
Chasing Fairness Under Distribution Shift: A Model Weight Perturbation Approach

Zhimeng Jiang^{1*}, Xiaotian Han^{1*}, Hongye Jin¹, Guanchu Wang², Rui Chen³, Na Zou¹, Xia Hu²
¹Texas A&M University, ²Rice University, ³Samsung Electronics America

Abstract

Fairness in machine learning has attracted increasing attention in recent years. The fairness methods improving algorithmic fairness for in-distribution data may not perform well under distribution shifts. In this paper, we first theoretically demonstrate the inherent connection between distribution shift, data perturbation, and model weight perturbation. Subsequently, we analyze the sufficient conditions to guarantee fairness (i.e., low demographic parity) for the target dataset, including fairness for the source dataset, and low prediction difference between the source and target datasets for each sensitive attribute group. Motivated by these sufficient conditions, we propose robust fairness regularization (RFR) by considering the worst case within the model weight perturbation ball for each sensitive attribute group. We evaluate the effectiveness of our proposed RFR algorithm on synthetic and real distribution shifts across various datasets. Experimental results demonstrate that RFR achieves better fairness-accuracy trade-off performance compared with several baselines. The source code is available at https://github.com/zhimengj0326/RFR_NeurIPS23.

1 Introduction

Previous research [1-4] has shown that a classifier trained on a specific source distribution will perform worse when testing on a different target distribution, due to distribution shift. Recently, many studies have focused on investigating the impact of distribution shift on machine learning models, where fairness performance degradation is even more significant than that of prediction performance [5]. The sensitivity of fairness over distribution shift challenges machine learning models in high stake applications, such as criminal justice [6], healthcare [7], and job marketing [8]. Thus the transferability of the fairness performance under distribution shift is a crucial consideration for real-world applications.

To achieve the fairness of the model (already achieved fairness on the source dataset) on the target dataset, we first reveal that distribution shift is equivalent to model weight perturbation, and then seek to achieve fairness under distribution shift via model weight perturbation. Specifically, *i*) we reveal the inherent connection between distribution shift, data perturbation, and model weight perturbation. We theoretically demonstrate that any distribution shift can be equivalent to data perturbation and model weight perturbation in terms of loss value. In other words, the effect of distribution shift on model training can be attributed to data or model perturbation. *ii*) Given the established connection between the distribution shift and model weight perturbation, we next tackle fairness under the distribution shift problem. We first investigate demographic parity relation between source and target datasets. Achieving fair prediction (e.g., low demographic parity) in the source dataset is insufficient for fair prediction in the target dataset. More importantly, the average prediction difference between source and target datasets with the same sensitive attribute also matters for achieving fairness in target datasets.

*Equal contribution.

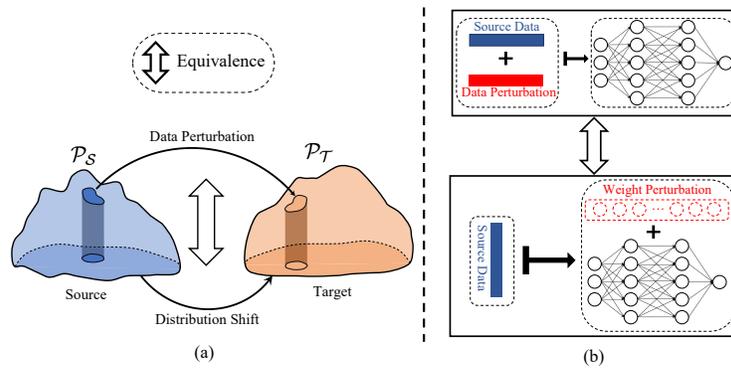


Figure 1: The overview of distribution shift understanding. The left part demonstrates distribution shift can be transformed as data perturbation, while the right part shows that data perturbation and model weight perturbation are equivalent.

Motivated by the established equivalence connection, we propose robust fairness regularization (RFR) to enforce average prediction robustness over model weight. In this way, the well-trained model can tackle the distribution shift problem in terms of demographic parity. Considering the expensive computation for the inner maximization problem in RFR, we accelerate RFR by obtaining the approximate closed-form model weight perturbation using first-order Taylor expansion. In turn, an efficient RFR algorithm, trading the inner maximization problem into two forward and backward propagations for each model update, is proposed to achieve robust fairness under distribution shift. Our contributions are highlighted as follows:

- We theoretically reveal the inherent connection between distribution shift, data perturbation, and model weight perturbation. In other words, distribution shift and model perturbation are equivalent in terms of loss value for any model architecture and loss function.
- We first analyze the sufficient conditions to guarantee fairness transferability under distribution shift. Based on the established connection, we propose RFR explicitly pursuing sufficient conditions for robust fairness by considering the worst case of model perturbation.
- We evaluate the effectiveness of RFR on various real-world datasets with synthetic and real distribution shifts. Experiments results demonstrate that RFR can mostly achieve better fairness-accuracy tradeoff with both synthetic and real distribution shifts.

2 Understanding Distribution Shift

In this section, we first provide the notations used in this paper. And then, we theoretically understand the relations between distribution shift, data perturbation, and model weight perturbation, as shown in Figure 1.

2.1 Notations

We consider source dataset \mathcal{D}_S and target dataset \mathcal{D}_T , defined as a probability distribution \mathcal{P}_S and \mathcal{P}_T for samples $S \in \mathcal{S}$ and $T \in \mathcal{T}$, respectively. Each sample defines values for three random variables: features X with arbitrary domain \mathcal{X} , binary sensitive attribute $A \in \mathcal{A} = \{0, 1\}$, and label Y arbitrary domain \mathcal{Y} , i.e., $\mathcal{S} = \mathcal{T} = \mathcal{X} \times \mathcal{A} \times \mathcal{Y}$. We denote δ as data perturbation on source domain \mathcal{S} , where $\delta_X(X)$ and $\delta_Y(Y)$ data perturbation of features and labels. Using $\mathcal{P}(\cdot)$ to denote the space of probability distributions over some domain, we denote the space of distributions over examples as $\mathcal{P}(\mathcal{S})$. We use $\|\cdot\|_p$ as L_p norm. Let $f_\theta(x)$ be the output of the neural networks parameterized with θ to approximate the true label y . We define $l(f_\theta(x), y)$ as the loss function for neural network training, and the optimal model parameters θ^* training on source data \mathcal{S} is given by $\theta^* = \arg \min_{\theta} \mathcal{R}_S$, where $\mathcal{R}_S = \mathbb{E}_{(X,Y) \sim \mathcal{P}_S} [l(f_\theta(X), Y)]$. Since the target distribution \mathcal{P}_T maybe be different with source distribution \mathcal{P}_S , the well-trained model $f_{\theta^*}(\cdot)$ trained on source dataset \mathcal{S} typically does not perform well in target dataset \mathcal{T} .

2.2 Distribution Shift is Data Perturbation

Deep neural networks are immensely challenged by data quality or dynamic environments [9, 10]. The well-trained deep neural network model may not perform well if existing feature/label noise in training data or distribution shifts among training and target environments. In this subsection, we reveal the inherent equivalence between distribution shift and data perturbation for any neural networks $f_\theta(\cdot)$, source distribution \mathcal{P}_S , and target distribution \mathcal{P}_T using optimal transport. We first provide the formal definition of optimal transport as follows:

Definition 1 (Optimal Transport [11]). *Considering two distributions with probability distribution P_S and P_T , and cost function moving from s to t as $c(s, t)$, optimal transport between probability distribution P_S and P_T is given by*

$$\gamma^*(s, t) = \arg \inf_{\gamma \in \Gamma(P_S, P_T)} \iint c(s, t) \gamma(s, t) ds dt, \quad (1)$$

where the distribution set $\Gamma(P_S, P_T)$ is the collection of all possible transportation plans and given by $\Gamma(P_S, P_T) = \left\{ \gamma(s, t) > 0, \int \gamma(s, t) dt = P_S(s), \int \gamma(s, t) ds = P_T(t) \right\}$. In other words, the distribution set consists of all possible joint distributions with margin distribution P_S and P_T .

Based on the definition of optimal transport, we demonstrate that distribution shift is equivalent to data perturbation in the following theorem:

Theorem 2.1. *For any two different distributions with probability distribution P_S and P_T , adding data perturbation δ ² on source data S can make perturbed source data and target data with the same distribution, where the distribution of data perturbation δ is given by*

$$\mathcal{P}(\delta) = \int_S \gamma^*(s, s + \delta) ds. \quad (2)$$

Additionally, for any positive $p > 0$, data perturbation δ with minimal power $\mathbb{E}[\|\delta\|_p^p]$ is given by Eq. (2) if optimal transport plan $\gamma^*(\cdot, \cdot)$ is calculated based on Eq. (1) with cost function $c(s, t) = \|s - t\|_p^p$.

Proof sketch. Given two different distributions, there are many possible perturbations to move one distribution to another. Optimal transport can select the optimal feasible perturbations or transportation plan in terms of specific objectives (e.g., perturbations power). Given the optimal transportation plan, we can derive the distribution of perturbations based on basic probability theory.

Theorem 2.1 demonstrates the equivalence between distribution shift and data perturbation, where data perturbation distribution is dependent on optimal transport between source and target distribution. The intuition is that such distribution shift can be achieved via optimal transportation (i.e., data perturbation). Based on Theorem 2.1, we have the following Corollary 2.2 on the equivalent of neural network behavior for source and target data:

Corollary 2.2. *Given source and target datasets with probability distribution \mathcal{P}_S and \mathcal{P}_T , there exists data perturbation δ so that the training loss of any neural network $f_\theta(\cdot)$ for target distribution equals that for source distribution with data perturbation δ , i.e.,*

$$\mathbb{E}_{(X, Y) \sim \mathcal{P}_T} [l(f_\theta(X), Y)] = \mathbb{E}_{\delta_X(X), \delta_Y(Y)} \mathbb{E}_{(X, Y) \sim \mathcal{P}_S} [l(f_\theta(X + \delta_X(X)), Y + \delta_Y(Y))]. \quad (3)$$

In other words, for the model trained with loss minimization on source data, the deteriorated performance on the target dataset stems from the perturbation of features and labels. The proof sketch is based on Theorem 2.1 since adding a perturbation on the source dataset can be consistent with the target dataset distribution. Therefore, the conclusion can hold for any loss function that is only dependent on the dataset.

2.3 Data Perturbation Equals Model Weight Perturbation

Although we understand that the distribution shift can be attributed to the data perturbation of features and labels in source data, it is still unclear how to tackle the distribution shift issue. A natural solution

²Data perturbation δ includes the perturbation for features $\delta_X(X)$ and labels $\delta_Y(Y)$.

is to adopt adversarial training to force the well-trained model to be robust over data perturbation. However, it is complicated to generate data perturbation on features and labels simultaneously, and many well-developed adversarial training methods are mainly designed for adversarial feature perturbation. Fortunately, we show that model weight perturbation is equivalent to data perturbation by Theorem 2.3.

Theorem 2.3. *Considering the source dataset with distribution \mathcal{P}_S , suppose the source dataset is perturbed with data perturbation δ , and the neural network is given by $f_\theta(\cdot)$, for general case, there exists model weight perturbation $\Delta\theta$ so that the training loss on perturbed source dataset is the same with that for model weight perturbation $\Delta\theta$ on source distribution:*

$$\mathbb{E}_{\delta_X(X), \delta_Y(Y)} \mathbb{E}_{(X,Y) \sim \mathcal{P}_S} [l(f_\theta(X + \delta_X(X)), Y + \delta_Y(Y))] = \mathbb{E}_{(X,Y) \sim \mathcal{P}_S} [l(f_{\theta+\Delta\theta}(X), Y)]. \quad (4)$$

Proof sketch. The loss under data perturbation and weight perturbation can be analyzed using first-order Taylor expansion. The critical step is to find the condition of the equivalence for the first order term of data perturbation and weight perturbation. Fortunately, the existence of such conditions can be easily proved using linear algebra.

Theorem 2.3 demonstrates that the training loss on perturbed data distribution (but fixed model weight) equals the training loss on perturbed model weight (but original data distribution). In other words, the training loss fluctuation from data perturbation can equal model weight perturbation. Furthermore, we conclude that chasing a robust model over data perturbation can be achieved via model weight perturbation, i.e., finding a “flattened” local minimum in terms of the target objective is sufficient for robustness.

3 Methodology

In this section, we first analyze the sufficient condition to achieve robust fairness over distribution shift in terms of demographic parity. Based on the analysis, we propose a simple yet effective robust fairness regularization via explicit adopting group model weight perturbation. Note that the proposed robust fairness regularization involves a computation-expensive maximization problem, we further accelerate model training via first-order Taylor expansion.

3.1 Robust Fairness Analysis

We consider binary sensitive attribute $A \in \{0, 1\}$ and demographic parity as fairness metric, i.e., the average prediction gap of model $f_\theta(\cdot)$ for different sensitive attribute groups in target dataset $\Delta DP_{\mathcal{T}} = |\mathbb{E}_{\mathcal{T}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_1}[f_\theta(\mathbf{x})]|$, where \mathcal{T}_0 and \mathcal{T}_1 represent the target datasets for sensitive attribute $A = 0$ and $A = 1$, respectively. However, only source dataset \mathcal{S} is available for neural network training. In other words, even though demographic parity on source dataset $\Delta DP_{\mathcal{S}} = |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})]|$ is low, demographic parity on target dataset $\Delta DP_{\mathcal{T}}$ may not be guaranteed to be low due to distribution shift.

To investigate robust fairness over distribution shift, we try to reveal the connection between demographic parity for source and target datasets. We bound the demographic parity difference for source and target datasets as follows:

$$\begin{aligned} DP_{\mathcal{T}} &\stackrel{(a)}{\leq} DP_{\mathcal{S}} + \left| \mathbb{E}_{\mathcal{T}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_1}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] + \mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})] \right| \\ &\stackrel{(b)}{\leq} DP_{\mathcal{S}} + |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_0}[f_\theta(\mathbf{x})]| + |\mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_1}[f_\theta(\mathbf{x})]|, \end{aligned} \quad (5)$$

where inequality (a) and (b) hold due to $a - b \leq |a - b|$ and $||a - b| - |a' - b'|| \leq |a - a'| + |b - b'|$, respectively, for any a, a', b, b' . In other words, in order to minimize demographic parity for target dataset $DP_{\mathcal{T}}$, the objective of $DP_{\mathcal{S}}$ minimization is insufficient. The minimization of prediction difference for source and target datasets given sensitive attribute groups $A = 0$ and $A = 1$, defined as $\Delta_0 = |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_0}[f_\theta(\mathbf{x})]|$ and $\Delta_1 = |\mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_1}[f_\theta(\mathbf{x})]|$, are also beneficial to achieve robust fairness over distribution shift in terms of demographic parity.

The bound in Eq. (5) is tight when both condition (a) $DP_{\mathcal{S}} \leq DP_{\mathcal{T}}$ and (b) maximum and minimum of value set $\min \{ \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})], \mathbb{E}_{\mathcal{T}_0}[f_\theta(\mathbf{x})] \}$, $\min \{ \mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})], \mathbb{E}_{\mathcal{T}_1}[f_\theta(\mathbf{x})] \}$ both are from source or target distribution. Even though conditions (a) and (b) may not hold for the neural network model, we

would like to that our goal is not to obtain a tight bound for demographic parity on target distribution. Instead, we aim to find sufficient conditions to guarantee low demographic parity and such low demographic parity can be achieved in model training without any target distribution information. The proposed upper bound Eq. (5) actually reveals sufficient conditions, which can be achieved by our proposed RFR algorithm.

3.2 Robust Fairness Regularization

Motivated by Section 3.1 and distribution shift understanding in Section 2, we develop a robust fairness regularization to achieve robust fairness over distribution shift. Section 3.1 demonstrates that fair prediction on the target dataset requires fair prediction on the source dataset and low prediction difference between the source and target dataset for each sensitive attribute, i.e., low $\Delta_0 = |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_0}[f_\theta(\mathbf{x})]|$ and $\Delta_1 = |\mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{T}_1}[f_\theta(\mathbf{x})]|$. Based on Theorem 2.3 there exists ϵ_0 so that the following equality holds:

$$\Delta_0 = |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\epsilon_0} \mathbb{E}_{\mathcal{S}_0}[f_{\theta+\epsilon_0}(\mathbf{x})]|. \quad (6)$$

Note that the distribution shift is unknown, we consider the worst case for model weight perturbation ϵ_0 within L_p -norm perturbation ball with radius ρ as follows:

$$\Delta_0 = |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\epsilon_0} \mathbb{E}_{\mathcal{S}_0}[f_{\theta+\epsilon_0}(\mathbf{x})]| \leq \max_{\|\epsilon_0\|_p \leq \rho} |\mathbb{E}_{\mathcal{S}_0}[f_{\theta+\epsilon_0}(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})]|, \quad (7)$$

where $\|\cdot\|_p$ represents L_p norm, ρ and p are hyperparameters. Note that the feasible region of model weight perturbation ϵ_0 is symmetric, and the neural network prediction is locally linear around parameter θ , we have $\mathbb{E}_{\mathcal{S}_0}[f_{\theta+\epsilon_0}(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] \approx -(\mathbb{E}_{\mathcal{S}_0}[f_{\theta-\epsilon_0}(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})])$ due to the local linearity. In other words, the absolute operation can be removed if we consider the maximization problem in a symmetric feasible region since there are always non-negative value for the pair perturbation ϵ_0 and $-\epsilon_0$. Therefore, we can further bound Δ_0 as

$$\begin{aligned} \Delta_0 &\leq \max_{\|\epsilon_0\|_p \leq \rho} |\mathbb{E}_{\mathcal{S}_0}[f_{\theta+\epsilon_0}(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})]| \\ &\approx \max_{\|\epsilon_0\|_p \leq \rho} \mathbb{E}_{\mathcal{S}_0}[f_{\theta+\epsilon_0}(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] \triangleq \mathcal{L}_{RFR, \mathcal{S}_0}, \end{aligned} \quad (8)$$

Similarly, we can bound the prediction difference for source and target distribution with sensitive group $A = 1$ as follows:

$$\Delta_1 \leq \max_{\|\epsilon_1\|_p \leq \rho} \mathbb{E}_{\mathcal{S}_1}[f_{\theta+\epsilon_1}(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})] \triangleq \mathcal{L}_{RFR, \mathcal{S}_1}. \quad (9)$$

Therefore, demographic parity for source and target distribution relation is given by $DP_{\mathcal{T}} \leq DP_{\mathcal{S}} + \mathcal{L}_{RFR, \mathcal{S}_0} + \mathcal{L}_{RFR, \mathcal{S}_1}$, we propose *robust fairness regularization* (RFR) to achieve robust fairness as

$$\mathcal{L}_{RFR} = \mathcal{L}_{RFR, \mathcal{S}_0} + \mathcal{L}_{RFR, \mathcal{S}_1}. \quad (10)$$

It is worth noting that our proposed RFR is agnostic to the training loss function and model architectures. Additionally, we follow [12] to accelerate model training using sharpness-aware minimization via trading maximization problem can be simplified as two forward and two backward propagations. More details on training acceleration are in Appendix E.

3.3 The Proposed Method

In this subsection, we introduce how to use the proposed RFR to achieve robust fairness, i.e., the fair model (low $\Delta DP_{\mathcal{S}}$) trained on the source dataset will also be fair on the target dataset (low $\Delta DP_{\mathcal{T}}$). Considering binary sensitive attribute $A \in \{0, 1\}$ and binary classification problem $Y \in \{0, 1\}$, the classification loss is denoted as

$$\mathcal{L}_{CLF} = \mathbb{E}_{\mathcal{S}}[-Y f_\theta(X) - (1 - Y)(1 - f_\theta(X))]. \quad (11)$$

To achieve fairness, we consider demographic parity as fairness regularization, i.e.,

$$\mathcal{L}_{DP} = |\mathbb{E}_{\mathcal{S}_0}[f_\theta(\mathbf{x})] - \mathbb{E}_{\mathcal{S}_1}[f_\theta(\mathbf{x})]|, \quad (12)$$

Table 1: Performance Comparison with Baselines on Synthetic Dataset. (α, β) control distribution shift intensity, and $(0, 1)$ represents no distribution shift. The best/second-best results are highlighted in **boldface**/underlined, respectively.

(α, β)	Methods	Adult			ACS-I			ACS-E		
		Acc (%) \uparrow	Δ_{DP} (%) \downarrow	Δ_{EO} (%) \downarrow	Acc (%) \uparrow	Δ_{DP} (%) \downarrow	Δ_{EO} (%) \downarrow	Acc (%) \uparrow	Δ_{DP} (%) \downarrow	Δ_{EO} (%) \downarrow
(1.0, 2.0)	MLP	82.09 \pm 0.05	15.11 \pm 0.04	14.33 \pm 0.05	77.95 \pm 0.52	3.51 \pm 0.59	3.77 \pm 0.55	80.95 \pm 0.10	1.10 \pm 0.06	1.43 \pm 0.06
	REG	80.60 \pm 0.05	3.79 \pm 0.06	3.27 \pm 0.08	77.77 \pm 0.09	2.28 \pm 0.32	2.59 \pm 0.23	80.44 \pm 0.07	0.86\pm0.09	1.05 \pm 0.10
	ADV	78.80 \pm 0.68	0.83 \pm 0.26	0.79 \pm 0.14	75.72 \pm 0.63	1.96 \pm 0.38	2.00 \pm 0.35	79.39 \pm 0.15	<u>1.09\pm0.26</u>	0.95 \pm 0.26
	FCR	79.06 \pm 0.09	9.98\pm0.06	9.47\pm0.07	76.99 \pm 0.47	<u>2.94\pm0.34</u>	<u>2.95\pm0.28</u>	79.74 \pm 0.11	0.97 \pm 0.21	1.00 \pm 0.22
	RFR	78.84 \pm 0.09	0.44\pm0.05	0.12\pm0.06	74.15 \pm 0.81	1.84\pm0.27	1.60\pm0.33	80.08 \pm 0.08	0.71\pm0.10	0.06\pm0.11
(1.5, 3.0)	MLP	82.05 \pm 0.05	15.16 \pm 0.09	14.33 \pm 0.09	77.85 \pm 0.25	3.73 \pm 0.53	3.70 \pm 0.56	80.42 \pm 0.10	1.14 \pm 0.07	1.10 \pm 0.07
	REG	80.64 \pm 0.08	3.74 \pm 0.11	3.23 \pm 0.10	77.87 \pm 0.18	2.25 \pm 0.28	2.37 \pm 0.27	80.21 \pm 0.13	0.72\pm0.04	0.75 \pm 0.03
	ADV	78.71 \pm 0.41	1.07 \pm 0.87	0.87 \pm 0.96	75.79 \pm 0.68	2.22 \pm 0.53	2.44 \pm 0.48	79.58 \pm 0.13	1.07 \pm 0.19	1.26 \pm 0.18
	FCR	79.05 \pm 0.12	<u>10.01\pm0.07</u>	<u>9.51\pm0.06</u>	77.06 \pm 0.68	3.39 \pm 0.33	3.10 \pm 0.36	79.59 \pm 0.26	1.17 \pm 0.24	1.08 \pm 0.23
	RFR	78.91 \pm 0.03	0.46\pm0.10	0.16\pm0.09	74.19 \pm 0.58	1.82\pm0.29	2.17\pm0.32	80.47 \pm 0.03	0.72\pm0.04	0.71\pm0.05
(3.0, 6.0)	MLP	82.07 \pm 0.05	15.23 \pm 0.14	14.45 \pm 0.15	77.89 \pm 0.45	3.35 \pm 0.36	3.47 \pm 0.41	80.30 \pm 0.04	1.17 \pm 0.04	1.13 \pm 0.04
	REG	80.62 \pm 0.07	3.72 \pm 0.05	3.21 \pm 0.04	78.19 \pm 0.12	1.60\pm0.48	1.84\pm0.44	80.36 \pm 0.09	0.70\pm0.09	0.68 \pm 0.11
	ADV	78.97 \pm 0.49	1.28\pm0.74	1.09\pm0.50	75.71 \pm 0.68	2.28 \pm 0.39	2.24 \pm 0.41	79.66 \pm 0.16	1.34 \pm 0.14	<u>1.16\pm0.13</u>
	FCR	79.03 \pm 0.13	10.00 \pm 0.05	9.50 \pm 0.05	76.71 \pm 0.39	2.97 \pm 0.34	3.28 \pm 0.31	79.89 \pm 0.22	1.06 \pm 0.14	1.14 \pm 0.18
	RFR	80.15 \pm 0.07	<u>1.75\pm0.15</u>	<u>1.30\pm0.14</u>	74.22 \pm 0.56	<u>1.80\pm0.26</u>	<u>1.89\pm0.24</u>	80.28 \pm 0.12	<u>0.74\pm0.04</u>	0.51\pm0.04

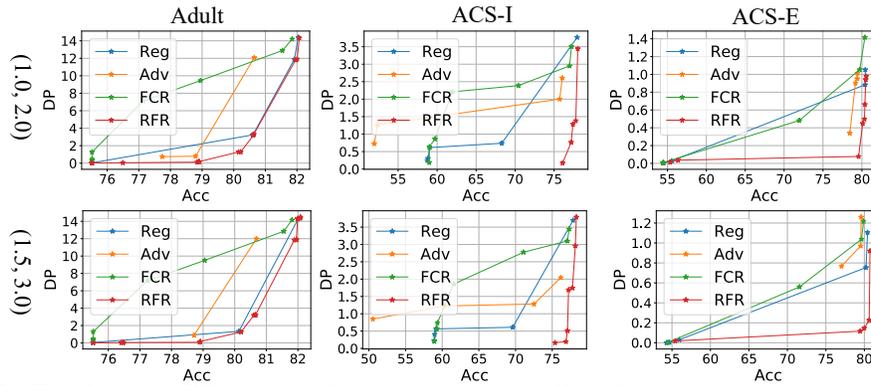


Figure 2: The fairness (DP) and prediction (Acc) trade-off performance on three datasets with different synthetic distribution shifts. The units for x- and y-axis are percentages (%).

Additionally, we also adopt \mathcal{L}_{RFR} as robust fairness regularization. The overall objective is given as follows:

$$\mathcal{L}_{all} = \mathcal{L}_{CLF} + \lambda \cdot (\mathcal{L}_{DP} + \mathcal{L}_{RFR}), \quad (13)$$

where λ is a hyperparameter to balance the model performance and fairness. \mathcal{L}_{CLF} is the loss function for the downstream task, \mathcal{L}_{DP} is to ensure the demographic parity on the source dataset, and \mathcal{L}_{RFR} is to ensure the transferability of fairness from the source dataset to the target dataset.

In the proposed RFR algorithm, we directly apply approximation to accelerate model training. Note that the optimal model weight perturbation is dependent on the current model weight θ , two forward and backward propagations are required to calculate the final gradient.

4 Experiments

In this section, we conduct experiments to evaluate the effectiveness of our proposed RFR, aiming to answer the three research questions. **Q1:** How effective is RFR to achieve fairness under distribution shift for synthetic distribution shift? **Q2:** How effective is RFR for real spatial and temporal distribution shift? **Q3:** How sensitive is RFR to the key hyperparameter λ ?

4.1 Experimental Setting

In this subsection, we present the experimental setting, including datasets, evaluation metrics, baseline methods, distribution shift generation, and implementing details.

Datasets. We adopt the following datasets in our experiments. **UCI Adult** [13] dataset contains information about 45,222 individuals with 15 attributes from the 1994 US Census. We consider

gender as the sensitive attribute and the task is to predict whether the income of the person is higher than \$50k or not. **ACS-Income** [14] is extracted from the American Community Survey (ACS) Public Use Microdata Sample (PUMS) with 3, 236, 107 samples. We choose gender as the sensitive attribute. Similar to the task in UCI Adult, the task is to predict whether the individual income is above \$50k. **ACS-Employment** [14] also derives from ACS PUMS, and we also use gender as the sensitive attribute. The task is to predict whether an individual is employed.

Evaluation Metrics. We use accuracy to evaluate the prediction performance for the downstream task. For fairness metrics, we adopt two common-used quantitative group fairness metrics to measure the prediction bias [15, 16], i.e., *demographic parity* $\Delta_{DP} = |\mathbb{P}(\hat{Y} = 1|A = 0) - \mathbb{P}(\hat{Y} = 1|A = 1)|$ and *equal opportunity* $\Delta_{EO} = |\mathbb{P}(\hat{Y} = 1|A = 0, Y = 1) - \mathbb{P}(\hat{Y} = 1|A = 1, Y = 1)|$, where A represents sensitive attribute, Y and \hat{Y} represent the ground-truth label and predicted label, respectively.

Baselines. In our experiments, we consider vanilla multi-layer perceptron (MLP) and two widely adopted in-processing debiasing methods, including fairness regularization (REG), adversarial debiasing (ADV), and fair consistency regularization (FCR). **MLP** directly uses vanilla 3-layer MLP with 50 hidden unit and ReLU activation function [17] to minimize cross-entropy loss with source dataset. In the experiments, we adopt the same model architecture for all other methods (i.e., REG and ADV). **REG** adds a fairness-related metric as a regularization term in the objective function to mitigate the prediction bias [18, 19]. Specifically, we directly adopt demographic parity as a regularization term, and the objective function is $\mathcal{L}_{CLF} + \lambda\mathcal{L}_{DP}$, where \mathcal{L}_{CLF} and \mathcal{L}_{DP} are defined in Eqs. (11) and (12). **ADV** [20] employs a two-player game to mitigate bias, in which a classification network is trained to predict labels based on input features, and an adversarial network takes the output of the classification network as input and aims to identify which sensitive attribute group the sample belongs to. **FCR** [21] aims to minimize and balance consistency loss across groups.

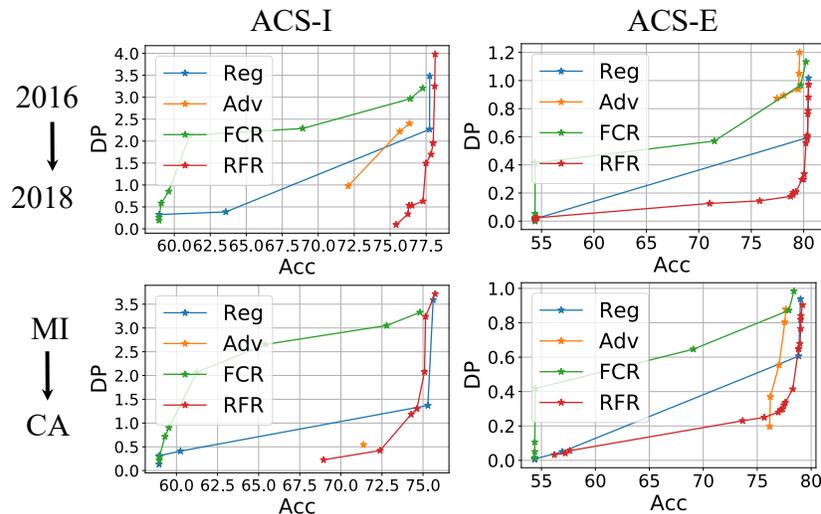


Figure 3: DP and Acc trade-off performance on three real-world datasets with temporal (Top) and spatial (Bottom) distribution shift. The trade-off curve close to the right bottom corner means better trade-off performance. The units for x- and y-axis are percentages (%).

Synthetic and Real Distribution Shift. We adopt synthetic and real distribution shifts. For synthetic distribution shift, we follow work [22, 23] to generate distribution shift via biased sampling. Specifically, we adopt applying principal component analysis (PCA) [24] to retrieve the first principal component \mathcal{C} from input attributes. Subsequently, we estimate the mean $\mu(\mathcal{C})$ and standard deviation $\sigma(\mathcal{C})$, and set Gaussian distribution $\mathcal{N}_{\mathcal{S}}(\mu(\mathcal{C}), \sigma(\mathcal{C}))$ to randomly sampling of target dataset. As for source dataset, we choose different parameters α and β to generate distribution shift using another Gaussian distribution $\mathcal{N}_{\mathcal{T}}(\mu(\mathcal{C}) + \alpha, \frac{\sigma(\mathcal{C})}{\beta})$ for randomly sampling. The source and target datasets are constructed by sampling without replacement. For real distribution shift, we adopt the sampling

and pre-processing approaches following *Folktables* [14] to generate spatial and temporal distribution shift via data partition based on different US states and year from 2014 to 2018.

Implementation Details. We run the experiments 5 times and report the average performance for each method. We adopt Adam optimizer with 10^{-5} learning rate and 0.01 weight decay for all models. For baseline ADV, we alternatively train classification and adversarial networks with 70 and 30 epochs, respectively. The hyperparameters for ADV are set as $\{0.0, 1.0, 10.0, 100.0, 500.0\}$. For adding regularization, we adopt the hyperparameters set $\{0.0, 0.5, 1.0, 10.0, 30.0, 50.0\}$.

4.2 Experimental Results on Synthetic Distribution Shift

In this experiment, we evaluate the effectiveness of our proposed Robust Fairness Regularization (RFR) method with synthetic distribution shifts via biased sampling with different parameters (α, β) . We compare the performance of RFR with several other baseline methods, including standard training, adversarial training, and fairness regularization methods. The results of performance value are presented in Table 1, and the results of fairness-accuracy tradeoff are presented in Figure 2. We have the following observations:

- The results in Table 1 demonstrate that RFR consistently outperforms the baselines in terms of fairness for small distribution shifts, achieving a better balance between fairness and accuracy. For example, RFR achieves an improvement of 47.0% on metric ΔDP and 84.8% on metric ΔEO compared to the second-best method in the Adult dataset with (1.0, 2.0)-synthetic distribution shift. Furthermore, we observed that the outperforms of RFR compared with baselines decrease as the distribution shift intensity increases. The reason is that the approximation RFR involved Taylor expansion over model perturbation and is effective with mild distribution shifts.
- The results in Figure 2 show that RFR achieved a better fairness-accuracy tradeoff compared to the baseline methods for mild distribution shift. We observed that our proposed method achieved a better Pareto frontier compared to the existing methods, and the bias can be mitigated with tiny accuracy drop.

The experimental results show that our method can effectively address the fairness problem under mild distribution shift while maintaining high accuracy, outperforming the existing state-of-the-art baseline methods.

4.3 Experimental Results on Real Distribution Shift

In this experiment, we evaluate the performance of RFR on multiple real-world datasets with real distribution shifts. We use ACS dataset in the experiment. The distribution of this dataset varies across different time periods or geographic locations, which also causes fairness performance degradation under the distribution shift. The results are presented in Table 2, and the results of the fairness-accuracy tradeoff are presented in Figure 3. The results show that our method consistently outperforms the baselines across all datasets, achieving a better balance between fairness and accuracy. Specifically, we have the following observations:

- Table 2 demonstrates that RFR consistently outperforms the baselines across all datasets in terms of fairness-accuracy tradeoff. This suggests that our method is effective in achieving robust fairness under real temporal and spatial distribution shifts. For example, RFR achieved an improvement of 34.9% on metric ΔDP and 34.4% on metric ΔEO compared to the second-best method in ACS-I dataset with temporal distribution shift.
- Figure 3 shows that RFR can achieve better fairness-accuracy tradeoff than baselines, i.e., RFR is particularly effective in addressing the fairness problem on the ACS dataset with source/target data varying across different time periods or geographic locations. This is important for many practical applications, where fairness is a critical requirement, such as in credit scoring, and loan approval.
- The variance for spatial shift on ACS-I is higher than that of temporal shift on all methods, which indicates neural networks easily converge to different local minima for spatial distribution shift. The proposed methods can achieve better fairness results for most cases and the tradeoff results clearly demonstrate the effectiveness.

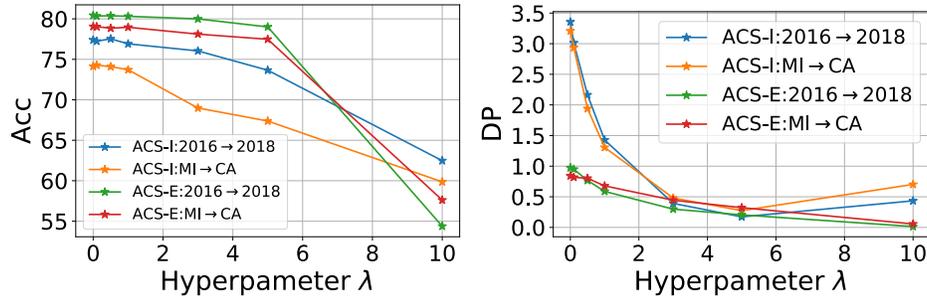


Figure 4: The hyperparameters study for λ . The **Left** subfigure and **Right** subfigure show the results of accuracy and fairness performance, respectively.

Overall, the experimental results on temporal and spatial distribution shifts further support the effectiveness of our proposed method RFR and its potential for practical applications in diverse settings.

Table 2: Performance comparison with baselines on real temporal (the year 2016 to the year 2018) and spatial (Michigan State to California State) distribution shift. The best and second-best results are highlighted with **bold** and underline, respectively.

Real	Methods	ACS-I			ACS-E		
		Acc (%) \uparrow	Δ_{DP} (%) \downarrow	Δ_{EO} (%) \downarrow	Acc (%) \uparrow	Δ_{DP} (%) \downarrow	Δ_{EO} (%) \downarrow
2016 \rightarrow 2018	MLP	77.75 \pm 0.44	3.26 \pm 0.38	3.48 \pm 0.41	80.46 \pm 0.05	1.07 \pm 0.10	1.02 \pm 0.10
	REG	77.74 \pm 0.62	2.09 \pm 0.21	2.27 \pm 0.24	80.37 \pm 0.12	0.77 \pm 0.08	0.74 \pm 0.08
	ADV	75.94 \pm 0.40	2.41 \pm 0.49	2.53 \pm 0.55	79.62 \pm 0.14	1.17 \pm 0.14	1.10 \pm 0.14
	FCR	76.40 \pm 0.45	2.81 \pm 0.30	2.96 \pm 0.30	79.59 \pm 0.38	0.95 \pm 0.42	0.91 \pm 0.34
	RFR	77.49 \pm 0.32	1.36 \pm 0.17	1.49 \pm 0.17	80.36 \pm 0.05	0.61 \pm 0.11	0.58 \pm 0.10
MI \rightarrow CA	MLP	75.62 \pm 0.80	5.22 \pm 0.86	3.60 \pm 0.34	79.02 \pm 0.20	0.73 \pm 0.07	0.94 \pm 0.05
	REG	75.52 \pm 0.78	2.88 \pm 0.44	2.17 \pm 0.22	75.34 \pm 1.11	0.42 \pm 0.09	0.61 \pm 0.11
	ADV	73.38 \pm 1.07	1.04 \pm 0.58	0.54 \pm 0.38	77.56 \pm 0.41	0.61 \pm 0.18	0.80 \pm 0.13
	FCR	74.28 \pm 0.35	5.06 \pm 0.62	3.67 \pm 0.51	77.96 \pm 0.22	0.44 \pm 0.14	0.67 \pm 0.38
	RFR	74.63 \pm 0.45	<u>1.35</u> \pm 0.39	<u>1.30</u> \pm 0.24	78.84 \pm 0.21	<u>0.44</u> \pm 0.09	0.65 \pm 0.07

4.4 Hyperparameter Study

In this experiment, we investigate the sensitivity of the hyperparameter λ in Equation $\mathcal{L}_{all} = \mathcal{L}_{CLF} + \lambda(\mathcal{L}_{DP} + \mathcal{L}_{RFR})$ for spatial and temporal distribution shift across different datasets. Specifically, we tune the hyperparameter as $\lambda = \{0.0, 0.1, 0.5, 1.0, 3.0, 5.0, 10.0\}$. From the results in Figure 4, it is seen that the accuracy and demographic parity are both sensitive to hyperparameter λ , which implies the capability of accuracy-fairness control. With the increase of λ , accuracy decreases while demographic parity also decreases. When λ is smaller than 5, accuracy drops slowly while demographic parity drops faster. Such observation represents that an appropriate hyperparameter can mitigate prediction bias while preserving comparable prediction performance.

5 Related Work

In this section, we present two lines of related work, including fairness in machine learning and distribution shift.

Fairness in Machine Learning. Fairness [25?–31] is a legal requirement for machine learning models for various high-stake real-world predictions, such as healthcare [7, 32, 33], education [34–36], and job market [37, 38]. Achieving fairness, either from a data or model perspective [39–42], in machine learning is a challenging problem. As such, there has been an increasing interest in both the industrial and research community to develop algorithms to mitigate these issues and ensure that machine learning models make fair and unbiased decisions. Extensive efforts led to the development of various techniques and metrics for fairness and proposed various definitions of fairness, such as group fairness [43–47, 40], individual fairness [48–53], and counterfactual fairness [54–56]. In this paper, we focus on group fairness, and the widely used methods to achieve group fairness are fair regularization and adversarial debias method. [19, 18] proposed to add a fairness regularization term

to the objective function to achieve group fairness, and [20] proposes to jointly train a classification network and an adversarial network to mitigate the bias for different demographic groups to achieve group fairness. Overall, ensuring fairness in machine learning is a critical and ongoing research area that will continue to be an important focus for the development of responsible machine learning.

Distribution Shift. Previous work [1-4, 57] reveals that a classifier trained on a source distribution will perform worse on a given target distribution because of the distribution shift, and recently extensive works have explored the influence of distribution shift in model prediction. Moreover, distribution shift can significantly affect the fairness performance of machine learning models. The sensitivity of fairness to the distribution shift is notorious for the legal requirement. The degraded performance of fair models under a distribution shift would trigger new bias and discrimination issues. There are some works [23, 58, 21, 59, 60] that solve fairness under various distribution shifts. For example, [23] explore the fairness under covariate shift, where the inputs change with the in-distribution label. [60] proposes Shifty algorithms to hold fairness guarantees when the dataset in the deployment environment is out-of-distribution of the training datasets (distribution shift). Our work is different from those prior works from model weight perturbation perspective.

6 Conclusion

This paper aims to solve the fairness problem under the distribution shifts from the model weight perturbation perspective. We first establish a theoretical connection between distribution shift, data perturbation, and model weight perturbation, which allowed us to conclude that distribution shift and model perturbation are equivalent. We then propose a sufficient condition for ensuring fairness transference under distribution shift. To explicitly chase such sufficient conditions, we introduce Robust Fairness Regularization (RFR) method based on the established understanding, to achieve robust fairness. Our experiments on both synthetic and real distribution shifts demonstrate the effectiveness of RFR in achieving a better fairness-accuracy tradeoff compared to existing baselines. We believe that our understanding of distribution shift is valuable and intriguing to the development of robust machine learning models, and the proposed RFR approach can be of great practical value to build fair and robust machine learning models in real-world applications.

7 Acknowledgement

The authors thank the anonymous reviewers for their helpful comments. The work is in part supported by NSF grants NSF IIS-2224843, IIS-1939716, and IIS-1900990. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- [1] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.
- [2] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021.
- [3] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *Advances in Neural Information Processing Systems*, 33:18583–18599, 2020.
- [4] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. *Advances in Neural Information Processing Systems*, 33:11539–11551, 2020.
- [5] L Elisa Celis, Anay Mehrotra, and Nisheeth Vishnoi. Fair classification with adversarial perturbations. *Advances in Neural Information Processing Systems*, 34:8158–8171, 2021.
- [6] Songül Tolan, Marius Miron, Emilia Gómez, and Carlos Castillo. Why machine learning may lead to unfairness: Evidence from risk assessment for juvenile justice in catalonia. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, pages 83–92, 2019.
- [7] Muhammad Aurangzeb Ahmad, Arpit Patel, Carly Eckert, Vikas Kumar, and Ankur Teredesai. Fairness in machine learning for healthcare. In *Proceedings of the 26th ACM SIGKDD*, pages 3529–3530, 2020.
- [8] Jeff Johnson, Donald M Truxillo, Berrin Erdogan, Talya N Bauer, and Leslie Hammer. Perceptions of overall fairness: are effects on job performance moderated by leader-member exchange? *Human Performance*, 22(5), 2009.
- [9] Ainhize Barrainkua, Paula Gordaliza, Jose A Lozano, and Novi Quadrianto. A survey on preserving fairness guarantees in changing environments. *arXiv preprint arXiv:2211.07530*, 2022.
- [10] Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre-Alvise Rebuffi, Ira Ktena, Krishnamurthy Dvijotham, and Ali Taylan Cemgil. A fine-grained analysis on distribution shift. In *International Conference on Learning Representations*, 2022.
- [11] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [12] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.
- [13] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [14] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair machine learning. *NeurIPS*, 2021.
- [15] Christos Louizos, Kevin Swersky, Yujia Li, Max Welling, and Richard Zemel. The variational fair autoencoder. *arXiv preprint arXiv:1511.00830*, 2015.
- [16] Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*, 2017.
- [17] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010.
- [18] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 2012.

- [19] Ching-Yao Chuang and Youssef Mroueh. Fair mixup: Fairness via interpolation. In *ICLR*, 2020.
- [20] Gilles Louppe, Michael Kagan, and Kyle Cranmer. Learning to pivot with adversarial networks. *NeurIPS*, 30, 2017.
- [21] Bang An, Zora Che, Mucong Ding, and Furong Huang. Transferring fairness under distribution shifts via fair consistency regularization. *arXiv preprint arXiv:2206.12796*, 2022.
- [22] Arthur Gretton, Alex Smola, Jiayuan Huang, Marcel Schmittfull, Karsten Borgwardt, and Bernhard Schölkopf. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4):5, 2009.
- [23] Ashkan Rezaei, Anqi Liu, Omid Memarrast, and Brian D Ziebart. Robust fairness under covariate shift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9419–9427, 2021.
- [24] Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010.
- [25] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.
- [26] Dana Pessach and Erez Shmueli. A review on fairness in machine learning. *ACM Computing Surveys (CSUR)*, 55(3):1–44, 2022.
- [27] Michael Madaio, Lisa Egede, Hariharan Subramonyam, Jennifer Wortman Vaughan, and Hanna Wallach. Assessing the fairness of ai systems: Ai practitioners’ processes, challenges, and needs for support. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1):1–26, 2022.
- [28] Dexun Li and Pradeep Varakantham. Efficient resource allocation with fairness constraints in restless multi-armed bandits. In James Cussens and Kun Zhang, editors, *Uncertainty in Artificial Intelligence, Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, UAI 2022, 1-5 August 2022, Eindhoven, The Netherlands*, volume 180 of *Proceedings of Machine Learning Research*, pages 1158–1167. PMLR, 2022.
- [29] Kirtan Padh, Diego Antognini, Emma Lejal Glaude, Boi Faltings, and Claudiu Musat. Addressing fairness in classification with a model-agnostic multi-objective algorithm. In Cassio P. de Campos, Marloes H. Maathuis, and Erik Quaeghebeur, editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, pages 600–609. AUAI Press, 2021.
- [30] Zhimeng Jiang, Xiaotian Han, Chao Fan, Zirui Liu, Na Zou, Ali Mostafavi, and Xia Hu. Fmp: Toward fair graph message passing against topology bias. *arXiv preprint arXiv:2202.04187*, 2022.
- [31] Xiaotian Han, Zhimeng Jiang, Hongye Jin, Zirui Liu, Na Zou, Qifan Wang, and Xia Hu. Retiring Δ DP: New distribution-level metrics for demographic parity. *Transactions on Machine Learning Research*, 2023.
- [32] Margrét Vilborg Bjarnadóttir and David Anderson. Machine learning in healthcare: Fairness, issues, and challenges. In *Pushing the Boundaries: Frontiers in Impactful OR/OM Research*, pages 64–83. INFORMS, 2020.
- [33] Thomas Grote and Geoff Keeling. Enabling fairness in healthcare through machine learning. *Ethics and Information Technology*, 24(3):39, 2022.
- [34] Steinar Bøyum. Fairness in education—a normative analysis of oecd policy documents. *Journal of Education Policy*, 29(6):856–870, 2014.
- [35] Paolo Brunori, Vito Peragine, and Laura Serlenga. Fairness in education: The italian university before and after the reform. *Economics of Education Review*, 31(5):764–777, 2012.

- [36] René F Kizilcec and Hansol Lee. Algorithmic fairness in education. In *The ethics of artificial intelligence in education*, pages 174–202. Routledge, 2022.
- [37] Lily Hu and Yiling Chen. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference*, pages 1389–1398, 2018.
- [38] G Stoney Alder and Joseph Gilbert. Achieving ethics and fairness in hiring: Going beyond the law. *Journal of Business Ethics*, 68:449–464, 2006.
- [39] Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. Data-centric artificial intelligence: A survey. *arXiv preprint arXiv:2303.10158*, 2023.
- [40] Qizhang Feng, Zhimeng Jiang, Ruiquan Li, Yicheng Wang, Na Zou, Jiang Bian, and Xia Hu. Fair graph distillation. 2023.
- [41] Chia-Yuan Chang, Yu-Neng Chuang, Zhimeng Jiang, Kwei-Herng Lai, Anxiao Jiang, and Na Zou. Coda: Temporal domain generalization via concept drift simulator. *arXiv preprint arXiv:2310.01508*, 2023.
- [42] Chia-Yuan Chang, Yu-Neng Chuang, Kwei-Herng Lai, Xiaotian Han, Xia Hu, and Na Zou. Towards assumption-free bias mitigation. *arXiv preprint arXiv:2307.04105*, 2023.
- [43] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *NeurIPS*, 29, 2016.
- [44] Sahil Verma and Julia Rubin. Fairness definitions explained. In *2018 IEEE/ACM International Workshop on Software Fairness (Fairware)*, pages 1–7. IEEE, 2018.
- [45] Peizhao Li, Yifei Wang, Han Zhao, Pengyu Hong, and Hongfu Liu. On dyadic fairness: Exploring and mitigating bias in graph connections. In *International Conference on Learning Representations*, 2020.
- [46] Gaurush Hiranandani, Jatin Mathur, Harikrishna Narasimhan, and Oluwasanmi Koyejo. Quadratic metric elicitation for fairness and beyond. In James Cussens and Kun Zhang, editors, *Uncertainty in Artificial Intelligence, Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, UAI 2022, 1-5 August 2022, Eindhoven, The Netherlands*, volume 180 of *Proceedings of Machine Learning Research*, pages 811–821. PMLR, 2022.
- [47] Zhimeng Jiang, Xiaotian Han, Chao Fan, Fan Yang, Ali Mostafavi, and Xia Hu. Generalized demographic parity for group fairness. In *International Conference on Learning Representations*, 2022.
- [48] Mikhail Yurochkin, Amanda Bower, and Yuekai Sun. Training individually fair ml models with sensitive subspace robustness. *International Conference on Learning Representations*, 2020.
- [49] Debarghya Mukherjee, Mikhail Yurochkin, Moulinath Banerjee, and Yuekai Sun. Two simple ways to learn individual fairness metrics from data. In *ICML*, 2020.
- [50] Mikhail Yurochkin and Yuekai Sun. Sensei: Sensitive set invariance for enforcing individual fairness. In *International Conference on Learning Representations*, 2021.
- [51] Jian Kang, Jingrui He, Ross Maciejewski, and Hanghang Tong. Inform: Individual fairness on graph mining. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 379–389, 2020.
- [52] Debarghya Mukherjee, Felix Petersen, Mikhail Yurochkin, and Yuekai Sun. Domain adaptation meets individual fairness. and they get along. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [53] Shantanu Das, Swapnil Dhamal, Ganesh Ghalme, Shweta Jain, and Sujit Gujar. Individual fairness in feature-based pricing for monopoly markets. In James Cussens and Kun Zhang, editors, *Uncertainty in Artificial Intelligence, Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, UAI 2022, 1-5 August 2022, Eindhoven, The Netherlands*, volume 180 of *Proceedings of Machine Learning Research*, pages 486–495. PMLR, 2022.

- [54] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *NeurIPS*, 30, 2017.
- [55] Chirag Agarwal, Himabindu Lakkaraju, and Marinka Zitnik. Towards a unified framework for fair and stable graph representation learning. In *UAI 2021: Uncertainty in Artificial Intelligence*, 2021.
- [56] Aoqi Zuo, Susan Wei, Tongliang Liu, Bo Han, Kun Zhang, and Mingming Gong. Counterfactual fairness with partially known causal graph. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [57] Chia-Yuan Chang, Yu-Neng Chuang, Guanchu Wang, Mengnan Du, and Zou Na. Dispel: Domain generalization via domain-specific liberating. *arXiv preprint arXiv:2307.07181*, 2023.
- [58] Jessica Schrouff, Natalie Harris, Oluwasanmi Koyejo, Ibrahim Alabdulmohsin, Eva Schnider, Krista Opsahl-Ong, Alex Brown, Subhrajit Roy, Diana Mincu, Christina Chen, et al. Maintaining fairness across distribution shift: do we have viable solutions for real-world applications? *arXiv preprint arXiv:2202.01034*, 2022.
- [59] Harvineet Singh, Rina Singh, Vishwali Mhasawade, and Rumi Chunara. Fairness violations and mitigation under covariate shift. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 3–13, 2021.
- [60] Stephen Giguere, Blossom Metevier, Bruno Castro da Silva, Yuriy Brun, Philip Thomas, and Scott Niekum. Fairness guarantees under demographic shift. In *International Conference on Learning Representations*, 2022.
- [61] Jiongli Zhu, Sainyam Galhotra, Nazanin Sabri, and Babak Salimi. Consistent range approximation for fair predictive modeling. *Proceedings of the VLDB Endowment*, 16(11):2925–2938, 2023.
- [62] Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and Vincent YF Tan. Efficient sharpness-aware minimization for improved training of neural networks. *International Conference on Learning Representations*, 2022.
- [63] Maksym Andriushchenko and Nicolas Flammarion. Towards understanding sharpness-aware minimization. In *International Conference on Machine Learning*, pages 639–668. PMLR, 2022.
- [64] Yuzhe Lu, Zhenlin Wang, Runtian Zhai, Soheil Kolouri, Joseph Campbell, and Katia Sycara. Predicting out-of-distribution error with confidence optimal transport. *arXiv preprint arXiv:2302.05018*, 2023.
- [65] Yaodong Yu, Zitong Yang, Alexander Wei, Yi Ma, and Jacob Steinhardt. Predicting out-of-distribution error with the projection norm. In *International Conference on Machine Learning*, pages 25721–25746. PMLR, 2022.