

AIM-Fair: Advancing Algorithmic Fairness via Selectively Fine-Tuning Biased Models with Contextual Synthetic Data

Zengqun Zhao Ziquan Liu Yu Cao Shaogang Gong Ioannis Patras

Centre for Multimodal AI, Queen Mary University of London

{zengqun.zhao, ziquan.liu, yu.cao, s.gong, i.patras}@qmul.ac.uk

Abstract

Recent advances in generative models have sparked research on improving model fairness with AI-generated data. However, existing methods often face limitations in the diversity and quality of synthetic data, leading to compromised fairness and overall model accuracy. Moreover, many approaches rely on the availability of demographic group labels, which are often costly to annotate. This paper proposes AIM-Fair, aiming to overcome these limitations and harness the potential of cutting-edge generative models in promoting algorithmic fairness. We investigate a fine-tuning paradigm starting from a biased model initially trained on real-world data without demographic annotations. This model is then fine-tuned using unbiased synthetic data generated by a state-of-the-art diffusion model to improve its fairness. Two key challenges are identified in this fine-tuning paradigm, 1) the low quality of synthetic data, which can still happen even with advanced generative models, and 2) the domain and bias gap between real and synthetic data. To address the limitation of synthetic data quality, we propose Contextual Synthetic Data Generation (CSDG) to generate data using a text-to-image diffusion model (T2I) with prompts generated by a context-aware LLM, ensuring both data diversity and control of bias in synthetic data. To resolve domain and bias shifts, we introduce a novel selective fine-tuning scheme in which only model parameters more sensitive to bias and less sensitive to domain shift are updated. Experiments on CelebA and UTKFace datasets show that our AIM-Fair improves model fairness while maintaining utility, outperforming both fully and partially fine-tuned approaches to model fairness. The code is available at <https://github.com/zengqunzhao/AIM-Fair>.

1. Introduction

Recent research has raised significant concerns about fairness and bias in machine learning models [30]. These mod-

Corresponding author: Ziquan Liu {ziquan.liu@qmul.ac.uk}.

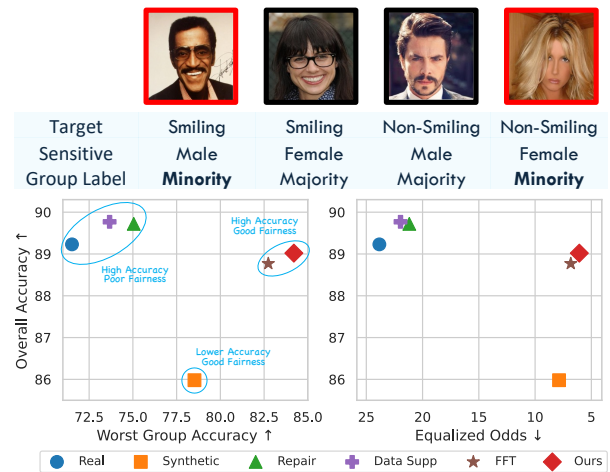


Figure 1. Facial attributes classification (FAC) on the CelebA dataset based on different training strategies, in which **Smiling** is the target attribute and **Gender** is the protected attribute. This shows different learning strategies result in variable model utility (Overall Accuracy) vs. model fairness (Worst Group Accuracy and Equalized Odds) on demographic groups. A model trained solely on real data if biased exhibits high accuracy but poor fairness scores. Conversely, models trained on balanced synthetic data show better fairness but poorer accuracy due to a “domain gap” between the real and synthetic data, and a lack of selective model fine-tuning when synthetic data is deployed. Strategies to repair imbalances in real data [9, 56] or to supplement real data with synthetic data [47] marginally increase accuracy but do little to improve fairness. Our method for selective fine-tuning of a pre-trained (biased) model with synthetic data not only preserves model accuracy but also substantially improves fairness, outperforming fully fine-tuning (FFT) in both model utility and fairness.

els often demonstrate varying performance across different demographic groups, leading to unfair outcomes. To mitigate the spurious correlation caused by learning from imbalanced (biased) data, regularizations were used as optimization objectives [1, 10, 12, 23]. Methods like distributionally robust optimization (DRO) [12] optimizes the worst-case performance, while invariant risk minimization (IRM) [1] learns unbiased representations with invariance to different environments. Influenced by the success of repre-




Figure 2. The effects of selective fine-tuning by layer-wise freezing from facial attribute classifications on the CelebA dataset, with Smiling as the target attribute and Male as the protected attribute: When only the fully connected (FC) layer is frozen the model shows improved worst demographic group accuracy and reduced equalized odds, i.e. more fair, but sacrifices some utility (overall accuracy). This indicates increased cross-domain generalisability (real and synthetic). Conversely, freezing only block 2 while fine-tuning the remaining parameters results in high overall accuracy but poorer fairness, i.e. further enhanced domain-bias specificity. When only block 1 is frozen, the model not only maintained equalized odds but also increased utility (overall accuracy) and worst group accuracies.

sensation learning [51], attempts were made to learn a fair feature representation invariant to protected facial attributes [6, 29, 34, 54]. However, these methods rely on demographic group labelling, often unavailable in practice [2, 4]. More recent advances in generative models have utilized generative augmentation for a more balanced training data distribution [8, 9, 17, 36, 40, 59]. These techniques primarily focus on image editing, either on real data [36, 40] or synthetic data [9], to create less biased training data. Yet, these methods still require additional group annotations of the data they edit. To address this limitation, DiGA [59] proposed to create an unbiased training set by editing spurious facial attributes to a random degree while keeping a target facial attribute unchanged without knowing each sample’s group label. However, DiGA increases training data size multiple times, leading to substantially higher training costs, and the quality of generated data is limited due to the unknown group labels.

Recent advancements in text-to-image generative models showcase impressive data fidelity [43], yet their potential for improving fairness through data expansion has not been fully explored by the fairness community. This raises the question: **can AI-generated synthetic data play a crucial role in mitigating biases within machine learning models?** This research question is particularly important given the successful application of AI-generated data in fine-tuning large language models (LLMs) [11], data augmentation [64], and long-tail recognition [47], alongside some reported limitations of synthetic data effectiveness [49]. This work presents a comprehensive empirical investigation into whether fine-tuning high-quality, balanced generative data from a contemporary text-to-image model can counteract model biases caused by training on imbalanced real data. We identify two key challenges in bias-correcting fine-tuning with synthetic data: (1) *A data-related challenge arising from linguistic ambiguity of the textual prompt and/or model misrepresentation [48, 57], which results in low-quality and low-diversity generated data.* (2) *A model learning challenge caused by both a domain shift (synthetic vs. real) and a bias shift (unbiased vs. biased) between*

the real and the synthetic data. Fine-tuning blindly on the synthetic will result in a model with decreased utility.

To address the data quality issue, we propose Contextual Synthetic Data Generation (CSDG). Existing Latent Diffusion Models (LDMs) [3, 36, 41, 45] often use vision-language models, such as CLIP, for cross-modality alignment. To yield more diverse and better fine-grained image generation, we formulate a contextual synthetic data generation strategy that uses more detailed text descriptions from expansive linguistic expressions of LLMs to condition an LDM with richer, context-driven text. As a result, the generated images cover more scenarios and provide more details, enhancing diversity and mitigating bias in real data. In contrast to other methods that edit or manipulate real data samples, a key strength of our method is that it does not require annotating the Protected Group Attributes of real data.

The model learning challenge is due to the existence of two types of shifts between the generated and the real data, namely, the desired bias shift and the undesired domain shift (image quality, realism, scene characteristics). To address this challenge we propose to locate and update model parameters that are more sensitive to the bias shift and less sensitive to the domain shift. Our observation is that some parameters are more sensitive to data distribution shift (cross-domain generalisability) – we call *domain-sensitive parameters*, while others are more sensitive to demographic group discrepancies (real data domain bias specificity), which we call *fairness-sensitive parameters*. This observation is supported by the experiments shown in Fig. 2. This is from fine-tuning a real data pre-trained model on balanced synthetic data while keeping parameters at different layers frozen. To identify which parameters to update, we propose a novel selection scheme in which the gradient differences between the updates by the real (biased) dataset and two synthetic datasets with one unbiased and another biased. Ranking the gradient differences between the synthetic-biased data and the synthetic-unbiased data reveals parameters sensitive to fairness (fairness-sensitive parameters). In contrast, inverse ranking the gradient differences between the real data and the synthetic-biased data

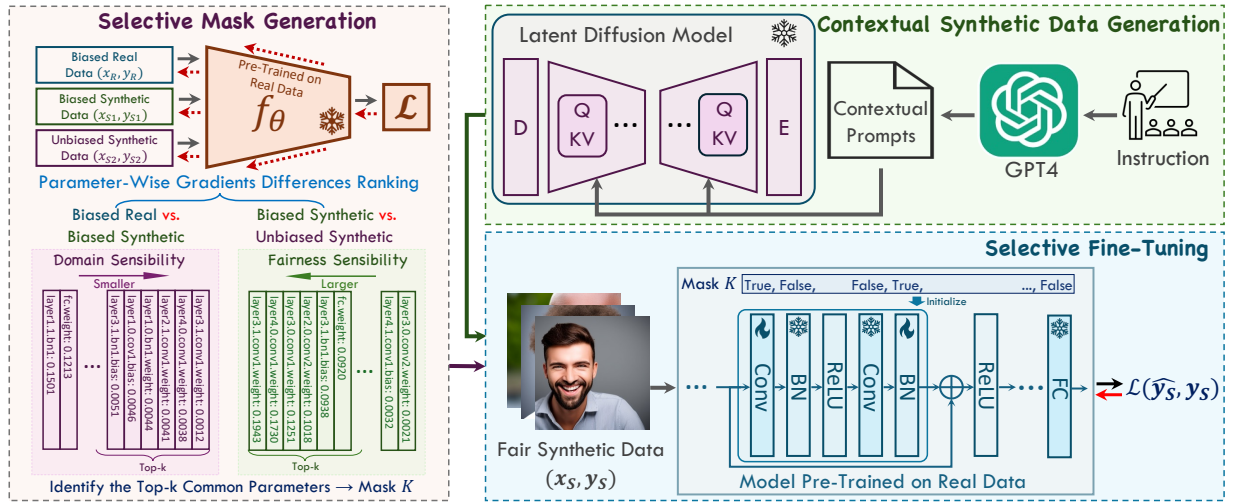


Figure 3. A selective fine-tuning model consisting of three parts: (1) Contextual Synthetic Data Generation (CSDG) for generating diverse images using GPT-4 generated prompts, (2) Selective Mask Generation (SMG) for creating a selection mask that determines which parameters are updated during fine-tuning, and (3) Selective Fine-Tuning (SFT) to enhance model fairness obtained from synthetic data whilst simultaneously to preserve model utility yielded from real data in pre-training.

discovers parameters less sensitive to domain shift (domain-insensitive parameters). In fine-tuning, a selection mask is constructed as the intersection of the top-k rankings. This selection mask is applied to initialize the pre-trained model, so that only the selected parameters are updated using the balanced synthetic data.

Fig. 1 presents a comparison of the results across various training strategies using real and/or synthetic data. Fig 3 shows an overview of our model, which consists of three parts: Contextual Synthetic Data Generation (CSDG), Selective Mask Generation (SMG), and Selective Fine-Tuning (SFT). CSDG uses a pre-trained and fixed LDM to generate high-quality images, where a series of contextual prompts serve as the conditions. SMG creates a selection mask that determines which parameters are updated during fine-tuning. SFT is designed to correct bias in the model. We initialize a pre-trained model using the parameter selection mask. Then the model is fine-tuned on balanced synthetic data to enhance its fairness while retaining the model utility. Our contributions are as follows:

- We investigate a fine-tuning paradigm that mitigates model bias stemming from unbalanced real data using synthetic data generated by a text-to-image (T2I) process without requiring demographic group annotations.
- We design contextual synthetic data generating by using a T2I diffusion mode with prompts generated by a context-aware LLM, ensuring both data diversity and control of bias in synthetic data.
- We introduce a selective fine-tuning method for fair model learning, which identifies domain-sensitive and fairness-sensitive parameters for improving model fairness and utility simultaneously.
- Our method outperforms both full and partial fine-tuning methods and achieves superior performance compared to state-of-the-art methods across several datasets.

2. Related Work

Mitigating Model Bias. Current methods for model fairness can be categorized into three types: *pre-processing*, *in-processing*, and *post-processing*. **Pre-processing** approaches modify sample distributions of protected variables or perform transformations to remove discrimination from the data. Many recent works use generative models to create balanced, unbiased datasets [36, 40, 58]. **In-processing** methods incorporate fairness metrics into model optimization to maximize both performance and fairness [1, 10, 12, 20, 23, 55]. Some studies explore mitigating bias without group annotations, such as Just Train Twice (JTT) [24], which upweights misclassified examples to improve worst-group performance, and Cross-Task Minimization (XRM) [37], which trains twin classifiers to reduce spurious correlations. **Post-processing** methods apply transformations to model outputs to improve fairness, such as model calibration [25, 33, 38] and thresholding [5, 13] to align predicted positive outcomes with actual positive examples across groups. Our method falls within the *pre-processing* paradigm, but instead of simply balancing real data with synthetic data, we focus on using synthetic data to enhance model fairness.

Improving Model Fairness on Generative Data. Generative models have advanced rapidly in recent years [43, 46], with several studies exploring their use to improve model fairness through generative data [9, 17, 59]. Much of the prior research has focused on generating counterfactual samples to assess fairness [8, 17], and generative methods have also made strides in bias mitigation by creating balanced, unbiased datasets [36, 40, 58]. Instead of generating counterfactual samples based on real data, D’Incà *et al.* [9] use a diffusion model for uncontrolled image generation, followed by manipulation of the synthetic images in a semantic space. However, these approaches require additional

group annotations. Zhang *et al.* [59] use generative models to identify spurious attributes that may bias the model, then edit each image to modify these attributes. They train a fair FAC model on the augmented dataset, enhancing its invariance to these spurious variations. While this approach does not require protected group labels for training the FAC model, annotations are still needed to train the generative models that detect and edit spurious attributes. Additionally, generating accurate counterfactual images remains challenging, as the edits are applied randomly to images without known protected attributes. In contrast, our method uses generative data from a text-to-image LDM, enabling greater control over the diversity of the generated data and allowing a broader range of scenarios and variations.

Model Fine-Tuning. A common approach to transfer learning in the presence of distribution shifts is to fine-tune the last few layers of a pre-trained model, retaining the learned features while adapting to the new task [27, 65]. To prevent overfitting during this fine-tuning process, existing methods suggest using a smaller learning rate compared to the initial pretraining phase [21], freezing the early backbone layers and gradually unfreezing them [15], or applying different learning rates to each layer [42]. Lee *et al.* [22] introduced surgical fine-tuning, which showed that selectively fine-tuning a subset of layers can either match or outperform conventional fine-tuning strategies. Recent research [19] indicates that training a carefully selected subset of layers while keeping the remaining weights frozen at their initial values can lead to varying contributions from different layers across the network to overall performance. Our proposed method inherits some common findings from these works but we further investigated with parameters-wise selective fine-tuning, which details the parameters’ sensibility to the property from the pre-train data distribution and the property from the downstream data distribution.

3. Methods

This section first introduces how to generate contextual synthetic data with LLM-generated prompts. Then, we detail how to obtain the parameters selection mask. Finally, we present how to conduct model fair fine-tuning with a selection mask on balanced synthetic data. We provide the algorithm for the selective fine-tuning in the *Appendix*.

3.1. Contextual Synthetic Data Generation

Our method attempts to correct the biased model trained on imbalanced real data by fine-tuning on balanced synthetic data, so the quality of the synthetic data is crucial to the success of fair learning. Although current text-to-image models have demonstrated remarkable performance, several studies indicate that directly expressing the desired attributes in the prompt often results in sub-optimal outcomes due to linguistic ambiguity or model misrepresentation [48, 57]. For example, as shown in Fig. 4, to generate a face photo with



Figure 4. Generated images of *Smiling Male* conditioned on different prompts. Compared to the plain prompt, our contextual prompts enhance the diversity. More on UTKFace in *Appendix*.

both target and protected attributes, one might use a prompt like “Portrait face photo of a smiling male.” as suggested in [56]. However, such manually designed prompts may lead to stereotypical image generation, excluding certain attributes or minority groups, which in turn can introduce bias in other attributes, such as age or hairstyle [57].

The recent work in zero-shot [31] and supervised [62] classification suggest that leveraging additional contextual and detailed information can enhance vision-language alignment. Considering current text-to-image models [3, 36, 41, 45] often use vision-language models, such as CLIP [39], for cross-modality alignment, we propose a contextual synthetic data generation strategy that leverages the powerful linguistic capabilities of large language models, such as GPT-4, to condition an LDM with richer and context-driven text. As a result, the generated images cover more scenarios and provide more details, enhancing diversity and mitigating the bias in the characteristics not involved in the target and protected. The structure of the CSDG is shown in the upper right of Fig. 3.

Concretely, the instruction provided for the GPT-4 model following the format: “{Task}, {Number of Prompts}, {Target Attribute}, {Protected Attribute}, {Other Detailed Descriptions}, and {Prompt Format}”. For generating facial photos, for example, the other detailed descriptions contain specific facial features, hair characteristics, eye-related details, and head orientation or angle. Our text-to-image generation model is a pre-trained latent diffusion model—Stable Diffusion (SD) [43], which reverses noises applied to the latent embedding of images. SD contains a variational autoencoder (VAE) [28], a CLIP text encoder [39], and a U-Net [44]. During the inference, the random Gaussian noise $\varepsilon_t \sim \mathcal{N}(0, \mathbf{I})$ and the contextual prompt features $c = (w_1, w_2, \dots, w_n)$ encoded via a CLIP text

encoder $\Psi(\cdot)$ will be fed into the U-Net to condition the denoising process, where n is the number of the contextual prompts. We provide detailed instruction and contextual prompts in the *Appendix*. The empirical results in Tab. 5 demonstrate that the generated contextual synthetic data mitigate domain shift and improve model fairness.

3.2. Selective Mask Generation

As recent studies [50, 53, 64] have mentioned, synthetic data often fail to align with the real data distribution due to domain shift, suggesting that fine-tuning the pre-trained model directly on the balanced synthetic data may not learn the desired properties effectively. This is caused by the fact that there is both a domain shift (synthetic vs. real) and a bias shift (unbiased vs. biased) between the real and the synthetic data, and fine-tuning with the domain shift will result in a model with decreased utility. The empirical finding shown in Fig. 2 demonstrates that when fine-tuning a real data pre-trained model on balanced synthetic data, some parameters are more sensitive to the data distribution shift, called *domain-sensitive parameters*, while some are more sensitive to group discrepancies, called *fairness-sensitive parameters*. To discover the parameters' sensibility towards different scenarios, we propose to construct three distinct datasets with different distributions to elicit the model responses by calculating the parameter-wise gradients.

We aim to fine-tune a pre-trained model $f_\theta(x)$ to improve fairness while accounting for domain differences between real and synthetic data, hence, we want to find the parameters which are more sensitive to fairness while less sensitive to domain shift. Specifically, we start by constructing three different datasets: biased real data $\{(x_i^{(R)}, y_i^{(R)})\}_{i=1}^{N_R} \in \mathcal{D}_R$ representing training data with inherent biases, biased synthetic data $\{(x_i^{(S_1)}, y_i^{(S_1)})\}_{i=1}^{N_{S_1}} \in \mathcal{D}_{S_1}$ mirroring the unfairness of the real data, and unbiased synthetic data $\{(x_i^{(S_2)}, y_i^{(S_2)})\}_{i=1}^{N_{S_2}} \in \mathcal{D}_{S_2}$ designed to be fair, mitigating biases. N_R , N_{S_1} , and N_{S_2} is the image number of the corresponding dataset.

We then compute the gradients \mathbf{g}_R , \mathbf{g}_{S_1} , and \mathbf{g}_{S_2} of the loss function $\mathcal{L}(\theta; x, y)$, which is defined from a binary softmax loss. Rather than considering fine-grained, scalar-wise parameters, we select parameters at the level of weights and biases, denoted by θ , which includes weights $W^{(l)}$ and biases $b^{(l)}$ for the Convolution Layer, Batch Normalization, and Fully Connected Layer, where l refers to the layer index. Then, for the model parameters θ on each dataset:

$$\mathbf{g}_R = \frac{1}{N_R} \sum_{i=1}^{N_R} \nabla_\theta \mathcal{L}(\theta; x_i^{(R)}, y_i^{(R)}) \quad (1)$$

$$\mathbf{g}_{S_1} = \frac{1}{N_{S_1}} \sum_{i=1}^{N_{S_1}} \nabla_\theta \mathcal{L}(\theta; x_i^{(S_1)}, y_i^{(S_1)}) \quad (2)$$

$$\mathbf{g}_{S_2} = \frac{1}{N_{S_2}} \sum_{i=1}^{N_{S_2}} \nabla_\theta \mathcal{L}(\theta; x_i^{(S_2)}, y_i^{(S_2)}) \quad (3)$$

where ∇ is the gradient operator. We then calculate the gradient differences to capture how each parameter behaves across different data distributions. For parameter θ_j :

$$\begin{aligned} \Delta_{1,j} &= |\mathbf{g}_{R,j} - \mathbf{g}_{S_1,j}| \\ \Delta_{2,j} &= |\mathbf{g}_{S_1,j} - \mathbf{g}_{S_2,j}| \end{aligned} \quad (4)$$

where Δ_1 measures parameters sensitive to the domain shift while Δ_2 identifies parameters crucial for fairness, j denotes the index of all the parameters. To obtain the parameters which are less affected by domain differences and more impactful for fairness, we conduct a ranking in ascending order on Δ_1 (smaller differences first) and a ranking in descending order on Δ_2 (larger differences first):

$$R_1 = \text{argsort}(\Delta_1); R_2 = \text{argsort}(-\Delta_2) \quad (5)$$

Then we can find the intersections from the *top-k* parameters are both less sensitive to the domain gap and significant for fairness by $K = K_1 \cap K_2$, where

$$\begin{aligned} K_1 &= \{\theta_j \mid j \in R_1[1 : k]\} \\ K_2 &= \{\theta_j \mid j \in R_2[1 : k]\} \end{aligned} \quad (6)$$

Finally, the selection mask can be obtained by:

$$M_j = \begin{cases} \text{True}, & \text{if } \theta_j \in K \\ \text{False}, & \text{otherwise} \end{cases} \quad (7)$$

3.3. Selective Fine-Tuning on Synthetic Data

Our goal is to fine-tune the pre-trained model $f_\theta(x)$ on the balanced synthetic dataset \mathcal{D}_{S_2} following the ERM framework, while only updating selected parameters as indicated by the mask M . Specifically, given the model $f_\theta(x)$ pre-trained on the biased real data, balanced synthetic data $\{(x_i^{(S)}, y_i^{(S)})\}_{i=1}^{N_{S_2}}$, and the parameters selection mask M , during optimization, we apply the mask to the gradients so that only the selected parameters are updated:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \left(M \odot \nabla_\theta \mathcal{L} \left(f_{\theta^{(t)}}(x_i^{(S_2)}), y_i^{(S_2)} \right) \right) \quad (8)$$

where $\theta^{(t)}$ are the parameters at iteration t , η is the learning rate, and \odot denotes element-wise multiplication. Notably, once the mask M is provided by SMG, it will be applied throughout the entire fine-tuning optimization process.

4. Experiments

This section presents the experimental setup and results. We begin by describing the datasets, followed by the implementation details. The main results and ablation studies are presented at the end.

4.1. Datasets

CelebA [26] contains over 200,000 facial images with 40 binary attribute annotations. Following the setting of the previous works [35, 58, 59], we set Male and Young to protected attributes and selected Smiling and Young as the target attribute which has the highest correlation with the protected attributes. For each experiment, we randomly sample a biased subset as a training dataset with a size of 20,000 images, where the majority group and minority group have

Table 1. Comparisons to other methods on the CelebA dataset under settings of varied target and protected attributes.

Methods	T: Smiling, P: Male			T: Smiling, P: Young			T: Young, P: Male		
	ACC (↑)	WST(↑)	EO (↓)	ACC (↑)	WST(↑)	EO (↓)	ACC (↑)	WST(↑)	EO (↓)
ERM [14]	88.20	70.10	25.30	88.30	71.50	15.60	77.70	42.00	52.00
CVaR DRO [23]	87.30	74.00	22.80	87.00	76.10	13.90	75.40	42.30	48.80
EIIL [7]	87.90	75.60	19.70	87.90	72.50	13.30	77.50	45.60	39.20
LfF [32]	87.10	77.50	17.00	85.30	72.90	14.30	77.40	44.20	43.60
JTT [24]	88.00	74.80	19.40	87.60	73.30	14.20	76.30	43.60	47.70
MAPLE [63]	88.10	72.00	19.60	88.10	73.60	13.60	76.30	46.20	43.50
DiGA [59]	88.40	81.90	7.40	89.10	78.50	9.50	80.00	51.30	33.30
AIM-Fair (Ours)	89.02	84.20	6.07	90.21	87.78	5.44	78.19	54.89	28.18

90% and 10% of the sample size respectively. We report performance on the whole original test dataset. **UTK Face** [60] consists of over 20,000 facial images with three kinds of annotations: gender, age, and ethnicity. Following the experimental setup in the previous works [17, 18, 59], we define a binary spurious attribute “Ethnicity” based on whether the facial image is white or not. The task is to predict the Gender. We randomly sample a biased subset of 10,000 images, with the same bias degree as CelebA. We also construct a balanced and unbiased test dataset consisting of 3,200 images.

4.2. Fairness Metrics

Our goal is to learn a fair and accurate model. For model utility, we present the *overall accuracy (ACC)* and *group accuracy*. For fairness, follow previous work [35, 58, 59] we use *equalized odds (EO)* [13], defined as:

$$\sum_{\forall y, \hat{y}} \left| P_{s_0}(\hat{Y} = \hat{y} | Y = y) - P_{s_1}(\hat{Y} = \hat{y} | Y = y) \right|, \quad (9)$$

where \sum is the averaged sum, Y target label, \hat{Y} classifier predictive label, and $s_0, s_1 \in S$ protected attributes. The *worst-group accuracy (WST)* defined as:

$$\min_{\forall y \in \mathcal{Y}, \forall s \in S} P_s(\hat{Y} = y | Y = y) \quad (10)$$

and *group standard deviation (STD)* [52] defined as:

$$\text{std}_{\forall y \in \mathcal{Y}, \forall s \in S} P_s(\hat{Y} = y | Y = y) \quad (11)$$

4.3. Implementation Details

Following previous work [59], we use ResNet-18 [14] as the backbone for all experiments. We use stable diffusion v2.1 as our latent diffusion model for generating images. We generated 10,000 images for each group and then randomly sampled N_S . To be consistent with the real data image number N_R , we set the $N_S = N_R/G$, where G is the group number. During the pre-training phase, we set the batch size to 128 and trained the model for 15 epochs. The initial learning rate was set to 0.01, which was then reduced by a factor of 0.01 at the 10th epoch. During the fine-tuning phase, we set the batch size to 128 and trained the model for 10 epochs. And the learning rate is searched from {0.4, 0.5, 0.6}. All models are optimized by the SGD optimizer and trained on the Tesla A100 GPU based on the open-source PyTorch platform. To obtain more stable and reliable results, we conducted all experiments 10 times with different

Table 2. Comparisons to other methods on the CelebA dataset (T=Smiling, P=Male) under settings of training set sizes.

Methods	Samples Ratio = 50%			Samples Ratio = 25%			Samples Ratio = 10%		
	ACC (↑)	WST(↑)	EO (↓)	ACC (↑)	WST(↑)	EO (↓)	ACC (↑)	WST(↑)	EO (↓)
ERM [14]	87.50	67.80	26.10	87.10	65.90	27.70	86.90	62.80	28.90
CVaR DRO [23]	86.60	72.90	22.10	86.60	72.30	22.40	85.50	69.10	27.30
EIIL [7]	86.20	71.30	22.50	85.90	69.60	25.40	86.80	64.20	26.70
LfF [32]	86.90	75.50	19.40	85.90	72.10	23.60	85.50	66.10	27.70
JTT [24]	87.30	72.90	20.10	86.70	71.10	20.60	86.80	67.10	23.10
MAPLE [63]	87.40	73.70	23.80	87.00	72.70	24.20	85.60	69.20	27.10
DiGA [59]	88.40	81.10	7.80	88.40	78.30	8.00	88.30	78.80	8.40
AIM-Fair (Ours)	88.85	83.99	6.28	88.89	82.54	7.59	87.90	81.73	8.16

random seeds and then used the average as the final result.

4.4. Comparisons to State of the Art

We compared our method against seven contemporary techniques, including the baseline ERM [14] method and six debiasing models all of which do not require protected attribute labels: Two regularization-based methods (CVaR DRO [23] and EIIL [7]); three reweighting-based methods (LfF [32], JTT [24], and MAPLE [63]); and generative-model-based method (DiGA [59]). Our comparison covers two settings: different target and protected attributes and varying numbers of target labels.

Tab. 1 shows that ERM achieves good accuracy but suffers from significant unfairness. Other debiasing methods, while improving fairness to some extent, generally sacrifice accuracy. DiGA [59] improves fairness while maintaining accuracy by using a generative model to edit real data and create more balanced training data. In contrast, our method generates images from scratch using a text-driven latent diffusion model. It is evident that our method outperforms the best of the existing models DiGA [59] in both fairness and accuracy in all categories, except “Young / Male” where we come close second to DiGA. We outperform all other methods consistently. We also conducted comparisons with other methods under smaller training set sizes, with the subsampling ratio of 50%, 25%, and 10%. Tab. 2 shows that our method outperforms all others consistently, except on ACC score for Lable Ratio at 10% where we come close second to DiGA. Critically, our method maintains robust model utility and fairness even with varying amounts of real training data, with these improvements attributed to the selective updating with balanced synthetic data.

4.5. Ablation Analysis

Comparisons of Different Training Strategies. To evaluate the effectiveness of the proposed selective fine-tuning, we first compare our method with several other strategies trained on different types of data. These include the baseline, which trains the model using conventional ERM on real data, training the model solely on synthetic data, supplementing real data with synthetic data [47], and balancing the real data [9, 56]. Furthermore, we compare common fine-tuning methods, such as linear probing and full fine-tuning.

The results in Tab. 3 indicate that both data supplement-

Table 3. Comparisons of varied training strategies on CelebA and UTKFace datasets.

Methods	Target	CelebA												UTKFace							
		Protected		T: Smiling; P: Male				Protected		T: Smiling; P: Young				Protected		T: Female; P: White					
		P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)	P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)	P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)		
Baseline [14]	T=0 T=1	83.57 95.36	98.50 71.52	89.23	71.52	23.84	10.65	78.72 94.37	96.09 86.42	90.19	78.72	17.37	6.92	79.58 95.75	96.16 83.96	88.86	79.58	16.59	7.26		
Trained on Synthetic Data	T=0 T=1	86.82 86.39	89.40 78.52	85.98	78.52	7.87	4.07	85.45 82.42	85.76 85.27	85.12	82.42	3.50	1.34	82.71 84.90	88.12 86.04	85.44	82.71	5.50	1.96		
Data Supplementation [47]	T=0 T=1	84.39 95.40	98.39 73.67	89.77	73.67	21.96	9.76	78.05 94.90	95.36 87.62	<u>90.35</u>	78.05	17.31	7.02	82.12 95.62	93.97 89.28	90.25	82.12	11.85	5.24		
Data Repairing [9, 56]	T=0 T=1	82.82 96.04	98.10 75.05	<u>89.72</u>	75.05	21.18	9.50	79.50 94.51	95.19 88.35	90.68	79.50	15.69	6.30	81.15 95.71	94.22 86.61	<u>89.42</u>	81.15	13.37	5.89		
Linear Prob	T=0 T=1	84.55 94.82	96.22 77.41	89.69	77.41	17.51	7.71	85.20 90.43	93.81 87.88	90.14	85.20	8.61	3.18	81.06 94.59	91.51 86.42	88.40	81.06	10.59	5.14		
Fully Fine-Tuning	T=0 T=1	88.81 89.55	91.62 82.74	88.77	<u>82.74</u>	<u>6.83</u>	<u>3.31</u>	85.98 86.02	90.06 86.81	87.85	<u>85.98</u>	<u>4.37</u>	<u>1.68</u>	83.00 90.74	88.04 89.83	87.90	<u>83.00</u>	<u>5.40</u>	2.99		
AIM-Fair (Ours)	T=0 T=1	88.16 90.25	91.40 84.20	89.02	84.20	6.07	2.74	88.19 88.98	93.64 87.78	90.21	87.78	5.44	2.34	84.26 90.62	89.08 89.25	88.30	84.26	4.81	<u>2.41</u>		

Table 4. Comparisons of varied partial fine-tuning strategies on CelebA and UTKFace datasets.

Methods	Target	CelebA												UTKFace							
		Protected		T: Smiling; P: Male				Protected		T: Smiling; P: Young				Protected		T: Female; P: White					
		P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)	P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)	P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)		
Best Random Selection	T=0 T=1	87.93 91.31	92.81 82.17	<u>89.33</u>	82.17	9.14	4.09	85.52 88.25	91.53 87.63	88.95	85.52	6.02	<u>2.16</u>	83.50 90.30	88.85 89.01	87.92	83.50	5.38	2.61		
Best Sub-Tuning [19] (By Updating One Block)	T=0 T=1	87.50 93.09	94.75 80.64	90.04	80.64	12.45	5.52	87.28 89.15	94.63 87.32	90.33	87.28	7.34	3.00	83.05 92.07	89.56 91.34	89.01	83.05	6.51	3.56		
Best Sub-Tuning [19] (By Freezing One Block)	T=0 T=1	88.69 90.48	91.09 83.48	89.06	<u>83.48</u>	<u>7.00</u>	<u>3.00</u>	86.55 87.08	91.89 86.55	88.65	86.55	5.34	2.25	83.58 88.96	88.34 88.81	87.42	83.58	4.76	2.23		
Selective Fine-Tuning (Cosine Similarity)	T=0 T=1	86.07 91.53	91.77 83.33	88.85	83.33	8.40	3.61	87.43 89.52	92.96 88.65	<u>90.26</u>	<u>87.43</u>	5.53	2.05	83.72 90.14	88.82 90.15	88.21	<u>83.72</u>	5.10	2.64		
AIM-Fair (Ours) (Absolute Difference)	T=0 T=1	88.16 90.25	91.40 84.20	89.02	84.20	6.07	2.74	88.19 88.98	93.64 87.78	90.21	87.78	<u>5.44</u>	2.34	84.26 90.62	89.08 89.25	<u>88.30</u>	84.26	<u>4.81</u>	<u>2.41</u>		

tation and data repairing do not significantly improve fairness. Data supplementation, which combines biased real data with balanced synthetic data for co-training, makes it difficult for the model to learn fairness properties from the synthetic data because of the domain gap. While data repairing can create balanced training data, the domain shift between real and synthetic data limits the effectiveness of fair learning. Common transfer learning methods, such as linear probing and full fine-tuning, also face challenges due to domain shifts. Specifically, linear probing is affected by discrepancies in feature representations, while full fine-tuning improves fairness at the cost of reduced model utility.

Comparisons of Different Partial Fine-Tuning. We also compare our method with other partial fine-tuning approaches, including random selection and sub-tuning [19]. For random selection, we set the updating ratios at 40%, 55%, 70%, and 85%, and select the best result as the final one. For sub-tuning, we perform block-wise fine-tuning and block-wise freezing (i.e., updating one block or freezing one block while updating the rest), also selecting the block with the best result.

Tab. 4 shows that fine-tuning only one block consistently yields the best accuracy, but at the cost of fairness, while freezing one block tends to improve fairness but often sacrifices model utility. This demonstrates that coarse-grained approaches, such as block-wise updating or freezing, struggle

Table 5. Classification results on CelebA dataset (T=Smiling, P=Male) under settings of different prompt types and numbers.

Prompt Types (Number)	Target	Protected		T: Smiling; P: Male				Protected		T: Smiling; P: Young			
		P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)	P=0	P=1	ACC (↑)	WST (↑)	EO (↓)	STD (↓)
Plain Prompt (1)	t=0 t=1	65.18 77.47	89.44 43.28	71.65	43.28	34.20	17.07	83.10 59.53	77.72 68.83	72.75	59.53	9.30	8.96
Contextual Prompts (25)	t=0 t=1	77.12 85.99	89.25 69.73	81.85	69.73	16.3	7.65	76.72 85.16	67.90 91.09	79.80	67.90	8.82	8.75
Contextual Prompts (50)	t=0 t=1	89.47 77.26	89.80 72.22	82.66	72.22	5.07	7.66	87.72 79.67	90.43 82.67	85.85	79.67	4.36	4.20
Contextual + Head Poses (50)	t=0 t=1	86.82 86.39	89.40 78.52	85.98	78.52	7.87	4.07	85.45 82.42	85.76 85.27	85.12	82.42	3.50	1.34

to achieve an optimal balance between model utility and fairness. In contrast, our method performs fine-grained, parameter-wise updating, which allows the model to enhance fairness while maintaining utility. As a result, our method achieves the best worst-group accuracy while preserving high overall accuracy.

Additionally, our method uses gradient differences to identify parameter sensitivity to varying data distributions. We also compared absolute gradient differences with the cosine similarity between gradients. The results in Tab. 4 show that using absolute gradient differences yields better results. The possible reason is that the cosine similarity only provides the direction of gradient disparity, whereas the absolute difference directly captures the magnitude of the gradient disparity.

Comparisons of Different Prompts. To evaluate the performance of the varied prompts used for generating images,




Figure 5. Classification results on the CelebA dataset (T=Smiling, P=Male) with different top-k values. The top-k values {40, 45, 50, 55, 60} correspond approximately to {64%, 72%, 80%, 88%, 96%} of the total model parameters.

we trained the model solely on the balanced synthetic data generated with corresponding prompts and tested it on the real test set. Tab. 5 shows the results of different prompt types and quantities, indicating that images generated with contextual prompts lead to significantly better performance in both model accuracy and fairness. We believe this improvement is due to the increased diversity and details in the synthetic images generated from contextual prompts. Moreover, as the number of contextual prompts increases, test performance improves as well. Specifically, incorporating head pose variation into the prompt further enhances both accuracy and fairness in the generated images, as the results shown in the last row of Tab. 5.

Evaluations of Varied Top-k Values. In our method, to identify the intersection parameters from the two gradient difference rankings, we use the top-k selection to ensure that the chosen parameters are more sensitive to fairness and less affected by domain shift. The results for different top-k values are shown in Fig. 5. When k is set to 55, our method achieves the best fairness performance while retaining overall accuracy. Intuitively, a lower k results in fewer updates, leading to better model utility but worse fairness, while a high k involves more parameter updates, which can negatively impact both accuracy and fairness.

Evaluations of Varied Bias Ratio for Synthetic Data Construction. To evaluate the model’s sensitivity to domain shift and group disparity, we construct three distinct datasets with different distributions. In our method, the bias ratio for the biased synthetic data distribution is treated as a hyperparameter, referenced by the error set of real training examples, as suggested in JTT [24]. To assess the impact of varying bias ratios in synthetic data, we explore different ratios. As shown in Tab. 6, a disparity between the synthetic data bias ratio and the real data bias ratio leads to worse performance in both accuracy and fairness. However, despite these differences in bias ratios, our method still outperforms the fully fine-tuning approach. This demonstrates that our

Table 6. Classification results on CelebA dataset (T=Smiling, P=Male) with varied bias ratio of the biased synthetic data.

Bias Ratio	Target	Protected		ACC (↑)	WST (↑)	EOD (↓)	STD (↓)
		P=0	P=1				
Fully Fine-Tuning	t=0 t=1	88.81 89.55	91.62 82.74	88.77	82.74	6.83	3.31
4 : 6	t=0 t=1	85.72 92.69	91.71 83.90	89.23	83.90	8.78	3.77
3 : 7	t=0 t=1	87.05 90.93	91.62 83.52	88.90	83.52	7.72	3.26
2 : 8	t=0 t=1	88.01 90.46	91.76 83.16	88.97	83.16	7.30	3.28
1 : 9	t=0 t=1	88.16 90.25	91.40 84.20	89.02	84.20	6.07	2.74

Table 7. Classification results on CelebA dataset (T=Smiling, P=Male) with different number of synthetic data.

Ratio To Real Data	Target	Protected		ACC (↑)	WST (↑)	EOD (↓)	STD (↓)
		P=0	P=1				
0.5	t=0 t=1	88.65 91.10	93.96 80.80	89.51	80.80	10.31	4.90
1	t=0 t=1	88.16 90.25	91.40 84.20	89.02	84.20	6.07	2.74
1.5	t=0 t=1	89.18 88.38	92.05 81.25	88.34	81.25	7.13	3.98
2	t=0 t=1	89.01 91.26	91.26 89.44	88.81	83.36	6.08	2.96

method is able to identify the model’s sensitivity under the bias ratios different to real data.

Evaluations of Different Number of Synthetic Data. We also evaluate fine-tuning with varying amounts of balanced synthetic data. As shown in Tab. 7, we use the real training data count as a reference and set different ratios to determine the number of synthetic data. The results indicate that when the amount of synthetic data matches that of real data, the model achieves the best fairness. Additionally, using half the amount of real data results in the best accuracy. We believe that using too much synthetic data during fine-tuning can lead to overfitting, while using too little data may fail to adequately debias the model.

5. Conclusion

In this work, we proposed a method to mitigate bias in machine learning models using synthetic data generated by a text-to-image process. By designing contextual synthetic data generation and selective fine-tuning, we enhance model fairness without requiring demographic group annotations. Our model updates selectively fairness-sensitive parameters, optimizing simultaneously model fairness and utility scores. Empirical results demonstrate that our method outperforms existing techniques, improving fairness while maintaining model utility performance. This work highlights the potential of synthetic data for creating fairer AI systems, offering a promising direction for future research in bias mitigation.

6. Acknowledgments

This research utilised Queen Mary’s Apocrita HPC facility, supported by QMUL Research-IT. Zengqun Zhao is funded by Queen Mary Principal’s PhD Studentships. Zengqun Zhao wants to thank Yining Wang and James Oldfield for the valuable comments and help, and Yiming Lin and Jie Shen for the valuable discussions.

References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 1, 3
- [2] Carolyn Ashurst and Adrian Weller. Fairness without demographic data: A survey of approaches. In *EAAMO*, pages 1–12, 2023. 2
- [3] Junsong Chen, YU Jincheng, GE Chongjian, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *ICLR*, 2023. 2, 4
- [4] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. In *ICML*, pages 1887–1898, 2020. 2
- [5] Sam Corbett-Davies, Johann D Gaebler, Hamed Nilforoshan, Ravi Shroff, and Sharad Goel. The measure and mismeasure of fairness. *JMLR*, 24(1):14730–14846, 2023. 3
- [6] Elliot Creager, David Madras, Jörn-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. Flexibly fair representation learning by disentanglement. In *ICML*, pages 1436–1445, 2019. 2
- [7] Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. Environment inference for invariant learning. In *ICML*, pages 2189–2200, 2021. 6
- [8] Saloni Dash, Vineeth N Balasubramanian, and Amit Sharma. Evaluating and mitigating bias in image classifiers: A causal perspective using counterfactuals. In *WACV*, pages 915–924, 2022. 2, 3
- [9] Moreno D’Incà, Christos Tzelepis, Ioannis Patras, and Nicu Sebe. Improving fairness using vision-language driven image augmentation. In *WACV*, pages 4695–4704, 2024. 1, 2, 3, 6, 7
- [10] Michele Donini, Luca Oneto, Shai Ben-David, John S Shawe-Taylor, and Massimiliano Pontil. Empirical risk minimization under fairness constraints. *NeurIPS*, 31, 2018. 1, 3
- [11] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 2
- [12] John C Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, 2021. 1, 3
- [13] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *NeurIPS*, 29, 2016. 3, 6
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6, 7
- [15] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. In *ACL*, pages 328–339, 2018. 4
- [16] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, pages 1–13, 2022. 13
- [17] Jungseock Joo and Kimmo Kärkkäinen. Gender slopes: Counterfactual fairness for computer vision models by attribute manipulation. In *FATE/MM*, pages 1–5, 2020. 2, 3, 6
- [18] Sangwon Jung, Sanghyuk Chun, and Taesup Moon. Learning fair classifiers with partially annotated group labels. In *CVPR*, pages 10348–10357, 2022. 6
- [19] Gal Kaplun, Andrey Gurevich, Tal Swisa, Mazon David, Shai Shalev-Shwartz, and Eran Malach. Less is more: Selective layer finetuning with sub-tuning. *arXiv preprint arXiv:2302.06354*, 2023. 4, 7
- [20] Adarsh Kappiyath, Abhra Chaudhuri, AJAY KUMAR JAISWAL, Ziquan Liu, Yunpeng Li, Xiatian Zhu, and Lu Yin. Sebra: Debiasing through self-guided bias ranking. In *ICLR*, pages 1–12, 2025. 3
- [21] Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *CVPR*, pages 2661–2671, 2019. 4
- [22] Yoonho Lee, Annie S Chen, Fahim Tajwar, Ananya Kumar, Huaxiu Yao, Percy Liang, and Chelsea Finn. Surgical fine-tuning improves adaptation to distribution shifts. In *ICLR*, 2022. 4
- [23] Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. Large-scale methods for distributionally robust optimization. *NeurIPS*, 33:8847–8860, 2020. 1, 3, 6
- [24] Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *ICML*, pages 6781–6792, 2021. 3, 6, 8
- [25] Lydia T Liu, Max Simchowitz, and Moritz Hardt. The implicit fairness criterion of unconstrained learning. In *ICML*, pages 4051–4060. PMLR, 2019. 3
- [26] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 5, 12
- [27] Ziquan Liu, Yi Xu, Xiangyang Ji, and Antoni B Chan. Twins: A fine-tuning framework for improved transferability of adversarial robustness and generalization. In *CVPR*, pages 16436–16446, 2023. 4
- [28] Romain Lopez, Jeffrey Regier, Michael I Jordan, and Nir Yosef. Information constraints on auto-encoding variational bayes. *NeurIPS*, 31, 2018. 4
- [29] David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and transferable representations. In *ICML*, pages 3384–3393, 2018. 2

- [30] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6):1–35, 2021. 1
- [31] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *ICLR*, 2022. 4
- [32] Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: De-biasing classifier from biased classifier. *NeurIPS*, 33:20673–20684, 2020. 6
- [33] James Oldfield, Markos Georgopoulos, Grigorios Chrysos, Christos Tzelepis, Yannis Panagakis, Mihalys Nicolaou, Jiankang Deng, and Ioannis Patras. Multilinear mixture of experts: Scalable expert specialization through factorization. *NeurIPS*, 37:53022–53063, 2025. 3
- [34] Sungho Park, Sunhee Hwang, Dohyung Kim, and Hyeran Byun. Learning disentangled representation for fair facial attribute classification via fairness-aware information alignment. In *AAAI*, pages 2403–2411, 2021. 2
- [35] Sungho Park, Jewook Lee, Pilhyeon Lee, Sunhee Hwang, Dohyung Kim, and Hyeran Byun. Fair contrastive learning for facial attribute classification. In *CVPR*, pages 10389–10398, 2022. 5, 6
- [36] Momchil Peychev, Anian Ruoss, Mislav Balunović, Maximilian Baader, and Martin Vechev. Latent space smoothing for individually fair representations. In *ECCV*, pages 535–554, 2022. 2, 3, 4
- [37] Mohammad Pezeshki, Diane Bouchacourt, Mark Ibrahim, Nicolas Ballas, Pascal Vincent, and David Lopez-Paz. Discovering environments with xrm. In *ICML*, 2024. 3
- [38] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. *NeurIPS*, 30, 2017. 3
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 4
- [40] Vikram V Ramaswamy, Sunnie SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *CVPR*, pages 9301–9310, 2021. 2, 3
- [41] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, pages 1–27, 2022. 2, 4
- [42] Youngmin Ro and Jin Young Choi. Autolr: Layer-wise pruning and auto-tuning of learning rates in fine-tuning of deep networks. In *AAAI*, pages 2486–2494, 2021. 4
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 2, 3, 4
- [44] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 4
- [45] 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. In *NeurIPS*, pages 36479–36494, 2022. 2, 4
- [46] 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. *NeurIPS*, 35:36479–36494, 2022. 3
- [47] Joonghyuk Shin, Minguk Kang, and Jaesik Park. Fill-up: Balancing long-tailed data with generative models. *arXiv preprint arXiv:2306.07200*, 2023. 1, 2, 6, 7
- [48] Robik Shrestha, Yang Zou, Qiuyu Chen, Zhiheng Li, Yusheng Xie, and Siqi Deng. Fairrag: Fair human generation via fair retrieval augmentation. In *CVPR*, pages 11996–12005, 2024. 2, 4
- [49] Ilya Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. Ai models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759, 2024. 2
- [50] Krishnakant Singh, Thanush Navaratnam, Jannik Holmer, Simone Schaub-Meyer, and Stefan Roth. Is synthetic data all we need? benchmarking the robustness of models trained with synthetic images. In *CVPR*, pages 2505–2515, 2024. 5
- [51] Tobias Uelwer, Jan Robine, Stefan Sylvius Wagner, Marc Höftmann, Eric Upschulte, Sebastian Konietzny, Maike Behrendt, and Stefan Harmeling. A survey on self-supervised representation learning. *arXiv preprint arXiv:2308.11455*, 2023. 2
- [52] Mei Wang and Weihong Deng. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *CVPR*, pages 9322–9331, 2020. 6
- [53] Yinong Oliver Wang, Younjoon Chung, Chen Henry Wu, and Fernando De la Torre. Domain gap embeddings for generative dataset augmentation. In *CVPR*, pages 28684–28694, 2024. 5
- [54] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *CVPR*, pages 8919–8928, 2020. 2
- [55] Yuyang Xue, Junyu Yan, Raman Dutt, Fasih Haider, Jing-shuai Liu, Steven McDonagh, and Sotirios A Tsaftaris. Bmft: Achieving fairness via bias-based weight masking fine-tuning. In *MICCAI Workshop on Fairness of AI in Medical Imaging*, pages 98–108, 2024. 3
- [56] Moon Ye-Bin, Nam Hyeon-Woo, Wonseok Choi, Nayeong Kim, Suha Kwak, and Tae-Hyun Oh. Exploiting synthetic data for data imbalance problems: Baselines from a data perspective. *arXiv preprint arXiv:2308.00994*, 2023. 1, 4, 6, 7
- [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 *ICCV*, pages 3969–3980, 2023. 2, 4
- [58] Fengda Zhang, Kun Kuang, Long Chen, Yuxuan Liu, Chao Wu, and Jun Xiao. Fairness-aware contrastive learning with partially annotated sensitive attributes. In *ICLR*, 2023. 3, 5, 6

- [59] Fengda Zhang, Qianpei He, Kun Kuang, Jiashuo Liu, Long Chen, Chao Wu, Jun Xiao, and Hanwang Zhang. Distributionally generative augmentation for fair facial attribute classification. In *CVPR*, pages 22797–22808, 2024. [2](#), [3](#), [4](#), [5](#), [6](#)
- [60] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, 2017. [6](#), [12](#)
- [61] Qihao Zhao, Yalun Dai, Hao Li, Wei Hu, Fan Zhang, and Jun Liu. Ltgc: Long-tail recognition via leveraging llms-driven generated content. In *CVPR*, pages 19510–19520, 2024. [13](#)
- [62] Zengqun Zhao and Ioannis Patras. Prompting visual-language models for dynamic facial expression recognition. In *BMVC*, pages 1–14, 2023. [4](#)
- [63] Xiao Zhou, Yong Lin, Renjie Pi, Weizhong Zhang, Renzhe Xu, Peng Cui, and Tong Zhang. Model agnostic sample reweighting for out-of-distribution learning. In *ICML*, pages 27203–27221, 2022. [6](#)
- [64] Yongchao Zhou, Hshmat Sahak, and Jimmy Ba. Using synthetic data for data augmentation to improve classification accuracy. In *ICML*, 2023. [2](#), [5](#)
- [65] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020. [4](#)