
Fair Generative Modeling via Weak Supervision

Kristy Choi^{1*} Aditya Grover^{1*} Trisha Singh² Rui Shu¹ Stefano Ermon¹

Abstract

Real-world datasets are often biased with respect to key demographic factors such as race and gender. Due to the latent nature of the underlying factors, detecting and mitigating bias is especially challenging for unsupervised machine learning. We present a weakly supervised algorithm for overcoming dataset bias for deep generative models. Our approach requires access to an additional small, unlabeled reference dataset as the supervision signal, thus sidestepping the need for explicit labels on the underlying bias factors. Using this supplementary dataset, we detect the bias in existing datasets via a density ratio technique and learn generative models which efficiently achieve the twin goals of: 1) data efficiency by using training examples from both biased and reference datasets for learning; and 2) data generation close in distribution to the reference dataset at test time. Empirically, we demonstrate the efficacy of our approach which reduces bias w.r.t. latent factors by an average of up to 34.6% over baselines for comparable image generation using generative adversarial networks.

1. Introduction

Increasingly, many applications of machine learning (ML) involve *data generation*. Examples of such production level systems include Transformer-based models such as BERT and GPT-3 for natural language generation (Vaswani et al., 2017; Devlin et al., 2018; Radford et al., 2019; Brown et al., 2020), Wavenet for text-to-speech synthesis (Oord et al., 2017), and a large number of creative applications such as Coconet used for designing the “first AI-powered Google Doodle” (Huang et al., 2017). As these generative applications become more prevalent, it becomes increasingly important to consider questions with regards to the potential

discriminatory nature of such systems and ways to mitigate it (Podesta et al., 2014). For example, some natural language generation systems trained on internet-scale datasets have been shown to produce generations that are biased towards certain demographics (Sheng et al., 2019).

A variety of socio-technical factors contribute to the discriminatory nature of ML systems (Barocas et al., 2018). A major factor is the existence of biases in the training data itself (Torralba et al., 2011; Tommasi et al., 2017). Since data is the fuel of ML, any existing bias in the dataset can be propagated to the learned model (Barocas & Selbst, 2016). This is a particularly pressing concern for generative models which can easily amplify the bias by generating more of the biased data at test time. Further, learning a generative model is fundamentally an unsupervised learning problem and hence, the bias factors of interest are typically latent. For example, while learning a generative model of human faces, we often do not have access to attributes such as gender, race, and age. Any existing bias in the dataset with respect to these attributes are easily picked by deep generative models. See Figure 1 for an illustration.

In this work, we present a weakly-supervised approach to learning *fair* generative models in the presence of dataset bias. Our source of weak supervision is motivated by the observation that obtaining multiple unlabelled (biased) datasets is relatively cheap for many domains in the big data era. Among these data sources, we may wish to generate samples that are close in distribution to a particular target (reference) dataset.¹ As a concrete example of such a reference, organizations such as the World Bank and biotech firms (23&me, 2016; Hong, 2016) typically follow several good practices to ensure representativeness in the datasets that they collect, though such methods are unscalable to large sizes. We note that neither of our datasets need to be labeled w.r.t. the latent bias attributes and the size of the reference dataset can be much smaller than the biased dataset. Hence, the level of supervision we require is weak.

Using a reference dataset to augment a biased dataset, our goal is to learn a generative model that best approximates the

^{*}Equal contribution ¹Department of Computer Science, Stanford University ²Department of Statistics, Stanford University. Correspondence to: Kristy Choi <kechoi@cs.stanford.edu>.

¹We note that while there may not be concept of a dataset devoid of bias, carefully designed representative data collection practices may be more accurately reflected in some data sources (Gebru et al., 2018) and can be considered as *reference* datasets.



Figure 1. Samples from a baseline BigGAN that reflect the gender bias underlying the true data distribution in CelebA. All faces above the orange line (67%) are classified as female, while the rest are labeled as male (33%).

desired, reference data distribution. Simply using the reference dataset alone for learning is an option, but this may not suffice since this dataset can be too small to learn an expressive model that accurately captures the underlying reference data distribution. Our approach to learning a *fair* generative model that is robust to biases in the larger training set is based on importance reweighting. In particular, we learn a generative model which reweights the data points in the biased dataset based on the ratio of densities assigned by the biased data distribution as compared to the reference data distribution. Since we do not have access to explicit densities assigned by either of the two distributions, we estimate the weights by using a probabilistic classifier (Sugiyama et al., 2012; Mohamed & Lakshminarayanan, 2016).

We test our weakly-supervised approach on learning generative adversarial networks on the CelebA dataset (Ziwei Liu & Tang, 2015). The dataset consists of attributes such as gender and hair color, which we use for designing biased and reference data splits and subsequent evaluation. We empirically demonstrate how the reweighting approach can offset dataset bias on a wide range of settings. In particular, we obtain improvements of up to 36.6% (49.3% for $\text{bias}=0.9$ and 23.9% for $\text{bias}=0.8$) for single-attribute dataset bias and 32.5% for multi-attribute dataset bias on average over baselines in reducing the bias with respect to the latent factors for comparable sample quality.

2. Problem Setup

2.1. Background

We assume there exists a true (unknown) data distribution $p_{\text{data}} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ over a set of d observed variables $\mathbf{x} \in \mathbb{R}^d$. In generative modeling, our goal is to learn the parameters $\theta \in \Theta$ of a distribution $p_{\theta} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ over the observed variables \mathbf{x} , such that the model distribution p_{θ} is close to p_{data} . Depending on the choice of learning algorithm, different approaches have been previously considered. Broadly, these include adversarial training e.g., GANs (Goodfellow et al., 2014) and maximum likelihood estimation (MLE) e.g.,

variational autoencoders (VAE) (Kingma & Welling, 2013; Rezende et al., 2014) and normalizing flows (Dinh et al., 2014) or hybrids (Grover et al., 2018). Our bias mitigation framework is agnostic to the above training approaches.

For generality, we consider expectation-based learning objectives, where $\ell(\cdot)$ is a per-example loss that depends on both examples \mathbf{x} drawn from a dataset \mathcal{D} and the model parameters θ :

$$\mathbb{E}_{\mathbf{x} \sim p_{\text{data}}}[\ell(\mathbf{x}, \theta)] \approx \frac{1}{T} \sum_{i=1}^T \ell(\mathbf{x}_i, \theta) := \mathcal{L}(\theta; \mathcal{D}) \quad (1)$$

The above expression encompasses a broad class of MLE and adversarial objectives. For example, if $\ell(\cdot)$ denotes the negative log-likelihood assigned to the point \mathbf{x} as per p_{θ} , then we recover the MLE training objective.

2.2. Dataset Bias

The standard assumption for learning a generative model is that we have access to a sufficiently large dataset \mathcal{D}_{ref} of training examples, where each $\mathbf{x} \in \mathcal{D}_{\text{ref}}$ is assumed to be sampled independently from a reference distribution $p_{\text{data}} = p_{\text{ref}}$. In practice however, collecting large datasets that are i.i.d. w.r.t. p_{ref} is difficult due to a variety of socio-technical factors. The sample complexity for learning high dimensional distributions can even be doubly-exponential in the dimensions in many cases (Arora et al., 2018), surpassing the size of the largest available datasets.

We can partially offset this difficulty by considering data from alternate sources related to the target distribution, e.g., images scraped from the Internet. However, these additional datapoints are not expected to be i.i.d. w.r.t. p_{ref} .

We characterize this phenomena as *dataset bias*, where we assume the availability of a dataset $\mathcal{D}_{\text{bias}}$, such that the examples $\mathbf{x} \in \mathcal{D}_{\text{bias}}$ are sampled independently from a biased (unknown) distribution p_{bias} that is different from p_{ref} , but shares the same support.

2.3. Evaluation

Evaluating generative models and fairness in machine learning are both open areas of research. Our work is at the intersection of these two fields and we propose the following metrics for measuring bias mitigation for data generation.

Sample Quality: We employ sample quality metrics e.g., Frechet Inception Distance (FID) (Heusel et al., 2017), Kernel Inception Distance (KID) (Li et al., 2017), etc. These metrics match empirical expectations w.r.t. a reference data distribution p and a model distribution p_θ in a predefined feature space e.g., the prefinal layer of activations of Inception Network (Szegedy et al., 2016). A lower score indicates that the learned model can better approximate p_{data} . For the fairness context in particular, we are interested in measuring the discrepancy w.r.t. p_{ref} even if the model has been trained to use both \mathcal{D}_{ref} and $\mathcal{D}_{\text{bias}}$. We refer the reader to Supplement B.2 for more details on evaluation with FID.

Fairness: Alternatively, we can evaluate bias of generative models specifically in the context of some sensitive latent variables, say $\mathbf{u} \in \mathbb{R}^k$. For example, \mathbf{u} may correspond to the age and gender of an individual depicted via an image \mathbf{x} . We emphasize that such attributes are *unknown* during training, and used only for evaluation at test time.

If we have access to a highly accurate predictor $p(\mathbf{u}|\mathbf{x})$ for the distribution of the sensitive attributes \mathbf{u} conditioned on the observed \mathbf{x} , we can evaluate the extent of bias mitigation via the discrepancies in the expected marginal likelihoods of \mathbf{u} as per p_{ref} and p_θ .

Formally, we define the fairness discrepancy f for a generative model p_θ w.r.t. p_{ref} and sensitive attributes \mathbf{u} :

$$f(p_{\text{ref}}, p_\theta) = |\mathbb{E}_{p_{\text{ref}}}[p(\mathbf{u}|\mathbf{x})] - \mathbb{E}_{p_\theta}[p(\mathbf{u}|\mathbf{x})]|_2. \quad (2)$$

In practice, the expectations in Eq. equation 2 can be computed via Monte Carlo averaging. Again the lower is the discrepancy in the above two expectations, the better is the learned model’s ability to mitigate dataset bias w.r.t. the sensitive attributes \mathbf{u} . We refer the reader to Supplement E for more details on the fairness discrepancy metric.

3. Bias Mitigation

We assume a learning setting where we are given access to a data source $\mathcal{D}_{\text{bias}}$ in addition to a dataset of training examples \mathcal{D}_{ref} . Our goal is to capitalize on both data sources $\mathcal{D}_{\text{bias}}$ and \mathcal{D}_{ref} for learning a model p_θ that best approximates the target distribution p_{ref} .

3.1. Baselines

We begin by discussing two baseline approaches at the extreme ends of the spectrum. First, one could completely

ignore $\mathcal{D}_{\text{bias}}$ and consider learning p_θ based on \mathcal{D}_{ref} alone. Since we only consider proper losses w.r.t. p_{ref} , global optimization of the objective in Eq. equation 1 in a well-specified model family will recover the true data distribution as $|\mathcal{D}_{\text{ref}}| \rightarrow \infty$. However, since \mathcal{D}_{ref} is finite in practice, this is likely to give poor sample quality even though the fairness discrepancy would be low.

On the other extreme, we can consider learning p_θ based on the full dataset consisting of both \mathcal{D}_{ref} and $\mathcal{D}_{\text{bias}}$. This procedure will be data efficient and could lead to high sample quality, but it comes at the cost of fairness since the learned distribution will be heavily biased w.r.t. p_{ref} .

3.2. Solution 1: Conditional Modeling

Our first proposal is to learn a generative model *conditioned* on the identity of the dataset used during training. Formally, we learn a generative model $p_\theta(\mathbf{x}|y)$ where $y \in \{0, 1\}$ is a binary random variable indicating whether the model distribution was learned to approximate the data distribution corresponding to \mathcal{D}_{ref} (i.e., p_{ref}) or $\mathcal{D}_{\text{bias}}$ (i.e., p_{bias}). By sharing model parameters θ across the two values of y , we hope to leverage both data sources. At test time, conditioning on y for \mathcal{D}_{ref} should result in fair generations.

As we demonstrate in Section 4 however, this simple approach does not achieve the intended effect in practice. The likely cause is that the conditioning information is too weak for the model to infer the bias factors and effectively distinguish between the two distributions. Next, we present an alternate two-phased approach based on density ratio estimation which effectively overcomes the dataset bias in a data-efficient manner.

3.3. Solution 2: Importance Reweighting

Recall a trivial baseline in Section 3.1 which learns a generative model on the union of $\mathcal{D}_{\text{bias}}$ and \mathcal{D}_{ref} . This method is problematic because it assigns equal weight to the loss contributions from each individual datapoint in our dataset in Eq. equation 1, regardless of whether the datapoint comes from $\mathcal{D}_{\text{bias}}$ or \mathcal{D}_{ref} . For example, in situations where the dataset bias causes a minority group to be underrepresented, this objective will encourage the model to focus on the majority group such that the overall value of the loss is minimized on average with respect to a biased empirical distribution i.e., a weighted mixture of p_{bias} and p_{ref} with weights proportional to $|\mathcal{D}_{\text{bias}}|$ and $|\mathcal{D}_{\text{ref}}|$.

Our key idea is to reweight the datapoints from $\mathcal{D}_{\text{bias}}$ during training such that the model learns to downweight over-represented data points from $\mathcal{D}_{\text{bias}}$ while simultaneously upweighting the under-represented points from \mathcal{D}_{ref} . The challenge in the unsupervised context is that we do not have direct supervision on which points are over- or under-

represented and by how much. To resolve this issue, we consider importance sampling (Horvitz & Thompson, 1952). Whenever we are given data from two distributions, w.l.o.g. say p and q , and wish to evaluate a sample average w.r.t. p given samples from q , we can do so by reweighting the samples from p by the ratio of densities assigned to the sampled points by p and q . In our setting, the distributions of interest are p_{bias} and p_{ref} respectively. Hence, an importance weighted objective for learning from $\mathcal{D}_{\text{bias}}$ is:

$$\begin{aligned} \mathbb{E}_{\mathbf{x} \sim p_{\text{ref}}}[\ell(\mathbf{x}, \theta)] &= \mathbb{E}_{\mathbf{x} \sim p_{\text{bias}}} \left[\frac{p_{\text{ref}}(\mathbf{x})}{p_{\text{bias}}(\mathbf{x})} \ell(\mathbf{x}, \theta) \right] \\ &\approx \frac{1}{T} \sum_{i=1}^T w(\mathbf{x}_i) \ell(\mathbf{x}_i, \theta) := \mathcal{L}(\theta, \mathcal{D}_{\text{bias}}) \end{aligned} \quad (3)$$

where $w(\mathbf{x}_i) := \frac{p_{\text{ref}}(\mathbf{x}_i)}{p_{\text{bias}}(\mathbf{x}_i)}$ is defined to be the importance weight for $\mathbf{x}_i \sim p_{\text{bias}}$.

Estimating density ratios via binary classification. To estimate the importance weights, we use a binary classifier as described below (Sugiyama et al., 2012).

Consider a binary classification problem with classes $Y \in \{0, 1\}$ with training data generated as follows. First, we fix a prior probability for $p(Y = 1)$. Then, we repeatedly sample $y \sim p(Y)$. If $y = 1$, we independently sample a datapoint $\mathbf{x} \sim p_{\text{ref}}$, else we sample $\mathbf{x} \sim p_{\text{bias}}$. Then, as shown in Friedman et al. (2001), the ratio of densities p_{ref} and p_{bias} assigned to an arbitrary point \mathbf{x} can be recovered via a Bayes optimal (probabilistic) classifier $c^* : \mathcal{X} \rightarrow [0, 1]$:

$$w(\mathbf{x}) = \frac{p_{\text{ref}}(\mathbf{x})}{p_{\text{bias}}(\mathbf{x})} = \gamma \frac{c^*(Y = 1|\mathbf{x})}{1 - c^*(Y = 1|\mathbf{x})} \quad (5)$$

where $c(Y = 1|\mathbf{x})$ is the probability assigned by the classifier to the point \mathbf{x} belonging to class $Y = 1$. Here, $\gamma = \frac{p(Y=0)}{p(Y=1)}$ is the ratio of marginals of the labels for two classes.

In practice, we do not have direct access to either p_{bias} or p_{ref} and hence, our training data consists of points sampled from the empirical data distributions defined uniformly over \mathcal{D}_{ref} and $\mathcal{D}_{\text{bias}}$. Further, we may not be able to learn a Bayes optimal classifier and denote the importance weights estimated by the learned classifier c for a point \mathbf{x} as $\hat{w}(\mathbf{x})$.

Our overall procedure is summarized in Algorithm 1. We use deep neural networks for parameterizing the binary classifier and the generative model. Given a biased and reference dataset along with the network architectures and other standard hyperparameters (e.g., learning rate, optimizer etc.), we first learn a probabilistic binary classifier (Line 2). The learned classifier can provide importance weights for the datapoints from $\mathcal{D}_{\text{bias}}$ via estimates of the density ratios

Algorithm 1 Learning Fair Generative Models

Input: $\mathcal{D}_{\text{bias}}, \mathcal{D}_{\text{ref}}$, Classifier and Generative Model Architectures & Hyperparameters

Output: Generative Model Parameters θ

- 1: \triangleright Phase 1: Estimate importance weights
 - 2: Learn binary classifier c for distinguishing $(\mathcal{D}_{\text{bias}}, Y = 0)$ vs. $(\mathcal{D}_{\text{ref}}, Y = 1)$
 - 3: Estimate importance weight $\hat{w}(\mathbf{x}) \leftarrow \frac{c(Y=1|\mathbf{x})}{c(Y=0|\mathbf{x})}$ for all $\mathbf{x} \in \mathcal{D}_{\text{bias}}$ (using Eq. 5)
 - 4: Set importance weight $\hat{w}(\mathbf{x}) \leftarrow 1$ for all $\mathbf{x} \in \mathcal{D}_{\text{ref}}$
 - 5:
 - 6: \triangleright Phase 2: Minibatch gradient descent on θ based on weighted loss
 - 7: Initialize model parameters θ at random
 - 8: Set full dataset $\mathcal{D} \leftarrow \mathcal{D}_{\text{bias}} \cup \mathcal{D}_{\text{ref}}$
 - 9: **while** training **do**
 - 10: Sample a batch of points B from \mathcal{D} at random
 - 11: Set loss $\mathcal{L}(\theta; \mathcal{D}) \leftarrow \frac{1}{|B|} \sum_{\mathbf{x}_i \in B} \hat{w}(\mathbf{x}_i) \ell(\mathbf{x}_i, \theta)$
 - 12: Estimate gradients $\nabla_{\theta} \mathcal{L}(\theta; \mathcal{D})$ and update parameters θ based on optimizer update rule
 - 13: **end while**
 - 14: **return** θ
-

(Line 3). For the datapoints from \mathcal{D}_{ref} , we do not need to perform any reweighting and set the importance weights to 1 (Line 4). Using the combined dataset $\mathcal{D}_{\text{bias}} \cup \mathcal{D}_{\text{ref}}$, we then learn the generative model p_{θ} where the minibatch loss for every gradient update weights the contributions from each datapoint (Lines 7-13).

For a practical implementation, it is best to account for some diagnostics and best practices while executing Algorithm 1. For density ratio estimation, we test that the classifier is calibrated on a held out set. This is a necessary (but insufficient) check for the estimated density ratios to be meaningful. If the classifier is miscalibrated, we can apply standard recalibration techniques such as Platt scaling before estimating the importance weights. Furthermore, while optimizing the model using a weighted objective, there can be an increased variance across the loss contributions from each example in a minibatch due to importance weighting. We did not observe this in our experiments, but techniques such as normalization of weights within a batch can potentially help control the unintended variance introduced within a batch (Sugiyama et al., 2012).

Theoretical Analysis. The performance of Algorithm 1 critically depends on the quality of estimated density ratios, which in turn is dictated by the training of the binary classifier. We define the expected negative cross-entropy (NCE)

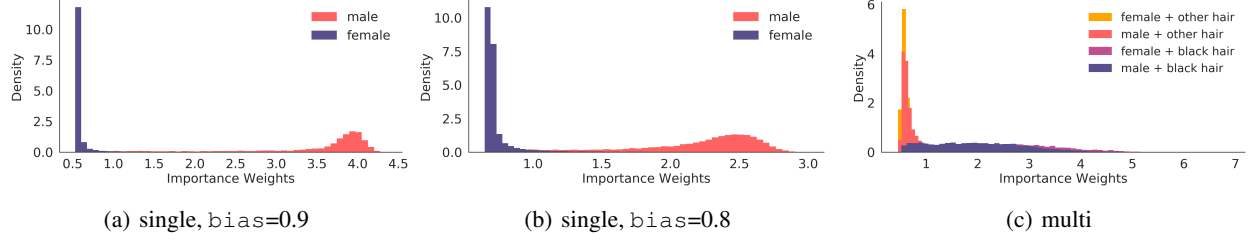


Figure 2. Distribution of importance weights for different latent subgroups. On average, The underrepresented subgroups are upweighted while the overrepresented subgroups are downweighted.

objective for a classifier c as:

$$NCE(c) := \frac{1}{\gamma + 1} \mathbb{E}_{p_{\text{ref}}(\mathbf{x})} [\log c(Y = 1|\mathbf{x})] + \frac{\gamma}{\gamma + 1} \mathbb{E}_{p_{\text{bias}}(\mathbf{x})} [\log c(Y = 0|\mathbf{x})]. \quad (6)$$

In the following result, we characterize the NCE loss for the Bayes optimal classifier.

Theorem 1. *Let \mathcal{Z} denote a set of unobserved bias variables. Suppose there exist two joint distributions $p_{\text{bias}}(\mathbf{x}, \mathbf{z})$ and $p_{\text{ref}}(\mathbf{x}, \mathbf{z})$ over $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$. Let $p_{\text{bias}}(\mathbf{x})$ and $p_{\text{bias}}(\mathbf{z})$ denote the marginals over \mathbf{x} and \mathbf{z} for the joint $p_{\text{bias}}(\mathbf{x}, \mathbf{z})$ and similar notation for the joint $p_{\text{ref}}(\mathbf{x}, \mathbf{z})$:*

$$p_{\text{bias}}(\mathbf{x}|\mathbf{z} = k) = p_{\text{ref}}(\mathbf{x}|\mathbf{z} = k) \quad \forall k \quad (7)$$

and $p_{\text{bias}}(\mathbf{x}|\mathbf{z} = k)$, $p_{\text{bias}}(\mathbf{x}|\mathbf{z} = k')$ have disjoint supports for $k \neq k'$. Then, the negative cross-entropy of the Bayes optimal classifier c^* is given as:

$$NCE(c^*) = \frac{1}{\gamma + 1} \mathbb{E}_{p_{\text{ref}}(\mathbf{z})} \left[\log \frac{1}{\gamma b(\mathbf{z}) + 1} \right] + \frac{\gamma}{\gamma + 1} \mathbb{E}_{p_{\text{bias}}(\mathbf{z})} \left[\log \frac{\gamma b(\mathbf{z})}{\gamma b(\mathbf{z}) + 1} \right]. \quad (8)$$

where $b(\mathbf{z}) = p_{\text{bias}}(\mathbf{z})/p_{\text{ref}}(\mathbf{z})$.

Proof. See Supplement A. \square

For example, as we shall see in our experiments in the following section, the inputs \mathbf{x} can correspond to face images, whereas the unobserved \mathbf{z} represents sensitive bias factors for a subgroup such as gender or ethnicity. The proportion of examples \mathbf{x} belonging a subgroup can differ across the biased and reference datasets with the relative proportions given by $b(\mathbf{z})$. Note that the above result only requires knowing these relative proportions and not the true \mathbf{z} for each \mathbf{x} . The practical implication is that under the assumptions of Theorem 1, we can check the quality of density ratios estimated by an arbitrary learned classifier c by comparing its empirical NCE with the theoretical NCE of the Bayes optimal classifier in Eq. 8 (see Section 4.1).

4. Empirical Evaluation

In this section, we are interested in investigating two broad questions empirically:

1. How well can we estimate density ratios for the proposed weak supervision setting?
2. How effective is the reweighting technique for learning fair generative models on the fairness discrepancy metric proposed in Section 2.3?

We further demonstrate the usefulness of our generated data in downstream applications such as data augmentation for learning a fair classifier in Supplement F.3.

Dataset. We consider the CelebA (Ziwei Liu & Tang, 2015) dataset, which is commonly used for benchmarking deep generative models and comprises of images of faces with 40 labeled binary attributes. We use this attribute information to construct 3 different settings for partitioning the full dataset into $\mathcal{D}_{\text{bias}}$ and \mathcal{D}_{ref} .

- **Setting 1 (single, bias=0.9):** We set \mathbf{z} to be a single bias variable corresponding to “gender” with values 0 (female) and 1 (male) and $b(\mathbf{z} = 0) = 0.9$.

Specifically, this means that \mathcal{D}_{ref} contains the same fraction of male and female images whereas $\mathcal{D}_{\text{bias}}$ contains 0.9 fraction of females and rest as males.

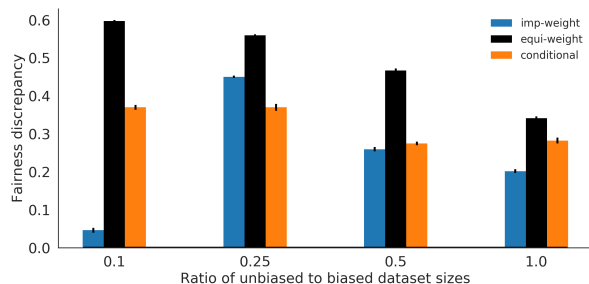
- **Setting 2 (single, bias=0.8):** We use same bias variable (gender) as Setting 1 with $b(\mathbf{z} = 0) = 0.8$.

- **Setting 3 (multi):** We set \mathbf{z} as two bias variables corresponding to “gender” and “black hair”. In total, we have 4 subgroups: females without black hair (00), females with black hair (01), males without black hair (10), and males with black hair (11). We set $b(\mathbf{z} = 00) = 0.437$, $b(\mathbf{z} = 01) = 0.063$, $b(\mathbf{z} = 10) = 0.415$, $b(\mathbf{z} = 11) = 0.085$.

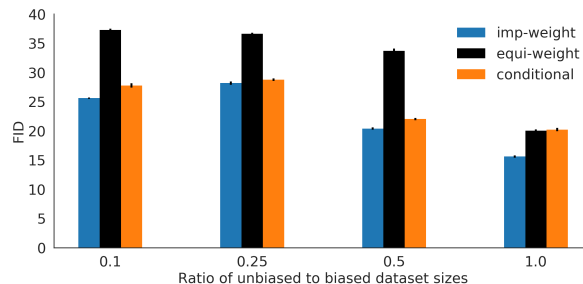
We emphasize that the attribute information is used only for designing controlled biased and reference datasets and faithful evaluation. Our algorithm does not explicitly require



(a) Samples generated via importance reweighting with subgroups separated by the orange line. For the 100 samples above, the classifier concludes 52 females and 48 males.



(b) Fairness Discrepancy



(c) FID

Figure 3. Single-Attribute Dataset Bias Mitigation for $\text{bias}=0.9$. Lower discrepancy and FID is better. Standard error in (b) and (c) over 10 independent evaluation sets of 10,000 samples each drawn from the models. We find that on average, *imp-weight* outperforms the *equi-weight* baseline by 49.3% and the *conditional* baseline by 25.0% across all reference dataset sizes for bias mitigation.

such labeled information. Additional information on constructing the dataset splits can be found in Supplement B.1.

Models. We train two classifiers for our experiments: (1) the attribute (e.g. gender) classifier which we use to assess the level of bias present in our final samples; and (2) the density ratio classifier. For both models, we use a variant of ResNet18 (He et al., 2016) on the standard train and validation splits of CelebA. For the generative model, we used a BigGAN (Brock et al., 2018) trained to minimize the hinge loss (Lim & Ye, 2017; Tran et al., 2017) objective. Additional details regarding the architectural design and hyperparameters in Supplement C.

4.1. Density Ratio Estimation via Classifier

For each of the three experiments settings, we can evaluate the quality of the estimated density ratios by comparing empirical estimates of the cross-entropy loss of the density ratio classifier with the cross-entropy loss of the Bayes optimal classifier derived in Eq. 8. We show the results in Table 1 for $\text{perc}=1.0$ where we find that the two losses are very close, suggesting that we obtain high-quality density ratio estimates that we can use for subsequently training fair generative models. In Supplement D, we show a more fine-grained analysis of the 0-1 accuracies and calibration of the learned models.

| Model | Bayes optimal | Empirical |
|---------------------------|---------------|-----------|
| single, $\text{bias}=0.9$ | 0.591 | 0.605 |
| single, $\text{bias}=0.8$ | 0.642 | 0.650 |
| multi | 0.619 | 0.654 |

Table 1. Comparison between the cross-entropy loss of the Bayes classifier and learned density ratio classifier.

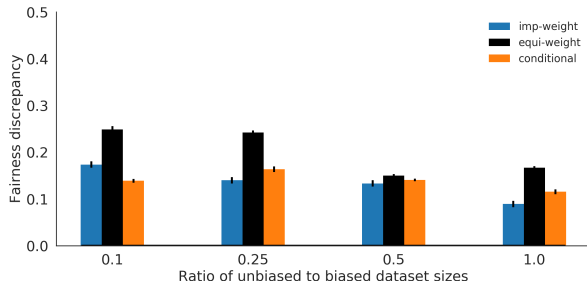
In Figure 2, we show the distribution of our importance weights for the various latent subgroups. We find that across all the considered settings, the underrepresented subgroups (e.g., males in Figure 2(a), 2(b), females with black hair in 2(c)) are upweighted on average (mean density ratio estimate > 1), while the overrepresented subgroups are downweighted on average (mean density ratio estimate < 1). Also, as expected, the density ratio estimates are closer to 1 when the bias is low (see Figure 2(a) v.s. 2(b)).

4.2. Fair Data Generation

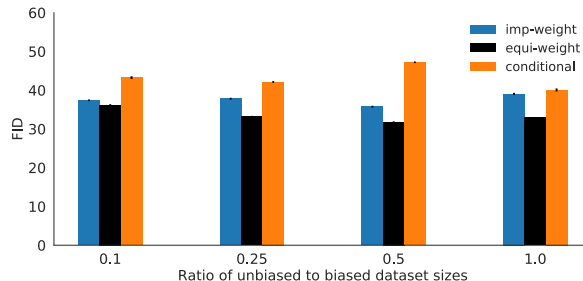
We compare our importance weighted approach against three baselines: (1) *equi-weight*: a BigGAN trained on the full dataset $\mathcal{D}_{\text{ref}} \cup \mathcal{D}_{\text{bias}}$ that weighs every point equally; (2) *reference-only*: a BigGAN trained on the reference dataset \mathcal{D}_{ref} ; and (3) *conditional*: a conditional



(a) Samples generated via importance reweighting. For the 100 samples above, the classifier concludes 37 females and 20 males without black hair, 22 females and 21 males with black hair.



(b) Fairness Discrepancy



(c) FID

Figure 4. Multi-Attribute Dataset Bias Mitigation. Standard error in (b) and (c) over 10 independent evaluation sets of 10,000 samples each drawn from the models. Lower discrepancy and FID is better. We find that on average, `imp-weight` outperforms the `equi-weight` baseline by 32.5% and the `conditional` baseline by 4.4% across all reference dataset sizes for bias mitigation.

BigGAN where the conditioning label indicates whether a data point \mathbf{x} is from $\mathcal{D}_{\text{ref}}(y = 1)$ or $\mathcal{D}_{\text{bias}}(y = 0)$. In all our experiments, the `reference-only` variant which only uses the reference dataset \mathcal{D}_{ref} for learning however failed to give any recognizable samples. For a clean presentation of the results due to other methods, we hence ignore this baseline in the results below and defer the reader to the supplementary material for further results.

We also vary the size of the balanced dataset \mathcal{D}_{ref} relative to the unbalanced dataset size $|\mathcal{D}_{\text{bias}}|$: `perc` = {0.1, 0.25, 0.5, 1.0}. Here, `perc` = 0.1 denotes $|\mathcal{D}_{\text{ref}}| = 10\%$ of $|\mathcal{D}_{\text{bias}}|$ and `perc` = 1.0 denotes $|\mathcal{D}_{\text{ref}}| = |\mathcal{D}_{\text{bias}}|$.

4.2.1. SINGLE ATTRIBUTE SPLITS

We train our attribute (gender) classifier for evaluation on the entire CelebA training set, and achieve a level of 98% accuracy on the held-out set. For each experimental setting, we evaluate bias mitigation based on the fairness discrepancy metric (Eq. 2) and also report sample quality based on FID (Heusel et al., 2017).

For the `bias` = 0.9 split, we show the samples generated via `imp-weight` in Figure 3a and the resulting fairness discrepancies in Figure 3b. Our framework generates samples that are slightly lower quality than `equi-weight`

baseline samples shown in Figure 1, but is able to produce almost identical proportion of samples across the two genders. Similar observations hold for `bias` = 0.8, as shown in Figure 8 in the supplement. We refer the reader to Supplement F.4 for corresponding results and analysis, as well as for additional results on the Shap3D dataset (Burgess & Kim, 2018).

4.2.2. MULTI-ATTRIBUTE SPLIT

We conduct a similar experiment with a multi-attribute split based on gender and the presence of black hair. The attribute classifier for the purpose of evaluation is now trained with a 4-way classification task instead of 2, and achieves an accuracy of roughly 88% on the test set.

Our model produces samples as shown in Figure 4a with the discrepancy metrics shown in Figures 4b, c respectively. Even in this challenging setup involving two latent bias factors, we find that the importance weighted approach again outperforms the baselines in almost all cases in mitigating bias in the generated data while admitting only a slight deterioration in image quality overall.

5. Related Work

Fairness & generative modeling. There is a rich body of work in fair ML, which focus on different notions of fairness (e.g. demographic parity, equality of odds and opportunity) and study methods by which models can perform tasks such as classification in a non-discriminatory way (Barocas et al., 2018; Dwork et al., 2012; Heidari et al., 2018; du Pin Calmon et al., 2018). Our focus is in the context of fair generative modeling. The vast majority of related work in this area is centered around fair and/or privacy preserving *representation learning*, which exploit tools from adversarial learning and information theory among others (Zemel et al., 2013; Edwards & Storkey, 2015; Louizos et al., 2015; Beutel et al., 2017; Song et al., 2018; Adel et al., 2019). A unifying principle among these methods is such that a discriminator is trained to perform poorly in predicting an outcome based on a protected attribute. Ryu et al. (2017) considers transfer learning of race and gender identities as a form of weak supervision for predicting other attributes on datasets of faces. While the end goal for the above works is classification, our focus is on data generation in the presence of dataset bias and we do not require explicit supervision for the protected attributes.

The most relevant prior works in data generation are FairGAN (Xu et al., 2018) and FairnessGAN (Sattigeri et al., 2019). The goal of both methods is to generate *fair* datapoints and their labels as a preprocessing technique. This allows for learning a useful downstream classifier and obscures information about protected attributes. Again, these works are not directly comparable to ours as we do not assume explicit supervision regarding the protected attributes during training, and our goal is fair generation given unlabelled biased datasets where the bias factors are latent. Another relevant work is DB-VAE (Amini et al., 2019), which utilizes a VAE to learn the latent structure of sensitive attributes, and in turn employs importance weighting based on this structure to mitigate bias in downstream classifiers. Contrary to our work, these importance weights are used to directly sample (rare) data points with higher frequencies with the goal of training a classifier (e.g. as in a facial detection system), as opposed to fair generation.

Importance reweighting. Reweighting datapoints is a common algorithmic technique for problems such as dataset bias and class imbalance (Byrd & Lipton, 2018). It has often been used in the context of fair classification (Calders et al., 2009), for example, (Kamiran & Calders, 2012) details *reweighting* as a way to remove discrimination without relabeling instances. For reinforcement learning, Doroudi et al. (2017) used an importance sampling approach for selecting fair policies. There is also a body of work on fair clustering (Chierichetti et al., 2017; Backurs et al., 2019; Bera et al., 2019; Schmidt et al., 2018) which ensure that the

clustering assignments are balanced with respect to some sensitive attribute.

Density ratio estimation using classifiers. The use of classifiers for estimating density ratios has a rich history of prior works across ML (Sugiyama et al., 2012). For deep generative modeling, density ratios estimated by classifiers have been used for expanding the class of various learning objectives (Nowozin et al., 2016; Mohamed & Lakshminarayanan, 2016; Grover & Ermon, 2018), evaluation metrics based on two-sample tests (Gretton et al., 2007; Bowman et al., 2015; Lopez-Paz & Oquab, 2016; Danihelka et al., 2017; Rosca et al., 2017; Im et al., 2018; Gulrajani et al., 2018), or improved Monte Carlo inference via these models (Grover et al., 2019; Azadi et al., 2018; Turner et al., 2018; Tao et al., 2018). Grover et al. (2019) use importance reweighting for mitigating model bias between p_{data} and p_{θ} .

Closest related is the proposal of Diesendruck et al. (2018) to use importance reweighting for learning generative models where training and test distributions differ, but explicit importance weights are provided for at least a subset of the training examples. We consider a more realistic, weakly-supervised setting where we estimate the importance weights using a small reference dataset. Finally, another related line of work in domain translation via generation considers learning via multiple datasets (Zhu et al., 2017; Choi & Jang, 2018; Grover et al., 2020) and it would be interesting to consider issues due to dataset bias in those settings in future work.

6. Discussion

Our work presents an initial foray into the field of fair image generation with weak supervision, and we stress the need for caution in using our techniques and interpreting the empirical findings. For scaling our evaluation, we proposed metrics that relied on a pretrained attribute classifier for inferring the bias in the generated data samples. The classifiers we considered are highly accurate on all subgroups, but can have blind spots especially when evaluated on generated data. For future work, we would like to investigate conducting human evaluations to mitigate such issues during evaluation (Grgic-Hlaca et al., 2018).

As another case in point, our work calls for rethinking sample quality metrics for generative models in the presence of dataset bias (Mitchell et al., 2019). On one hand, our approach *increases* the diversity of generated samples in the sense that the different subgroups are more balanced; at the same time, however, variation across other image features *decreases* because the newly generated underrepresented samples are learned from a smaller dataset of underrepresented subgroups. Moreover, standard metrics such as FID

even when evaluated with respect to a reference dataset, could exhibit a relative preference for models trained on larger datasets with little or no bias correction to avoid even slight compromises on perceptual sample quality.

More broadly, this work is yet another reminder that we must be mindful of the decisions made at each stage in the development and deployment of ML systems (Abebe et al., 2020). Factors such as the dataset used for training (Gebru et al., 2018; Sheng et al., 2019; Jo & Gebru, 2020) or algorithmic decisions such as the loss function or evaluation metric (Hardt et al., 2016; Buolamwini & Gebru, 2018; Kim et al., 2018; Liu et al., 2018; Hashimoto et al., 2018), among others, may have undesirable consequences. Becoming more aware of these downstream impacts will help to mitigate the potentially discriminatory nature of our present-day systems (Kaesler-Chen et al., 2020).

7. Conclusion

We considered the task of fair data generation given access to a (potentially small) reference dataset and a large biased dataset. For data-efficient learning, we proposed an importance weighted objective that corrects bias by reweighting the biased datapoints. These weights are estimated by a binary classifier. Empirically, we showed that our technique outperforms baselines by up to 34.6% on average in reducing dataset bias on CelebA without incurring a significant reduction in sample quality. We provide reference implementations in PyTorch (Paszke et al., 2017), and the codebase for this work is open-sourced at <https://github.com/ermongroup/fairgen>.

It would be interesting to explore whether even weaker forms of supervision would be possible for this task, e.g., when the biased dataset has a somewhat disjoint but *related* support from the small, reference dataset – this would be highly reflective of the diverse data sources used for training many current and upcoming large-scale ML systems (Ratner et al., 2017).

Acknowledgements

We are thankful to Hima Lakkaraju, Daniel Levy, Mike Wu, Chris Cundy, and Jiaming Song for insightful discussions and feedback. KC is supported by the NSF GRFP, Qualcomm Innovation Fellowship, and Stanford Graduate Fellowship, and AG is supported by the MSR Ph.D. fellowship, Stanford Data Science scholarship, and Lieberman fellowship. This research was funded by NSF (#1651565, #1522054, #1733686), ONR (N00014-19-1-2145), AFOSR (FA9550-19-1-0024), and Amazon AWS.

References

- 23&me. The real issue: Diversity in genetics research. Retrieved from <https://blog.23andme.com/ancestry/the-real-issue-diversity-in-genetics-research/>, 2016.
- Abebe, R., Barocas, S., Kleinberg, J., Levy, K., Raghavan, M., and Robinson, D. G. Roles for computing in social change. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 252–260, 2020.
- Adel, T., Valera, I., Ghahramani, Z., and Weller, A. One-network adversarial fairness. 2019.
- Amini, A., Soleimany, A. P., Schwarting, W., Bhatia, S. N., and Rus, D. Uncovering and mitigating algorithmic bias through learned latent structure. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 289–295, 2019.
- Arora, S., Risteski, A., and Zhang, Y. Do gans learn the distribution? some theory and empirics. In *International Conference on Learning Representations*, 2018.
- Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. Discriminator rejection sampling. *arXiv preprint arXiv:1810.06758*, 2018.
- Backurs, A., Indyk, P., Onak, K., Schieber, B., Vakilian, A., and Wagner, T. Scalable fair clustering. *arXiv preprint arXiv:1902.03519*, 2019.
- Barocas, S. and Selbst, A. D. Big data’s disparate impact. *Calif. L. Rev.*, 104:671, 2016.
- Barocas, S., Hardt, M., and Narayanan, A. *Fairness and Machine Learning*. fairmlbook.org, 2018. <http://www.fairmlbook.org>.
- Bera, S. K., Chakrabarty, D., and Negahbani, M. Fair algorithms for clustering. *arXiv preprint arXiv:1901.02393*, 2019.
- Beutel, A., Chen, J., Zhao, Z., and Chi, E. H. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*, 2017.
- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*, 2015.
- Brock, A., Donahue, J., and Simonyan, K. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. 2020.
- Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pp. 77–91, 2018.
- Burgess, C. and Kim, H. 3d shapes dataset. <https://github.com/deepmind/3dshapes-dataset/>, 2018.
- Byrd, J. and Lipton, Z. C. What is the effect of importance weighting in deep learning? *arXiv preprint arXiv:1812.03372*, 2018.
- Calders, T., Kamiran, F., and Pechenizkiy, M. Building classifiers with independency constraints. In *2009 IEEE International Conference on Data Mining Workshops*, pp. 13–18. IEEE, 2009.
- Chierichetti, F., Kumar, R., Lattanzi, S., and Vassilvitskii, S. Fair clustering through fairlets. In *Advances in Neural Information Processing Systems*, pp. 5029–5037, 2017.
- Choi, H. and Jang, E. Generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018.
- Danihelka, I., Lakshminarayanan, B., Uria, B., Wierstra, D., and Dayan, P. Comparison of maximum likelihood and gan-based training of real nvps. *arXiv preprint arXiv:1705.05263*, 2017.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Diesendruck, M., Elenberg, E. R., Sen, R., Cole, G. W., Shakkottai, S., and Williamson, S. A. Importance weighted generative networks. *arXiv preprint arXiv:1806.02512*, 2018.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Doroudi, S., Thomas, P. S., and Brunskill, E. Importance sampling for fair policy selection. *Grantee Submission*, 2017.
- du Pin Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N., and Varshney, K. R. Data pre-processing for discrimination prevention: Information-theoretic optimization and analysis. *IEEE Journal of Selected Topics in Signal Processing*, 12(5):1106–1119, 2018.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214–226. ACM, 2012.
- Edwards, H. and Storkey, A. Censoring representations with an adversary. *arXiv preprint arXiv:1511.05897*, 2015.
- Friedman, J., Hastie, T., and Tibshirani, R. *The elements of statistical learning*, volume 1. Springer series in statistics New York, NY, USA:, 2001.
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumeé III, H., and Crawford, K. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Gretton, A., Borgwardt, K. M., Rasch, M., Schölkopf, B., and Smola, A. J. A kernel method for the two-sample problem. In *Advances in neural information processing systems*, pp. 513–520, 2007.
- Grgic-Hlaca, N., Redmiles, E. M., Gummadi, K. P., and Weller, A. Human perceptions of fairness in algorithmic decision making: A case study of criminal risk prediction. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, pp. 903–912, 2018.
- Grover, A. and Ermon, S. Boosted generative models. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Grover, A., Dhar, M., and Ermon, S. Flow-gan: Combining maximum likelihood and adversarial learning in generative models. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Grover, A., Song, J., Agarwal, A., Tran, K., Kapoor, A., Horvitz, E., and Ermon, S. Bias correction of learned generative models using likelihood-free importance weighting. In *NeurIPS*, 2019.
- Grover, A., Chute, C., Shu, R., Cao, Z., and Ermon, S. Alignflow: Cycle consistent learning from multiple domains via normalizing flows. In *AAAI*, pp. 4028–4035, 2020.
- Gulrajani, I., Raffel, C., and Metz, L. Towards gan benchmarks which require generalization. 2018.

- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, pp. 3315–3323, 2016.
- Hashimoto, T. B., Srivastava, M., Namkoong, H., and Liang, P. Fairness without demographics in repeated loss minimization. *arXiv preprint arXiv:1806.08010*, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Heidari, H., Ferrari, C., Gummadi, K., and Krause, A. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Advances in Neural Information Processing Systems*, pp. 1265–1276, 2018.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pp. 6626–6637, 2017.
- Hong, E. 23andme has a problem when it comes to ancestry reports for people of color. *Quartz*. Retrieved from <https://qz.com/765879/23andme-has-a-race-problem-when-it-comes-to-ancestry-reports-for-non-whites>, 2016.
- Horvitz, D. G. and Thompson, D. J. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 1952.
- Huang, C.-Z. A., Cooijmnaas, T., Roberts, A., Courville, A., and Eck, D. Counterpoint by convolution. *ISMIR*, 2017.
- Im, D. J., Ma, H., Taylor, G., and Branson, K. Quantitatively evaluating gans with divergences proposed for training. *arXiv preprint arXiv:1803.01045*, 2018.
- Jo, E. S. and Gebru, T. Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 306–316, 2020.
- Kaerer-Chen, C., Dubois, E., Schüür, F., and Moss, E. Positionality-aware machine learning: translation tutorial. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 704–704, 2020.
- Kamiran, F. and Calders, T. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, pp. 2668–2677, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. Mmd gan: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems*, pp. 2203–2213, 2017.
- Lim, J. H. and Ye, J. C. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.
- Liu, L. T., Dean, S., Rolf, E., Simchowitz, M., and Hardt, M. Delayed impact of fair machine learning. *arXiv preprint arXiv:1803.04383*, 2018.
- Lopez-Paz, D. and Oquab, M. Revisiting classifier two-sample tests. *arXiv preprint arXiv:1610.06545*, 2016.
- Louizos, C., Swersky, K., Li, Y., Welling, M., and Zemel, R. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*, 2015.
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., and Gebru, T. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pp. 220–229, 2019.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pp. 271–279, 2016.
- Odena, A., Olah, C., and Shlens, J. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2642–2651. JMLR. org, 2017.
- Oord, A. v. d., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., Driessche, G. v. d., Lockhart, E., Cobo, L. C., Stimberg, F., et al. Parallel wavenet: Fast high-fidelity speech synthesis. *arXiv preprint arXiv:1711.10433*, 2017.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.
- Podesta, J., Pritzker, P., Moniz, E., Holdren, J., and Zients, J. Big data: seizing opportunities, preserving values. *Executive Office of the President, The White House*, 2014.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. 2019.

- Ratner, A., Bach, S. H., Ehrenberg, H., Fries, J., Wu, S., and Ré, C. Snorkel: Rapid training data creation with weak supervision. *Proceedings of the VLDB Endowment*, 11(3):269–282, 2017.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Rosca, M., Lakshminarayanan, B., Warde-Farley, D., and Mohamed, S. Variational approaches for auto-encoding generative adversarial networks. *arXiv preprint arXiv:1706.04987*, 2017.
- Ryu, H. J., Adam, H., and Mitchell, M. Inclusivefacenet: Improving face attribute detection with race and gender diversity. *arXiv preprint arXiv:1712.00193*, 2017.
- Sattigeri, P., Hoffman, S. C., Chenthamarakshan, V., and Varshney, K. R. Fairness gan: Generating datasets with fairness properties using a generative adversarial network. In *Proc. ICLR Workshop Safe Mach. Learn*, volume 2, 2019.
- Schmidt, M., Schwiegelshohn, C., and Sohler, C. Fair coresets and streaming algorithms for fair k-means clustering. *arXiv preprint arXiv:1812.10854*, 2018.
- Sheng, E., Chang, K.-W., Natarajan, P., and Peng, N. The woman worked as a babysitter: On biases in language generation. *arXiv preprint arXiv:1909.01326*, 2019.
- Song, J., Kalluri, P., Grover, A., Zhao, S., and Ermon, S. Learning controllable fair representations. *arXiv preprint arXiv:1812.04218*, 2018.
- Sugiyama, M., Suzuki, T., and Kanamori, T. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Tao, C., Chen, L., Hénao, R., Feng, J., and Duke, L. C. Chi-square generative adversarial network. In *International Conference on Machine Learning*, pp. 4894–4903, 2018.
- Tommasi, T., Patricia, N., Caputo, B., and Tuytelaars, T. A deeper look at dataset bias. In *Domain Adaptation in Computer Vision Applications*, pp. 37–55. Springer, 2017.
- Torralba, A., Efros, A. A., et al. Unbiased look at dataset bias. In *CVPR*, volume 1, pp. 7. Citeseer, 2011.
- Tran, D., Ranganath, R., and Blei, D. Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, pp. 5523–5533, 2017.
- Turner, R., Hung, J., Saatci, Y., and Yosinski, J. Metropolis-hastings generative adversarial networks. *arXiv preprint arXiv:1811.11357*, 2018.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Xu, D., Yuan, S., Zhang, L., and Wu, X. Fairgan: Fairness-aware generative adversarial networks. In *2018 IEEE International Conference on Big Data (Big Data)*, pp. 570–575. IEEE, 2018.
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., and Dwork, C. Learning fair representations. In *International Conference on Machine Learning*, pp. 325–333, 2013.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint*, 2017.
- Ziwei Liu, Ping Luo, X. W. and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.