi, Jungbeom Lee, Chaehun Shin, Sungwon Kim, Hyunwoo Kim, and Sungroh Yoon. Perception prioritized training of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11472–11481, 2022. 5
- [5] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. In *International Conference on Machine Learning*, pages 1887–1898. PMLR, 2020. 1, 2, 5
- [6] Ching-Yao Chuang, Varun Jampani, Yuanzhen Li, Antonio Torralba, and Stefanie Jegelka. Debiasing vision-language models via biased prompts. *arXiv preprint arXiv:2302.00070*, 2023. 3
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 8
- [8] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2, 3
- [9] Lukas Struppek Dominik Hintersdorf Patrick Schramowski Sasha Luccioni Kristian Kersting Felix Friedrich, Manuel Brack. Fair diffusion: Instructing text-to-image generation models on fairness. *arXiv preprint arXiv:2302.10893*, 2023. 8
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 8
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 5
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1
- [13] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. Magnet: Uniform sampling from deep generative network manifolds without retraining. In *International Conference on Learning Representations*, 2021. 1, 2, 5, 6
- [14] Niharika Jain, Alberto Olmo, Sailik Sengupta, Lydia Manikonda, and Subbarao Kambhampati. Imperfect imagination: Implications of gans exacerbating biases on facial data augmentation and snapchat selfie lenses. *arXiv preprint arXiv:2001.09528*, 2020. 2
- [15] Jungseock Joo. Fairface attribute model. <https://github.com/joojs/fairface>, 2021. 7
- [16] Cemre Efe Karakas, Alara Dirik, Eylül Yalçinkaya, and Pinar Yanardag. Fairstyle: Debiasing stylegan2 with style channel manipulations. In *European Conference on Computer Vision*, pages 570–586. Springer, 2022. 6
- [17] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 8
- [18] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 5
- [19] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 6, 7
- [20] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 6
- [21] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*, 2022. 2, 4, 5, 6, 7
- [22] Lucien Le Cam. *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012. 2

- [23] Ke Li and Jitendra Malik. Implicit maximum likelihood estimation. *arXiv preprint arXiv:1809.09087*, 2018. 2
- [24] Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. Stable bias: Evaluating societal representations in diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 1, 2
- [25] Vongani H Maluleke, Neerja Thakkar, Tim Brooks, Ethan Weber, Trevor Darrell, Alexei A Efros, Angjoo Kanazawa, and Devin Guillory. Studying bias in gans through the lens of race. In *European Conference on Computer Vision*, pages 344–360. Springer, 2022. 1, 2
- [26] Abhishek Mandal, Susan Leavy, and Suzanne Little. Measuring bias in multimodal models: Multimodal composite association score. In *International Workshop on Algorithmic Bias in Search and Recommendation*, pages 17–30. Springer, 2023. 1
- [27] Pietro Melzi, Christian Rathgeb, Ruben Tolosana, Ruben Vera-Rodriguez, Dominik Lawatsch, Florian Domin, and Maxim Schaubert. Gandifface: Controllable generation of synthetic datasets for face recognition with realistic variations. *arXiv preprint arXiv:2305.19962*, 2023. 1, 2
- [28] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022. 4
- [29] Soumik Mukhopadhyay, Matthew Gwilliam, Vatsal Agarwal, Namitha Padmanabhan, Archana Swaminathan, Srinidhi Hegde, Tianyi Zhou, and Abhinav Shrivastava. Diffusion models beat gans on image classification. *arXiv preprint arXiv:2307.08702*, 2023. 4
- [30] Junhyun Nam, Sangwoo Mo, Jaeho Lee, and Jinwoo Shin. Breaking the spurious causality of conditional generation via fairness intervention with corrective sampling. *Transactions on Machine Learning Research*, 2023. 2
- [31] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 3
- [32] Malsha V Perera and Vishal M Patel. Analyzing bias in diffusion-based face generation models. *arXiv preprint arXiv:2305.06402*, 2023. 1, 2
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastri, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [34] Vikram V Ramaswamy, Sunnie SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9301–9310, 2021. 2, 6
- [35] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [36] Nate Raw. Vit attribute model. <https://huggingface.co/nateraw/vit-age-classifier>, 2021. 7
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2, 5, 7, 8
- [38] Harrison Rosenberg, Shima Ahmed, Guruprasad V Ramesh, Ramya Korlakai Vinayak, and Kassem Fawaz. Unbiased face synthesis with diffusion models: Are we there yet? *arXiv preprint arXiv:2309.07277*, 2023. 1, 2
- [39] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. 2
- [40] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019. 7
- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 2
- [42] Pedro Sanchez and Sotirios A Tsaftaris. Diffusion causal models for counterfactual estimation. In *Conference on Causal Learning and Reasoning*, pages 647–668. PMLR, 2022. 3
- [43] Prasanna Sattigeri, Samuel C Hoffman, Vijil Chenthamarakshan, and Kush R Varshney. Fairness gan: Generating datasets with fairness properties using a

- generative adversarial network. *IBM Journal of Research and Development*, 63(4/5):3–1, 2019. 2
- [44] Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22522–22531, 2023. 1
- [45] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 5
- [46] Ashish Seth, Mayur Hemani, and Chirag Agarwal. Dear: Debiasing vision-language models with additive residuals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6820–6829, 2023. 3
- [47] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2020. 1, 3, 5
- [48] Shuhan Tan, Yujun Shen, and Bolei Zhou. Improving the fairness of deep generative models without retraining. *arXiv preprint arXiv:2012.04842*, 2020. 2
- [49] Christopher TH Teo, Milad Abdollahzadeh, and Ngai-Man Cheung. Fair generative models via transfer learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2429–2437, 2023. 2, 5
- [50] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective data augmentation with diffusion models. *arXiv preprint arXiv:2302.07944*, 2023. 1
- [51] Soobin Um and Changho Suh. A fair generative model using lecam divergence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10034–10042, 2023. 2
- [52] Boris van Breugel, Trent Kyono, Jeroen Berrevoets, and Mihaela van der Schaar. Decaf: Generating fair synthetic data using causally-aware generative networks. *Advances in Neural Information Processing Systems*, 34:22221–22233, 2021. 2
- [53] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 3
- [54] Julia Wolleb, Florentin Bieder, Robin Sandkühler, and Philippe C Cattin. Diffusion models for medical anomaly detection. In *International Conference on Medical image computing and computer-assisted intervention*, pages 35–45. Springer, 2022. 3
- [55] Depeng Xu, Shuhan Yuan, Lu Zhang, and Xintao Wu. Fairgan: Fairness-aware generative adversarial networks. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 570–575. IEEE, 2018. 1, 2
- [56] Ning Yu, Ke Li, Peng Zhou, Jitendra Malik, Larry Davis, and Mario Fritz. 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 16*, pages 377–393. Springer, 2020. 2
- [57] Cheng Zhang, Xuanbai Chen, Siqi Chai, Chen Henry Wu, Dmitry Lagun, Thabo Beeler, and Fernando De la Torre. Iti-gen: Inclusive text-to-image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3969–3980, 2023. 1, 3, 8
- [58] Beier Zhu, Yulei Niu, Saeil Lee, Minhoe Hur, and Hanwang Zhang. Debaised fine-tuning for vision-language models by prompt regularization. *arXiv preprint arXiv:2301.12429*, 2023. 3