

# Constructing Fair Latent Space for Intersection of Fairness and Explainability

Hyungjun Joo<sup>1,2</sup>, Hyeonggeun Han<sup>1,2</sup>, Sehwan Kim<sup>1</sup>, Sangwoo Hong<sup>1,2</sup>, Jungwoo Lee<sup>1,2,3\*</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, Seoul National University

<sup>2</sup>NextQuantum, Seoul National University

<sup>3</sup>HodooAI Labs

{joojh911, hygnhan, sehwan kim, tkddn0606, junglee}@snu.ac.kr

## Abstract

As the use of machine learning models has increased, numerous studies have aimed to enhance fairness. However, research on the intersection of fairness and explainability remains insufficient, leading to potential issues in gaining the trust of actual users. Here, we propose a novel module that constructs a fair latent space, enabling faithful explanation while ensuring fairness. The fair latent space is constructed by disentangling and redistributing labels and sensitive attributes, allowing the generation of counterfactual explanations for each type of information. Our module is attached to a pretrained generative model, transforming its biased latent space into a fair latent space. Additionally, since only the module needs to be trained, there are advantages in terms of time and cost savings, without the need to train the entire generative model. We validate the fair latent space with various fairness metrics and demonstrate that our approach can effectively provide explanations for biased decisions and assurances of fairness.

## 1 Introduction

With the rapid advancement of machine learning models, the demand for fairness has grown, especially in protecting sensitive features like age and gender from bias (Caton and Haas 2024). Although recent research has incorporated fairness metrics to ensure fairness, the sufficiency of these metrics alone in addressing stakeholders' concerns remains an open question (Fig. 1A). In the US, predictive policing tools have been criticized for disproportionately targeting individuals based on race and gender biases, sparking protests (Richardson, Schultz, and Crawford 2019). In this context, it is crucial for stakeholders to identify whether the decision is free from gender bias. Thus, presenting compelling evidence to stakeholders and practitioners becomes essential when addressing sensitive issues. Practitioners can adjust the model based on this evidence, while stakeholders can transparently rely on and utilize the model's decisions, grounded in the provided evidence (Fig. 1C-blue). In addition, when the model has been adjusted, providing assurances through explanations that decisions are not based on sensitive attributes (Fig. 1C-red) is essential for building trust in decision-making systems (Jacovi et al. 2021).

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

There has been research at the intersection of fairness and explainability, which attempts to explain the causes of unfairness through decomposing model disparities via features (Begley et al. 2020) or causal paths (Chiappa 2019; Plecko and Bareinboim 2023). However, these studies have several limitations. Firstly, they provide explanations for the overall fairness of the model, which are not sufficiently persuasive about the fairness of individual predictions. Secondly, although there are methods to explain the model's decisions (Qiang et al. 2022; Dash, Balasubramanian, and Sharma 2022), they do not fully reveal the underlying reasons because the explanations are based on variations in the model's outputs rather than the model's internal behavior.

To address the aforementioned limitations and provide a faithful explanation of decisions to each individual, we propose a novel framework that functions as a module for a generative model to construct a fair latent space. In the fair latent space, where sensitive attributes are disentangled from decision-making factors, we can elucidate why a predictor returns a specific outcome and ensure that this outcome is not influenced by sensitive attributes. This is achieved by manipulating these distinct attributes within the latent space and generating counterfactuals (Fig. 1B). As these counterfactual explanations (simulating alternative inputs with specific changes to the original) are generated directly by the model, they provide trustworthy explanations due to their inherent interpretability (Rudin 2019; Joo et al. 2023).

To construct fair latent space, we conduct the training of this module on the results of our theoretical exploration of mutual information. The fundamental concept involves adjusting latent representations to effectively disentangle and redistribute information associated with labels and sensitive attributes. By disentangling and optimizing information related to attributes, our method not only provides explanations but also consistently demonstrates a high level of fairness across various metrics. In addition, while recent generative models perform well, they entail significant training time and costs. Therefore, to reduce computational costs and training time, we propose constructing a fair latent space by training an invertible neural network (INN) on the latent space of pre-trained generative models, instead of retraining them from scratch (Fig. 1B). This approach leads to the development of a versatile module applicable to the generative model, which offers counterfactual reasoning.

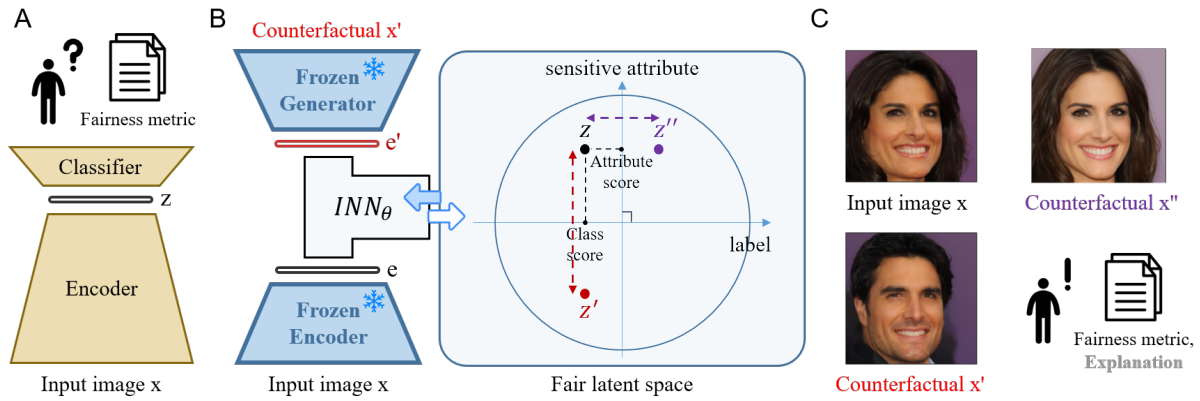


Figure 1: (A) Models aimed at enhancing fairness without any explanation. (B) The proposed model trains an invertible neural network based on a pre-trained generative model to construct a fair latent space where the information of labels and sensitive attributes is disentangled into separate dimensions. The Y-axis corresponds to the dimension of the sensitive attribute, while the X-axis corresponds to the dimension of the label. (C) Counterfactual explanations can be generated by adjusting values in the opposite direction within a fair latent space. Using an INN and a frozen generator,  $x'$  and  $x''$  are generated from  $z'$  and  $z''$ .

## 2 Related Work

**Fairness in machine learning** Fairness in machine learning has gained focus in recent research. One approach involves in-processing methods that incorporate fairness-aware regularization into the training objective (Donini et al. 2018; Sagawa\* et al. 2020; Han et al. 2024) or soften fairness constraints into score-based constraints for optimization (Zafar et al. 2017, 2019; Zhu et al. 2023). However, with research identifying the underlying biases in training data as the primary source of unfairness (Barocas and Selbst 2016; Tommasi et al. 2017; Gustafson et al. 2023), there have been notable efforts on direct interventions through augmentation to tackle this issue (Ramawamy, Kim, and Russakovsky 2021; Qiang et al. 2022; Zietlow et al. 2022; Zhang et al. 2024). Consequently, studies have investigated fairness evaluation using counterfactual samples generated by generative models (Denton et al. 2019; Joo and Kärkkäinen 2020; Dash, Balasubramanian, and Sharma 2022). However, existing counterfactual generation methods differ from ours in that they conduct an analysis by examining how the classifier handles counterfactual samples, rather than providing counterfactual explanations.

**Fair representation learning** As fairness becomes a crucial issue in the practical application of models, various approaches have been developed to learn fair representations through disentangling. Kim and Mnih (2018) propose a direct method that uses disentanglement learning, and Creager et al. (2019) isolate the latent space into sensitive and non-sensitive components. Other approaches enforce independence by learning target and sensitive codes to follow orthogonal priors (Sarhan et al. 2020) or by minimizing distance covariance, offering a non-adversarial alternative (Liu et al. 2022). However, these methodologies all rely on variational autoencoders, which are inherently limited in terms of generalization. Furthermore, because the entire generative model is trained, image reconstruction quality is com-

promised (Shao et al. 2022).

Research has also been conducted employing contrastive learning methods, which have proven highly successful in learning effective representations in recent years (Chen et al. 2020; Khosla et al. 2020). Similar to our approach, these methods learn representations by reducing the distance between positive samples and increasing it for negative samples. FSCL (Park et al. 2022) focuses on regulating the similarity between groups, ensuring it remains unaffected by sensitive attributes. Meanwhile, other approaches concentrate on enhancing the robustness of representation alignment (Zhang et al. 2022) or devising methods effective even with partially annotated sensitive attributes (Zhang et al. 2023). However, methods developed after FSCL assume unlabeled scenarios; as a result, FSCL has demonstrated the best performance in labeled situations.

## 3 Method

In this paper, our objective is to achieve a fair latent space in generative models, thereby gaining insights into the model's fairness through counterfactual explanations and fair classification. Therefore, we first conduct a theoretical analysis on separating information about labels and sensitive attributes in the latent space. Secondly, we connect this theoretical analysis to practical training methods.

### 3.1 Disentangling sensitive attributes from labels

Previously, many methods concentrated on being invariant to information associated with sensitive attributes, to depend solely on labels for fairness in classification (Kehrenberg et al. 2020; Ramawamy, Kim, and Russakovsky 2021; Park et al. 2022). On the other hand, our approach does not aim to exclude but rather to separate information regarding the sensitive attributes. We aim to disentangle data related to sensitive attributes from labels and assign them to distinct dimensions. By redistributing this data in the latent space,

we enrich the representation with label information, thus enhancing the fairness of classification. Therefore, our goal is to maximize the information associated with each assigned attribute within its respective dimension.

**Lens of the information bottleneck** Let  $S$  denote a sensitive attribute, such as race or gender, and  $Y$  a label. With an invertible network  $f_\theta$  and pre-trained generative model  $G = f_{dec} \circ f_{enc}$ , where  $f_{enc}$  is the encoder and  $f_{dec}$  is the decoder, the latent representation of image data  $X$  can be obtained as  $E = f_{enc}(X)$ . Concurrently, the latent representation of the invertible network is derived as  $Z = f_\theta(E)$ . Then, we allocate information pertaining to labels and attributes in separate dimensions, denoted as  $Z^Y$  and  $Z^S$ .

In the case of the label dimension, the objective is to maximize its information using the compact invertible model. This aligns with the objective of the Information Bottleneck (IB) (Tishby and Zaslavsky 2015), which aims to maximize the information between the representation and the target in situations where the model’s complexity is limited, as our training takes place in a compact invertible model. With the lens of the IB principle, we maximize the mutual information  $I(Z^Y, Y)$  subject to a complexity constraint specified as  $I(Z^Y, E) < b$  with a constant  $b$ . In this scenario, we can express our objective using the loss function  $L_{IB} = I(Z^Y, E) - \beta I(Z^Y, Y)$  with  $\beta > 1$ , as we focus more on the relationship between  $Y$  and  $Z^Y$ . Furthermore, substituting mutual information with entropy allows us to reformulate the loss function using the determinant of the covariance matrix  $C$  (Ahmed and Gokhale 1989), facilitating the transformation as follows.

**Theorem 1.** Let the representation  $Z^Y$  follow a Gaussian distribution, and  $\beta > 1$ . The information bottleneck-based loss  $L_{IB} = I(Z^Y, E) - \beta I(Z^Y, Y)$  can be reformulated as:

$$L_{IB} = \mathbb{E}_Y [\log \det(C_{Z^Y|Y})] - \lambda \log \det(C_{Z^Y}), \quad \lambda > 0. \quad (1)$$

Additional details for the proof are provided in the Appendix. To minimize  $L_{IB}$ , we focus on maximizing the second term  $\log \det(C_{Z^Y})$ , given that the first term remains constant when optimizing the representation  $Z^Y$ . If we set the dimension of fair representation as  $d$  and apply Jensen’s inequality, the second term  $\log \det(C_{Z^Y})$  can be rewritten as:

$$\sum_{i=1}^d \log(\lambda_i(C_{Z^Y})) \leq d \log\left(\frac{1}{d} \sum_{i=1}^d \lambda_i(C_{Z^Y})\right), \quad (2)$$

where  $\lambda_i(C_{Z^Y})$  is  $i$ -th eigenvalue of  $C_{Z^Y}$ . In Jensen’s inequality, equality holds when all values are equal as  $\lambda_i(C_{Z^Y}) = \lambda_j(C_{Z^Y})$  for  $\forall j \in [1, \dots, d]$ . Given that the covariance matrix is symmetric and positive semi-definite, it allows for diagonalization of  $C_{Z^Y}$  with maximum determinant through an orthogonal matrix  $Q$  (i.e.  $\det(Q) = \pm 1$ ), resulting in  $diag(c, c, \dots, c)$  for  $c = \lambda_i(C_{Z^Y})$ . Therefore, our objective is achieved when the covariance matrix is a diagonal matrix with identical diagonal entries.

**Connection between opposite sensitive attributes** In addition to maximizing mutual information along separate di-

mensions, our strategy involves directly mitigating the influence of sensitive attributes on decision-making processes. This is accomplished by ensuring that the label dimension corresponding to the label contains solely relevant information. Therefore, our additional goal is to maximize the mutual information between inputs with different sensitive attributes within the label’s dimension.

Considering a scenario with a binary sensitive attribute, for data with the same label  $y$ , we denote the data with a positive sensitive attribute as  $X_{s_0}^y$  and the data with a negative sensitive attribute as  $X_{s_1}^y$ . However, directly computing mutual information between two random variables is infeasible. To address this, we employ a widely-used approach that approximates the mutual information using noise-contrastive estimation (Oord, Li, and Vinyals 2018; Poole et al. 2019). Given the two random variables  $X_{s_0}^y$  and  $X_{s_1}^y$ , the mutual information lower bound is defined as follows.

$$I_{NCE} = \mathbb{E} \left[ \frac{1}{K} \sum_{i=1}^K \log \frac{e^{g(x_{s_0,i}^y, x_{s_1,i}^y)}}{\sum_{j=1}^K e^{g(x_{s_0,i}^y, x_{s_1,j}^y)}} \right] + \log(K), \quad (3)$$

where the expectation is over  $K$  independent samples.

To maximize the mutual information  $I(X_{s_0}^y, X_{s_1}^y)$ , our strategy necessitates numerically increasing  $g(x_{s_0,i}^y, x_{s_1,i}^y)$  while concurrently reducing  $g(x_{s_0,i}^y, x_{s_1,j}^y)$ . In our scenario, the representation vectors are derived from a combination of a pre-trained generative model and an invertible model. Consequently, the product between encoded representations  $g(x_0, x_1)$  can be formally denoted as  $z_0^T z_1 = f_\theta(f_{enc}(x_0))^T f_\theta(f_{enc}(x_1))$  at the dimension  $Z^Y$ .

As a result, the term to be increased can be denoted as  $(z_{s_0,i}^y)^T z_{s_1,i}^y$ , and the term to be decreased can be expressed as  $(z_{s_0,i}^y)^T z_{s_1,j}^y$ . Consequently, with the objective of transforming the covariance matrix into a scalar matrix, the term to be increased is represented as a negative squared  $L_2$  distance, whereas the term to be decreased is upper bounded by the  $L_2$  distance, in accordance with the Cauchy-Schwarz inequality. Detailed proofs are included in the Appendix, and these findings lead to the following theorem.

**Theorem 2.** Let the mutual information  $I(Z^Y, Y)$  be maximized within a network of constrained capacity. Then, maximizing the mutual information  $I(X_{s_0}^y, X_{s_1}^y)$  can be achieved by minimizing the  $L_2$  distance between samples from the groups  $X_{s_0}^y$  and  $X_{s_1}^y$ .

We can confirm that in our scenario, reducing the  $L_2$  distance enhances the mutual information between the two groups. This shift leads to an emphasis on accurate information rather than spurious attributes.

### 3.2 Constructing fair latent space

In this section, we elucidate how the theoretical analysis in Sec. 3.1 seamlessly translates into the proposed methodology, resulting in the losses illustrated in Fig. 2. Our approach involves training an invertible network  $f_\theta$ , thereby enabling the acquisition of a fair representation without incurring additional costs for a generative model.




Figure 2: Overview of our approach connecting theoretical analysis to practical implementation, comprising three main components. The distance loss  $L_{di}$  regulates distances to respond specifically to attributes. Furthermore, the diagonalizing loss  $L_{dg}$  and equalizing loss  $L_{eq}$  transform the covariance matrix into an identical diagonal matrix.

One of our objectives is to convert the covariance matrix of the representation into a scalar matrix, akin to dimension-contrastive methods in semi-supervised learning (Zbontar et al. 2021; Bardes, Ponce, and LeCun 2022), which decorrelate each dimension. Let  $B = \{X, Y, S\} = \{(x_i, y_i, s_i)\}_{i=1}^n$  represent the training batch, comprising images  $x_i$ , labels  $y_i$ , and sensitive attributes  $s_i$ . Subsequently, the representation corresponding to the label in the training batch, which is derived through the invertible network can be denoted as  $Z^Y = f_\theta(f_{enc}(X))^Y \in \mathbb{R}^{n \times d_y}$ . Then, the covariance matrix for each latent dimension of this representation can be defined as follows.

$$C(Z^Y) = \mathbb{E} [(Z^Y - \mathbb{E}(Z^Y))^T (Z^Y - \mathbb{E}(Z^Y))] \in \mathbb{R}^{d_y \times d_y} \quad (4)$$

To diagonalize the obtained covariance matrix, we subtract the term that is element-wise multiplied ( $\odot$ ) by the identity matrix  $I_{d_y}$ , isolating the off-diagonal elements. This process yields a diagonalizing loss term by summing the squared elements of the obtained matrix and incorporating a normalization factor of  $1/d_y$ .

$$L_{dg}(Z^Y) = \frac{1}{d_y} \|C(Z^Y) - C(Z^Y) \odot I_{d_y}\|_2^2. \quad (5)$$

Next, we further employ a loss function to equalize the diagonal elements. Instead of directly setting the diagonal elements to  $c$ , this approach is inspired by the observed performance improvements achieved by adjusting the scale through variance (Bardes, Ponce, and LeCun 2022). Given the covariance matrix  $C(Z^Y) \in \mathbb{R}^{d_y \times d_y}$ , we constrain the variance of the batch to  $c$  for each dimension. Utilizing the ReLU activation, which computes  $\max(0, x)$  for the input  $x$ , the loss is defined as follows.

$$L_{eq}(Z^Y) = \frac{1}{d_y} \sum_{j=1}^{d_y} \max(0, c - \sqrt{\text{Var}(z_{:,j}^Y)} + \epsilon), \quad (6)$$

where  $z_{:,j}^Y \in \mathbb{R}^n$  denotes the vector of the  $j$ -th dimension in the latent dimension of  $Z^Y$ ,  $\epsilon$  is a small constant to prevent gradient collapse, and  $\text{Var}(\cdot)$  represents the variance.

Another objective is to maximize the mutual information between groups with different values in unintended factors. From Thm. 2, our method minimizes  $L_2(x_{s_0}^y, x_{s_1}^y)$ , which is a feasible strategy for maximizing  $I(X_{s_0}^y, X_{s_1}^y)$  when focusing on  $Z^Y$ . Additionally, to enhance the influence of the intended factor, we incorporate a term into our loss function that increases the distance  $L_2(x_s^y, x_{s'}^y)$  between groups with the same sensitive attribute but different labels, where  $y \neq y'$ . This results in the following formulation.

$$L_{di}(Z^Y) = - \frac{1}{\sum M_{max}} \sum_{i=1}^n \sum_{j \neq i}^n M_{max} D(z_i^Y, z_j^Y) + \frac{1}{\sum M_{min}} \sum_{i=1}^n \sum_{j \neq i}^n M_{min} D(z_i^Y, z_j^Y), \quad (7)$$

where  $M_{max} = \delta(y_i, y_j)(1 - \delta(s_i, s_j))$  is the mask for selecting samples to maximize, and  $M_{min} = \delta(s_i, s_j)(1 - \delta(y_i, y_j))$  selects samples to minimize. Here,  $\delta(x, y)$  is the Kronecker delta, which equals 1 if  $x = y$  and 0 otherwise. The loss function uses  $D(x, y) = \log((\|x - y\|_2^2 + 1)/(\|x - y\|_2^2 + \epsilon))$ , a monotonically decreasing function with respect to the distance, instead of the  $L_2$  distance. This modification mitigates potential instability caused by unbounded values as distances grow infinitely. By employing a bounded function, we enhance training stability and focus on regions where the  $L_2$  distance is minimized.

The loss function for constructing a fair latent space in training the invertible network is derived by integrating these losses. Drawing inspiration from the technique of segregating information into distinct dimensions (Esser, Rombach, and Ommer 2020), we decompose the latent dimensions in the representation from the invertible network as  $Z = [Z^Y, Z^S] \in \mathbb{R}^d$  where  $Z^Y \in \mathbb{R}^{d_y}$  and  $Z^S \in \mathbb{R}^{d_s}$ . The loss function, based on theoretical justification and designed to construct a fair latent space, is defined as follows.

$$L_{fair}(Z^Y) = \lambda_{dg} L_{dg}(Z^Y) + \lambda_{eq} L_{eq}(Z^Y) + \lambda_{di} L_{di}(Z^Y). \quad (8)$$

Method	Y = a, S = m				Y = yo, S = m				Y = b, S = m			
	EO	DP	WGA	Acc	EO	DP	WGA	Acc	EO	DP	WGA	Acc
DiffAE	33.4	51.2	52.9	78.1	25.8	26.3	22.8	83.5	18.4	15.2	24.2	89.4
Ours	<b>5.9</b>	<b>25.5</b>	<b>70.4</b>	75.2	<b>3.4</b>	<b>13.2</b>	<b>73.1</b>	74.6	<b>1.6</b>	<b>6.7</b>	<b>76.0</b>	78.2
SimCLR	26.3	46.2	60.8	79.7	16.2	22.0	42.0	84.7	16.6	16.0	37.5	89.8
SupCon	28.0	47.9	60.3	79.9	20.0	23.2	32.2	85.1	16.2	14.6	32.3	90.4
FSQL	14.3	35.0	67.5	78.1	12.9	18.0	51.2	83.7	12.2	14.4	44.0	89.3

Table 1: Evaluation of the constructed latent space obtained with an invertible neural network in the CelebA. We measure EO and DP (the lower the better) and WGA (the higher the better) and average accuracy.  $a$ ,  $yo$ ,  $b$ , and  $m$  account for *attractive*, *young*, *bushy brows*, and *male*.

Method	CelebAHQ: Y = a, S = yo				UTK Face: Y = m, S = yo				CelebA: Y = a, S = m&yo			
	EO	DP	WGA	Acc	EO	DP	WGA	Acc	EO	DP	WGA	Acc
DiffAE	28.3	56.2	61.6	82.1	17.4	18.2	77.1	88.3	51.7	73.6	32.7	78.2
Ours	<b>13.2</b>	<b>41.1</b>	<b>68.4</b>	77.0	<b>8.5</b>	<b>9.3</b>	82.5	87.0	<b>17.4</b>	<b>45.9</b>	<b>62.3</b>	73.8
SimCLR	26.3	56.0	63.6	82.5	13.7	17.4	80.4	90.5	50.3	73.3	33.1	79.6
SupCon	24.8	55.0	64.6	82.7	13.0	16.6	<b>82.6</b>	90.7	49.5	72.1	30.6	79.9
FSQL	20.1	50.4	62.2	81.7	10.6	14.4	76.9	90.2	-	-	-	-

Table 2: Evaluation of the constructed latent space obtained with an invertible neural network in the various settings. We measure EO, DP, WGA, and average accuracy. The maximum values of EO and DP are reported across 4 groups defined by two sensitive attributes at CelebA.

### 3.3 Gaussianizing embeddings through INN

We train an invertible neural network (INN) instead of a convolutional network to ensure that transformations in the fair latent space can appropriately extend to the latent space of the generative model. In our method, the INN  $f_\theta$  maps generative representation  $e$  to a fair representation  $z$  with a forward mapping  $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$  and can also serve as an inverse mapping  $f_\theta^{-1} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ .

We use Normalizing Flows (NFs) for exact log-likelihood computation and precise inference. However, they often produce severe artifacts due to exploding inverses when generating images from out-of-distribution data (Behrmann et al. 2021; Hong, Park, and Chun 2023). To bypass this, we train NFs on high-level semantic representation (Kirichenko, Izmailov, and Wilson 2020). Our model computes the base distribution  $p_Z(z)$  with an encoder distribution  $p_E(e)$  in the latent space. Given the assumption in Thm. 1 that  $Z$  follows a Gaussian distribution, a standard Gaussian as the  $p_Z(z)$  completes our theoretical framework. We minimize the negative log-likelihood objective using the following equation, where the Jacobian of  $f_\theta$  is denoted as  $J_{f_\theta}$ .

$$L_g = -\frac{1}{n} \sum_{i=1}^n \left( \|f_\theta(e)\|^2 + \log |\det(J_{f_\theta}(e))| \right) \quad (9)$$

Furthermore, we train a one-layer fully connected classifier using  $Z^Y$  for labels  $Y$  and  $Z^S$  for sensitive attributes  $S$ , employing a cross-entropy loss denoted by  $L_{cls}$ . During training, the overall loss is derived by summing  $L_{fair}(Z^Y)$ ,  $L_{fair}(Z^S)$ ,  $L_g$ , and  $L_{cls}$ .

## 4 Experiment

We conducted experiments with the Diffusion Autoencoder (DiffAE) (Preechakul et al. 2022), a recent encoder-decoder

structure generative model, to evaluate our proposed approach and demonstrate its practical applicability. Additionally, we employed Glow (Kingma and Dhariwal 2018) as the invertible neural network.

### 4.1 Experimental details

**Fairness metrics** The most commonly used metrics for measuring group fairness are Demographic Parity (DP) (Chouldechova 2017) and Equalized Odds (EO) (Hardt, Price, and Srebro 2016). **DP** aims to equalize the rate of positive outcomes irrespective of the sensitive attribute. **EO** aims to equalize the true positive rate and false positive rate, which is appropriate for problems where negative outcomes are as important as positive outcomes, such as facial attribute classification. **Worst-Group-Accuracy (WGA)** has also been used as a fairness metric; Zietlow et al. (2022) argues that reaching fairness by performance degradation not only for the highest-scoring group but also for the lowest-scoring group is received substantial criticism in other areas (Brown 2003; Christiano and Braynen 2008; Doran 2001). Following Zhang et al. (2023), we defined EO and DP as follows.

$$EO = \sum_y \left| \mathbb{P}_{s^1}(\hat{Y} = y | Y = y) - \mathbb{P}_{s^0}(\hat{Y} = y | Y = y) \right| \quad (10)$$

$$DP = \left| \mathbb{P}_{s^1}(\hat{Y} = y_p) - \mathbb{P}_{s^0}(\hat{Y} = y_p) \right|, \quad y_p = \text{positive} \quad (11)$$

**Datasets** We conduct experiments on DiffAE using three datasets. In CelebA (Liu et al. 2015), which features 40 binary attribute labels, we designate  $S = \{male\}$  as the sensitive attribute and classify gender-dependent attributes (Ramaswamy, Kim, and Russakovsky 2021) such as  $Y = \{attractive, young, bushy brows\}$ . In CelebAHQ (Karras

Method	CelebA: $Y = a, S = m$			
	EO	DP	WGA	Acc
DiffAE	33.4	51.2	52.9	78.1
+INN	26.5	46.3	60.9	79.4
+ $L_{dg,eq}$	17.0	37.6	66.5	78.5
+ $L_{di}$	14.8	35.6	67.3	78.5
+ $L_g$ (Ours)	<b>5.9</b>	<b>25.5</b>	<b>70.4</b>	75.2

Table 3: Ablation study on the components of  $L_{fair}$ .

et al. 2018), a high-resolution face image dataset, we classify  $Y = \{attractive\}$  while setting  $S = \{young\}$  to verify applicability at high resolution. For UTK Face (Zhang, Song, and Qi 2017), which includes annotations for gender, age, and ethnicity, we establish the binary  $S = \{young\}$  based on an age threshold of 35 and conduct classification on  $Y = \{male\}$  following previous work (Zhang et al. 2023).

## 4.2 Fair latent space evaluation

We evaluated the fairness of the latent space by using fairness metrics. As one of our main objectives is to establish a fair latent space, we compared our approach with methods proposed for learning visual representations, such as SimCLR (Chen et al. 2020), SupCon (Khosla et al. 2020), and FSCL (Park et al. 2022). While there have been studies aiming to train fair representations even without the notion of fairness, to our knowledge, FSCL demonstrates state-of-the-art performance in terms of the group fairness metric for facial attribute classification.

Gender discrimination is one of the most important topics addressed in fairness. To address this issue within facial attribute classification, which is used in a wide range of practical applications including face verification and image search, we have designated gender as a sensitive attribute. We conducted experiments on the CelebA dataset, focusing on three attributes known to be related to gender (Ramswamy, Kim, and Russakovsky 2021).

The classification results are shown in Tab. 1. Firstly, our proposed method demonstrates significant performance improvements across all three metrics: EO, DP, and WGA. Through our proposed method, we observed a significant reduction of 86.8% in EO, 52.0% in DP, and a 39.9% increase in WGA, indicating a successful transition to a fair latent space from a previously biased one. Additionally, our method displayed a clear distinction from FSCL, which also aimed for fair representation.

Secondly, Tab. 2 demonstrates that our method extends beyond a single dataset, proving effective on large-scale image datasets like CelebAHQ. Its efficacy is further validated through experiments on the UTKFace dataset, assessing performance across diverse datasets. Additionally, experiments on the CelebA dataset, which includes two sensitive attributes, evaluate the method’s performance in scenarios involving multiple sensitive attributes.

## 4.3 Ablation study

In this section, we assess the alignment between our theoretical analysis and practical outcomes through an ablation

	INN				CelebA: $Y = a, S = m$			
	$De$	$L_{dg}$	$L_{eq}$	$L_{di}$	EO	DP	WGA	Acc
	-	-	-	-	26.5	46.3	60.9	79.4
	-	✓	-	-	25.6	45.5	61.1	79.6
	✓	-	-	-	21.6	40.4	60.8	79.5
	✓	✓	-	-	19.5	39.5	64.5	78.3
	✓	✓	✓	-	17.0	37.6	66.5	78.5
	✓	-	-	✓	32.0	50.0	55.5	78.0
	✓	✓	-	✓	25.8	40.3	54.0	72.2
	✓	✓	✓	✓	14.8	35.6	67.3	78.5

Table 4: Ablation study to confirm the necessity of the theoretical assumptions. We denote decomposition as  $De$ .

study. Our method’s theoretical design is developed incrementally, with each stage building upon the previous one. As a result, performance improves with the addition of each component, as shown in Tab. 3, which confirms the effectiveness of the approaches discussed in Sec. 3.1 for ensuring fairness within the latent space. Additional results for other datasets and settings are provided in the Appendix.

We further confirm the necessity of assumptions underlying each theoretical analysis in Sec. 3.1. Initially, we assess the necessity of decomposing ( $De$ ) representation dimensions into labels and sensitive attributes, by comparing scenarios in which the covariance matrix is diagonalized without such decomposition. As shown in Tab. 4, applying  $L_{dg}$  without  $De$  leads to the maximization of entangled information, which has a minimal impact on fairness. Furthermore, to ensure that information is optimally structured when transforming the covariance matrix into a scalar matrix, we compare the scenarios involving  $L_{dg}$  and  $L_{dg} + L_{eq}$  with the application of  $De$ . The second set of rows in Tab. 4 emphasizes the importance of transforming the diagonal matrix to ensure identical diagonal entries. Finally, we investigate the importance of maximizing mutual information between the label  $Y$  and the representation  $Z^Y$ , as assumed in Thm. 2. The third set of rows in Tab. 4 shows that without this assumption,  $L_{di}$  fails to perform accurately.

## 4.4 Explaining the fairness by counterfactual

In the previous section, we observed an increase in the worst group’s accuracy after applying our framework. In this section, we demonstrate how this improvement is reflected in the explanations provided by our framework. With the constructed fair latent space, our method can generate counterfactual explanations by adjusting the dimensions corresponding to the label or sensitive attribute and observing the model’s behavior. Specifically, the vector of each classifier weight  $h$  in the latent space is the best choice for the basis corresponding to the label or sensitive attribute, as these classifiers most clearly represent how the model differentiates information. Hence, we base our counterfactual explanation on the representations obtained by moving the representation  $z = f_{\theta}(f_{enc}(x))$  in the direction of the classifier weight vector  $\hat{h} = \frac{h}{\|h\|}$ .

A counterfactual explanation of the label reveals the factors that influenced the model’s decision, while a counterfac-

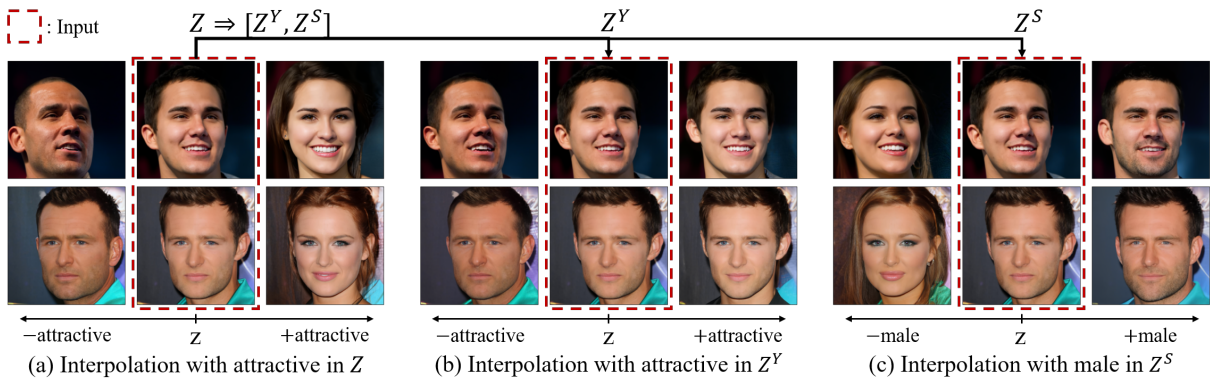


Figure 3: Counterfactual explanations with samples initially misclassified as unattractive by the original model. The x-axis indicates changes in the latent space based on the direction of classifier  $\hat{h}$ . In the original model, (a) counterfactuals of attractiveness reveal a clear correlation with gender. After constructing a fair latent space by isolating  $Z^S = \text{male}$ , we can observe (b) counterfactuals of attractiveness that exhibit no gender bias, and (c) counterfactuals across genders with equal attractiveness.

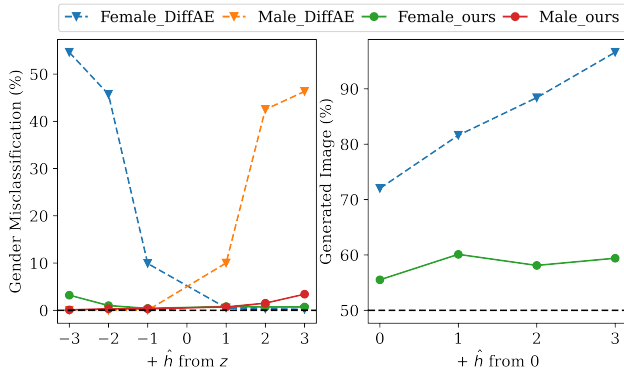


Figure 4: (Left) Gender misclassification rates when representations obtained from the CelebA HQ test dataset are shifted along the unit vector of the *attractive* classifier. (Right) Gender distribution after generating 1,000 images by shifting the mean of a standard Gaussian distribution along the unit vector of the *attractive* classifier.

tual explanation of the sensitive attribute shows whether the model would make the same decision under different sensitive attributes. Consider the *attractive* classification problem. If the model misclassifies data labeled as *attractive*, practitioners would want to identify the factors causing the misclassification. As shown in Fig. 3(a), practitioners can use label-based counterfactual generation to discern that the gender factor predominantly influences the determination of attractiveness. By designating *male* as the sensitive attribute and using our method, practitioners can exclude gender information when determining attractiveness. Indeed, when gender information is excluded, inputs that were previously misclassified are accurately classified, as evidenced in Fig. 3(b). With this explanation, stakeholders can verify the independence of the gender factor in decision. Furthermore, stakeholders can confirm the excluded information assigned to dimension  $Z^S$ , which represents inputs classified as at-

tractive with the same scores, as shown in Fig. 3(c).

To evaluate the explanation quantitatively, we conducted gender classification on the generated images using the CLIP model (ViT/B-32) with two classifier prompts: ‘photo of a male, man, or boy’ and ‘photo of a female, woman, or girl’, following previous works (Cho, Zala, and Bansal 2023; Shrestha et al. 2024). Please refer to the Appendix for more details about the experimental setup. When we transformed the representations obtained from the CelebA HQ test dataset according to the unit vector of attractiveness-classifier  $\hat{h}$ , the original model exhibited a clear correlation as shown in Fig. 4(L): as the representation became more attractive, males were increasingly misclassified, whereas females were similarly misclassified as it became less attractive. In contrast, with the fair latent space by our proposed method, we observed that this correlation between attractiveness and gender was nearly eliminated.

Additionally, to verify whether the obtained fair latent space maintains fairness during the image generation process, we gradually transformed the representation from the origin according to  $\hat{h}$ , adding noise from a standard Gaussian distribution to generate 1,000 images, and then evaluated their gender ratio. As shown in Fig. 4(R), the original model exhibited a slope of 8.06 in a linear regression, indicating a consistent correlation, while our method showed a slope of 0.97, demonstrating that it is possible to enhance attractiveness while maintaining the gender ratio.

## 5 Conclusions & Limitations

We propose a module that enhances the understanding of the model’s fairness by offering explanations to individuals, accompanied by fair decisions supported through a fair latent space. However, since our model involves integrating a module with a frozen generative model, there is a limitation in that the average accuracy depends on the performance of the pre-trained generative model. Despite this limitation, we have demonstrated the practical application of provided explanations and the fairness of constructed latent space.

## Acknowledgements

This work is in part supported by the National Research Foundation of Korea (NRF, RS-2024-00451435(20%), RS-2024-00413957(15%)), Institute of Information & communications Technology Planning & Evaluation (IITP, 2021-0-01059(20%), 2021-0-00106(20%), 2021-0-00180(20%), RS-2021-II212068(5%)) grant funded by the Ministry of Science and ICT (MSIT), Institute of New Media and Communications(INMAC), and the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University in 2024.

## References

- Ahmed, N. A.; and Gokhale, D. 1989. Entropy expressions and their estimators for multivariate distributions. *IEEE Transactions on Information Theory*, 35(3): 688–692.
- Bardes, A.; Ponce, J.; and LeCun, Y. 2022. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. In *International Conference on Learning Representations*.
- Barocas, S.; and Selbst, A. D. 2016. Big data’s disparate impact. *Calif. L. Rev.*, 104: 671.
- Begley, T.; Schwedes, T.; Frye, C.; and Feige, I. 2020. Explainability for fair machine learning. *arXiv preprint arXiv:2010.07389*.
- Behrmann, J.; Vicol, P.; Wang, K.-C.; Grosse, R.; and Jacobsen, J.-H. 2021. Understanding and mitigating exploding inverses in invertible neural networks. In *International Conference on Artificial Intelligence and Statistics*, 1792–1800. PMLR.
- Brown, C. 2003. Giving up levelling down. *Economics & Philosophy*, 19(1): 111–134.
- Caton, S.; and Haas, C. 2024. Fairness in machine learning: A survey. *ACM Computing Surveys*, 56(7): 1–38.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chiappa, S. 2019. Path-specific counterfactual fairness. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 7801–7808.
- Cho, J.; Zala, A.; and Bansal, M. 2023. Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3043–3054.
- Chouldechova, A. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2): 153–163.
- Christiano, T.; and Braynen, W. 2008. Inequality, injustice and levelling down. *Ratio*, 21(4): 392–420.
- Creager, E.; Madras, D.; Jacobsen, J.-H.; Weis, M.; Swersky, K.; Pitassi, T.; and Zemel, R. 2019. Flexibly fair representation learning by disentanglement. In *International conference on machine learning*, 1436–1445. PMLR.
- Dash, S.; Balasubramanian, V. N.; and Sharma, A. 2022. Evaluating and mitigating bias in image classifiers: A causal perspective using counterfactuals. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 915–924.
- Denton, E.; Hutchinson, B.; Mitchell, M.; Gebru, T.; and Zaldivar, A. 2019. Image counterfactual sensitivity analysis for detecting unintended bias. *arXiv preprint arXiv:1906.06439*.
- Donini, M.; Oneto, L.; Ben-David, S.; Shawe-Taylor, J. S.; and Pontil, M. 2018. Empirical risk minimization under fairness constraints. *Advances in neural information processing systems*, 31.
- Doran, B. 2001. Reconsidering the levelling-down objection against egalitarianism. *Utilitas*, 13(1): 65–85.
- Esser, P.; Rombach, R.; and Ommer, B. 2020. A disentangling invertible interpretation network for explaining latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9223–9232.
- Gustafson, L.; Rolland, C.; Ravi, N.; Duval, Q.; Adcock, A.; Fu, C.-Y.; Hall, M.; and Ross, C. 2023. Facet: Fairness in computer vision evaluation benchmark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20370–20382.
- Han, H.; Kim, S.; Joo, H.; Hong, S.; and Lee, J. 2024. Mitigating Spurious Correlations via Disagreement Probability. *arXiv preprint arXiv:2411.01757*.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.
- Hong, S.; Park, I.; and Chun, S. Y. 2023. On the robustness of normalizing flows for inverse problems in imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10745–10755.
- Jacovi, A.; Marasović, A.; Miller, T.; and Goldberg, Y. 2021. Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 624–635.
- Joo, H.; Kim, J.; Han, H.; and Lee, J. 2023. Distributional Prototypical Methods for Reliable Explanation Space Construction. *IEEE Access*, 11: 34821–34834.
- Joo, J.; and Kärkkäinen, K. 2020. Gender slopes: Counterfactual fairness for computer vision models by attribute manipulation. In *Proceedings of the 2nd international workshop on fairness, accountability, transparency and ethics in multimedia*, 1–5.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *International Conference on Learning Representations*.
- Kehrenberg, T.; Bartlett, M.; Thomas, O.; and Quadrianto, N. 2020. Null-sampling for interpretable and fair representations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, 565–580. Springer.

- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- Kim, H.; and Mnih, A. 2018. Disentangling by factorising. In *International conference on machine learning*, 2649–2658. PMLR.
- Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31.
- Kirichenko, P.; Izmailov, P.; and Wilson, A. G. 2020. Why normalizing flows fail to detect out-of-distribution data. *Advances in neural information processing systems*, 33: 20578–20589.
- Liu, J.; Li, Z.; Yao, Y.; Xu, F.; Ma, X.; Xu, M.; and Tong, H. 2022. Fair representation learning: An alternative to mutual information. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1088–1097.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, 3730–3738.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Park, S.; Lee, J.; Lee, P.; Hwang, S.; Kim, D.; and Byun, H. 2022. Fair contrastive learning for facial attribute classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10389–10398.
- Plecko, D.; and Bareinboim, E. 2023. Causal Fairness for Outcome Control. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Poole, B.; Ozair, S.; Van Den Oord, A.; Alemi, A.; and Tucker, G. 2019. On variational bounds of mutual information. In *International Conference on Machine Learning*, 5171–5180. PMLR.
- Preechakul, K.; Chatthee, N.; Wizadwongsa, S.; and Suwanajakorn, S. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10619–10629.
- Qiang, Y.; Li, C.; Brocanelli, M.; and Zhu, D. 2022. Counterfactual Interpolation Augmentation (CIA): A Unified Approach to Enhance Fairness and Explainability of DNN. In *IJCAI*, 732–739.
- Ramaswamy, V. V.; Kim, S. S.; and Russakovsky, O. 2021. Fair attribute classification through latent space de-biasing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9301–9310.
- Richardson, R.; Schultz, J. M.; and Crawford, K. 2019. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *NYUL Rev. Online*, 94: 15.
- Rudin, C. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, 1(5): 206–215.
- Sagawa\*, S.; Koh\*, P. W.; Hashimoto, T. B.; and Liang, P. 2020. Distributionally Robust Neural Networks. In *International Conference on Learning Representations*.
- Sarhan, M. H.; Navab, N.; Eslami, A.; and Albarqouni, S. 2020. Fairness by learning orthogonal disentangled representations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, 746–761. Springer.
- Shao, H.; Yang, Y.; Lin, H.; Lin, L.; Chen, Y.; Yang, Q.; and Zhao, H. 2022. Rethinking controllable variational autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19250–19259.
- Shrestha, R.; Zou, Y.; Chen, Q.; Li, Z.; Xie, Y.; and Deng, S. 2024. FairRAG: Fair human generation via fair retrieval augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11996–12005.
- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (itw)*, 1–5. IEEE.
- Tommasi, T.; Patricia, N.; Caputo, B.; and Tuytelaars, T. 2017. A deeper look at dataset bias. *Domain adaptation in computer vision applications*, 37–55.
- Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. P. 2019. Fairness constraints: A flexible approach for fair classification. *Journal of Machine Learning Research*, 20(75): 1–42.
- Zafar, M. B.; Valera, I.; Rogriguez, M. G.; and Gummadi, K. P. 2017. Fairness constraints: Mechanisms for fair classification. In *Artificial intelligence and statistics*, 962–970. PMLR.
- Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow twins: Self-supervised learning via redundancy reduction. In *International conference on machine learning*, 12310–12320. PMLR.
- Zhang, F.; He, Q.; Kuang, K.; Liu, J.; Chen, L.; Wu, C.; Xiao, J.; and Zhang, H. 2024. Distributionally Generative Augmentation for Fair Facial Attribute Classification. *arXiv preprint arXiv:2403.06606*.
- Zhang, F.; Kuang, K.; Chen, L.; Liu, Y.; Wu, C.; and Xiao, J. 2023. Fairness-aware Contrastive Learning with Partially Annotated Sensitive Attributes. In *The Eleventh International Conference on Learning Representations*.
- Zhang, M.; Sohoni, N. S.; Zhang, H. R.; Finn, C.; and Re, C. 2022. Correct-N-Contrast: a Contrastive Approach for Improving Robustness to Spurious Correlations. In *International Conference on Machine Learning*, 26484–26516. PMLR.
- Zhang, Z.; Song, Y.; and Qi, H. 2017. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5810–5818.
- Zhu, Z.; Yao, Y.; Sun, J.; Li, H.; and Liu, Y. 2023. Weak proxies are sufficient and preferable for fairness with missing sensitive attributes. In *International Conference on Machine Learning*, 43258–43288. PMLR.

Zietlow, D.; Lohaus, M.; Balakrishnan, G.; Kleindessner, M.; Locatello, F.; Schölkopf, B.; and Russell, C. 2022. Leveling down in computer vision: Pareto inefficiencies in fair deep classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10410–10421.