



## OPEN ACCESS

## EDITED BY

Qi Li,  
Iowa State University, United States

## REVIEWED BY

Chaochao Chen,  
Zhejiang University, China  
Mengdi Huai,  
Iowa State University, United States

## \*CORRESPONDENCE

Kai Shu  
✉ kshu@iit.edu

## SPECIALTY SECTION

This article was submitted to  
Data Mining and Management,  
a section of the journal  
Frontiers in Big Data

RECEIVED 20 September 2022

ACCEPTED 07 December 2022

PUBLISHED 04 January 2023

## CITATION

Liang Y, Chen C, Tian T and Shu K  
(2023) Fair classification via domain  
adaptation: A dual adversarial learning  
approach. *Front. Big Data* 5:1049565.  
doi: 10.3389/fdata.2022.1049565

## COPYRIGHT

© 2023 Liang, Chen, Tian and Shu.  
This is an open-access article  
distributed under the terms of the  
[Creative Commons Attribution License  
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or  
reproduction in other forums is  
permitted, provided the original  
author(s) and the copyright owner(s)  
are credited and that the original  
publication in this journal is cited, in  
accordance with accepted academic  
practice. No use, distribution or  
reproduction is permitted which does  
not comply with these terms.

# Fair classification *via* domain adaptation: A dual adversarial learning approach

Yueqing Liang<sup>1</sup>, Canyu Chen<sup>1</sup>, Tian Tian<sup>2</sup> and Kai Shu<sup>1\*</sup>

<sup>1</sup>Department of Computer Science, Illinois Institute of Technology, Chicago, IL, United States,

<sup>2</sup>Stuart School of Business, Illinois Institute of Technology, Chicago, IL, United States

Modern machine learning (ML) models are becoming increasingly popular and are widely used in decision-making systems. However, studies have shown critical issues of ML discrimination and unfairness, which hinder their adoption on high-stake applications. Recent research on fair classifiers has drawn significant attention to developing effective algorithms to achieve fairness and good classification performance. Despite the great success of these fairness-aware machine learning models, most of the existing models require sensitive attributes to pre-process the data, regularize the model learning or post-process the prediction to have fair predictions. However, sensitive attributes are often incomplete or even unavailable due to privacy, legal or regulation restrictions. Though we lack the sensitive attribute for training a fair model in the target domain, there might exist a similar domain that has sensitive attributes. Thus, it is important to exploit auxiliary information from a similar domain to help improve fair classification in the target domain. Therefore, in this paper, we study a novel problem of exploring domain adaptation for fair classification. We propose a new framework that can learn to adapt the sensitive attributes from a source domain for fair classification in the target domain. Extensive experiments on real-world datasets illustrate the effectiveness of the proposed model for fair classification, even when no sensitive attributes are available in the target domain.

## KEYWORDS

fair machine learning, adversarial learning, domain adaptation, trustworthiness, transfer learning

## 1. Introduction

The recent development of machine learning models has been increasingly used for high-stake decision making such as filtering loan applicants (Hamid and Ahmed, 2016), deploying police officers (Mena, 2011), etc. However, a prominent concern is when a learned model has bias against some specific demographic group such as race or gender. For example, a recent report shows that software used by schools to filter student applications may be biased toward a specific race group<sup>1</sup>; and COMPAS (Redmond, 2011), a tool for crime prediction, is shown to be more likely to assign a higher

1 <https://www.fastcompany.com/90342596/schools-are-quietly-turning-to-ai-to-help-pick-who-gets-in-what-could-go-wrong>

risk score to African-American offenders than to Caucasians with the same profile. Bias in algorithms can emanate from unrepresentative or incomplete training data or flawed information that reflects historical inequalities (Du et al., 2020), which can lead to unfair decisions that have a collective and disparate impact on certain groups of people. Therefore, it is important to ensure fairness in machine learning models.

Recent research on fairness machine learning has drawn significant attention to developing effective algorithms to achieve fairness and maintain good prediction performance (Barocas et al., 2017). However, the majority of these models require sensitive attributes to preprocess the data, regularize the model learning or post-process the prediction to have fair predictions. For example, Kamiran and Calders (2012) propose to assign the weights of each training sample differently by reweighting to ensure fairness before model training. In addition, Madras et al. (2018) utilize adversarial training on sensitive attributes and prediction labels for debiasing classification results. However, sensitive attributes are often incomplete or even unavailable due to privacy, legal, or regulation restrictions. For example, by the law of the US Consumer Financial Protection Bureau (CFPB), creditors may not request or collect information about an applicant's race, color, sex, etc.<sup>2</sup>.

Although we lack the sensitive attribute for training a fair model in the target domain, there might exist a similar domain that has sensitive attributes, which paves us a way to adapt the fairness from a related source domain to the target. For example, as shown in Figure 1, when historical features and risk labels are available for credit risk assessment, while user sensitive attributes (i.e., gender) are unavailable when expanding into new markets, we may learn models from its existing market data (source) to score people in the new market (target).

Though there are extensive works on domain adaptation, they are overwhelmingly on classification or knowledge transfer (Tan et al., 2018; Zhuang et al., 2020); while the work on fairness classification through domain adaptation is rather limited.

Very few recent works study a *shallow scenario of transfer* between different groups of sensitive attributes in the same domain for fair classification (Coston et al., 2019; Schumann et al., 2019). We study a general cross-domain fair classification problem in terms of a specific sensitive attribute, which is not explored yet to our best knowledge.

Therefore, we propose a dual-adversarial learning framework to *learn to adapt* sensitive attributes for fair classification in the target domain. In essence, we investigate the following challenges: (1) how to adapt the sensitive attributes to the target domain by transferring knowledge from the source domain; and (2) how to predict the labels accurately

and satisfy fairness criteria. It is non-trivial since there are domain discrepancies across domains and no anchors are explicitly available for knowledge transfer. Our solutions to these challenges result in a novel framework called FAIRDA for fairness classification *via* domain adaptation. Our main contributions are summarized as follows:

- We study a novel problem of fair classification *via* domain adaptation.
- We propose a new framework FAIRDA that simultaneously transfers knowledge to adapt sensitive attributes and learns a fair classifier in the target domain with a dual adversarial learning approach.
- We conduct theoretical analysis demonstrating that fairness in the target domain can be achieved with estimated sensitive attributes derived from the source domain.
- We perform extensive experiments on real-world datasets to demonstrate the effectiveness of the proposed method for fair classification without sensitive attributes.

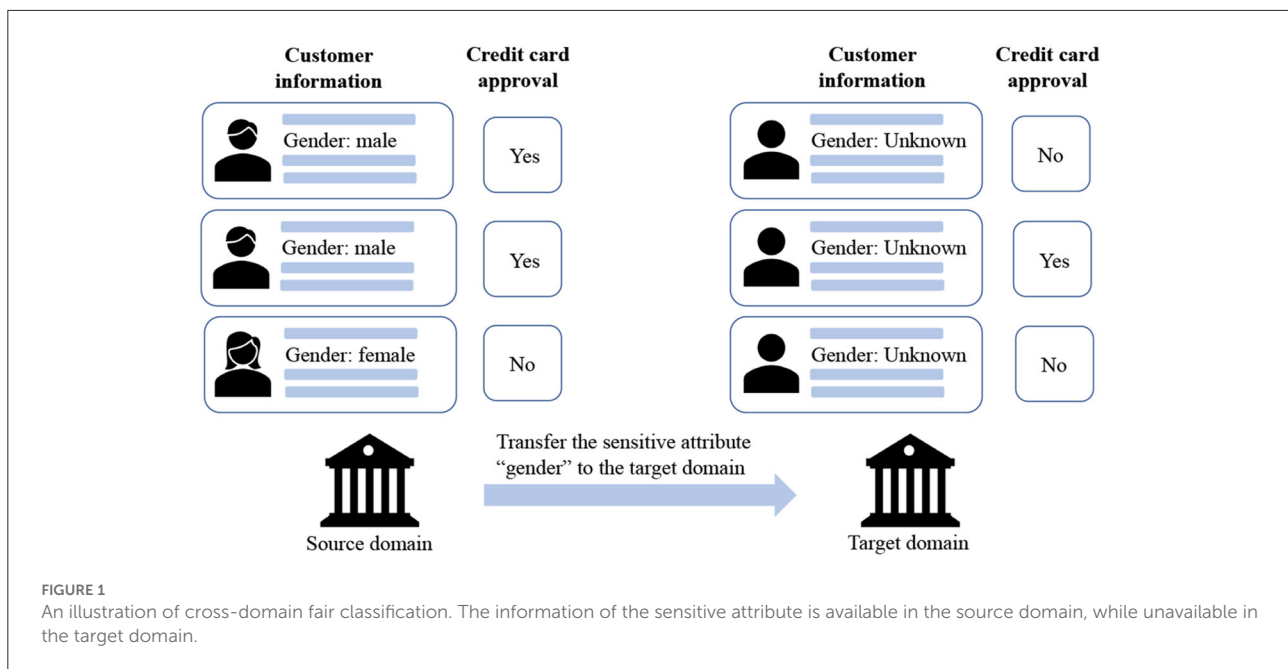
## 2. Related work

In this section, we briefly review the related works on fairness machine learning and deep domain adaptation.

### 2.1. Fairness in machine learning

Recent research on fairness in machine learning has drawn significant attention to developing effective algorithms to achieve fairness and maintain good prediction performance, which generally focus on individual fairness (Kang et al., 2020) or group fairness (Hardt et al., 2016; Zhang et al., 2017). Individual fairness requires the model to give similar predictions to similar individuals (Caton and Haas, 2020; Cheng et al., 2021). In group fairness, similar predictions are desired among multiple groups categorized by a specific sensitive attribute (e.g., gender). Other niche notions of fairness include subgroup fairness (Kearns et al., 2018) and Max-Min fairness (Lahoti et al., 2020), which aims to maximize the minimum expected utility across groups. In this work, we focus on group fairness. To improve group fairness, debiasing techniques have been applied at different stages of a machine learning model: (1) *Pre-processing* approaches (Kamiran and Calders, 2012) apply dedicated transformations on the original dataset to remove intrinsic discrimination and obtain unbiased training data prior to modeling; (2) *In-processing* approaches (Bechavod and Ligett, 2017; Agarwal et al., 2018) tackle this issue by incorporating fairness constraints or fairness-related objective functions to the design of machine learning models; and (3) *Post-processing*

<sup>2</sup> <https://www.consumerfinance.gov/rules-policy/regulations/1002/5/>



approaches (Dwork et al., 2018) revise the biased prediction labels by debiasing mechanisms.

Although the aforementioned methods can improve group fairness, they generally require the access of sensitive attributes, which is often infeasible. Very few recent works study fairness with limited sensitive attributes. For example, Zhao et al. (2022) and Gupta et al. (2018) explore and use the related features as proxies of the sensitive attribute to achieve better fairness results. Lahoti et al. (2020) proposes an adversarial reweighted method to achieve the Rawlsian Max-Min fairness objective which aims at improving the accuracy for the worst-case protected group. However, these methods may require domain knowledge to approximate sensitive attributes or are not suitable for ensuring group fairness. In addition, though Madras et al. (2018) mentioned fairness transfer, what they study is the cross-task transfer (Tan et al., 2021), which is different from our setting.

In this paper, we study the novel problem of cross-domain fairness classification which aims to estimate the sensitive attributes of the target domain to achieve the group fairness.

## 2.2. Deep domain adaption

Domain adaptation (Pan and Yang, 2009) aims at mitigating the generalization bottleneck introduced by domain shift. With the rapid growth of deep neural networks, deep domain adaptation has drawn much attention lately. In general, deep domain adaptation methods aim to learn a domain-invariant feature space that can reduce the discrepancy between the source and target domains. This goal is accomplished

either by transforming the features from one domain to be closer to the other domain, or by projecting both domains into a domain-invariant latent space (Shu et al., 2019). For instance, TLDA (Zhuang et al., 2015) is a deep autoencoder-based model for learning domain-invariant representations for classification. Inspired by the idea of Generative Adversarial Network (GAN) (Goodfellow et al., 2014), researchers also propose to perform domain adaptation in an adversarial training paradigm (Ganin et al., 2016; Tzeng et al., 2017; Shu et al., 2019). By exploiting a domain discriminator to distinguish the domain labels while learning deep features to confuse the discriminator, DANN (Ganin et al., 2016) achieves superior domain adaptation performance. ADDA (Tzeng et al., 2017) learns a discriminative representation using labeled source data and then maps the target data to the same space through an adversarial loss. Recently, very few works apply transfer learning techniques for fair classification (Coston et al., 2019; Schumann et al., 2019). Although sensitive attributes are not available in the target domain, there may exist some publicly available datasets which can be used as auxiliary sources. However, these methods mostly consider a shallow scenario of transfer between different sensitive attributes in the same dataset.

In this paper, we propose a new domain adaptive approach based on dual adversarial learning to achieve fair classification in the target domain.

## 3. Problem statement

We first introduce the notations of this paper, and then give the formal problem definition. Let  $\mathcal{D}^1 = \{\mathcal{X}^1, \mathcal{A}^1, \mathcal{Y}^1\}$  denote

the data in the source domain, where  $\mathcal{X}^1, \mathcal{A}^1$ , and  $\mathcal{Y}^1$  represent the set of data samples, sensitive attributes, and corresponding labels. Let  $\mathcal{D}^2 = \{\mathcal{X}^2, \mathcal{Y}^2\}$  be the target domain, where the sensitive attributes of the target domain are unknown. The data distributions in the source and target domains are similar but different with the domain discrepancy. Following existing work on fair classification (Barocas et al., 2017; Mehrabi et al., 2019), we evaluated fairness performance using metrics such as equal opportunity and demographic parity. Without loss of generality, we consider binary classification. Equal opportunity requires that the probability of positive instances with arbitrary sensitive attributes  $A$  being assigned to a positive outcome are equal:  $\mathbb{E}(\hat{Y} | A = a, Y = 1) = \mathbb{E}(\hat{Y} | A = b, Y = 1)$ , where  $\hat{Y}$  is predicted label. Demographic parity requires the behavior of the prediction model to be fair to different sensitive groups. Concretely, it requires that the positive rate across sensitive attributes are equal:  $\mathbb{E}(\hat{Y} | A = a) = \mathbb{E}(\hat{Y} | A = b), \forall a, b$ . The problem of fair classification with domain adaptation is formally defined as follows:

**Problem Statement:** Given the training data  $\mathcal{D}^1$  and  $\mathcal{D}^2$  from the source and target domain, learn an effective classifier for the target domain while satisfying the fairness criteria such as demographic parity.

## 4. FAIRDA: Fair classification with domain adaptation

In this section, we present the details of the proposed framework for fair classification with domain adaptation. As shown in Figure 2, our framework consists of two major modules: (1) an adversarial domain adaptation module that estimates sensitive attributes for the target domain; and (2) an adversarial debiasing module to learn a fair classifier in the target domain.

Specifically, first, the adversarial domain adaptation module contains a sensitive attribute predictor  $f_{\theta_A}$  that describes the modeling of predicting sensitive attributes and a domain classifier  $f_{\theta_d}$  that illustrates the process of transferring knowledge to the target domain for estimating sensitive attributes. In addition, the adversarial debiasing module consists of a label predictor  $f_{\theta_Y}$  that models the label classification in the target domain, and a bias predictor  $f_{\theta_a}$  to illustrate the process of differentiating the estimated sensitive attributes from the target domain data.

### 4.1. Estimating target sensitive attributes

Since the sensitive attributes in the target domain are unknown, it is necessary to estimate the target sensitive

attributes to build fair classifiers. Recent advancement in unsupervised domain adaptation has shown promising results to adapt a classifier on a source domain to the unlabeled data to the target domain (Ganin and Lempitsky, 2015). Therefore, we propose to investigate an unsupervised domain adaptation framework to infer the sensitive attributes in the target domain. Specifically, to minimize the prediction error of sensitive attributes in the source domain, with the following objective function:

$$\min_{\theta_A} \mathcal{L}_A = \mathbb{E}_{(X^1, A^1) \in \mathcal{D}^1} \ell(A^1, f_{\theta_A}(h_1(X^1))) \tag{1}$$

Where  $f_{\theta_A}$  is used to predict sensitive attributes,  $\ell$  denotes the loss function to minimize prediction error, such as cross-entropy loss, and the embedding function  $h_1(\cdot)$  is to learn the representation for predicting sensitive attributes. To ensure that  $f_{\theta_A}$  trained on  $\mathcal{D}^1$  can well estimate the sensitive attributes of  $\mathcal{D}^2$ . We further introduce a domain classifier  $f_{\theta_d}$  to make the representation of source and target domains in the same feature space. Specifically,  $f_{\theta_d}$  aims to differentiate if a representation is from the source or the target; while  $h_1(\cdot)$  aim to learn a domain-invariant representation that can fool  $f_{\theta_d}$ , i.e.,

$$\min_{\theta_{h_1}} \max_{\theta_d} \mathcal{L}_d = \mathbb{E}_{X^1 \in \mathcal{D}^1} [\log(f_{\theta_d}(h_1(X^1)))] + \mathbb{E}_{X^2 \in \mathcal{D}^2} [\log(1 - (f_{\theta_d}(h_1(X^2))))] \tag{2}$$

Where  $\theta_d$  are the parameters of the domain classifier. The overall objective function for this module is a min-max game between the sensitive attributes predictor  $f_{\theta_A}$  and the domain classifier  $f_{\theta_d}$  as follows:

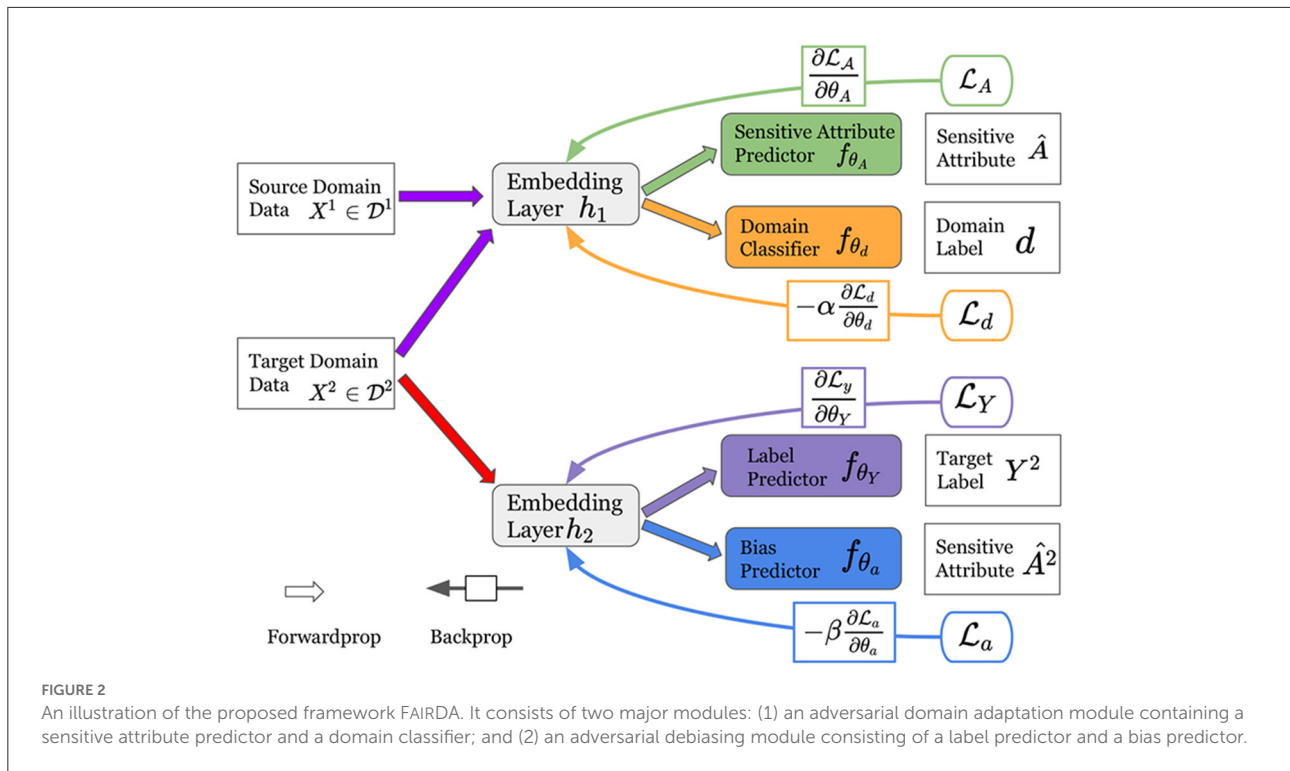
$$\min_{\theta_A, \theta_{h_1}} \max_{\theta_d} \mathcal{L}_1 = \mathcal{L}_A - \alpha \mathcal{L}_d \tag{3}$$

Where  $\alpha$  controls the importance of the adversarial learning of the domain classifier. To estimate the sensitive attributes in the target domain, we will utilize the following function:  $\hat{A}^2 = f_{\theta_A}(h_1(X^2))$ .

### 4.2. Adversarial debiasing for fair classification

To learn a fair classifier for the target domain, we will leverage the derived sensitive attributes  $\hat{A}^2$  and the labels  $Y^2$  in the target domain. Specifically, we aim to learn a representation that can predict the labels accurately, while being irrelevant to sensitive attributes. To this end, we propose to feed the representation into a label predictor and a bias predictor. Specifically, the label predictor is to minimize the predicted errors of the labels with the following objective function:

$$\min_{\theta_Y} \mathcal{L}_Y = \mathbb{E}_{(X^2, Y^2) \in \mathcal{D}^2} \ell(Y^2, f_{\theta_Y}(h_2(X^2))) \tag{4}$$



**FIGURE 2**  
An illustration of the proposed framework FAIRDA. It consists of two major modules: (1) an adversarial domain adaptation module containing a sensitive attribute predictor and a domain classifier; and (2) an adversarial debiasing module consisting of a label predictor and a bias predictor.

Where  $f_{\theta_Y}$  is to predict the labels,  $h_2(\cdot)$  is an embedding layer to encode the features into a latent representation space, and  $\ell$  is a cross-entropy loss. Furthermore, to learn fair representations and make fair predictions in the target domain, we incorporate a bias predictor to predict the sensitive attributes; while  $h_2(\cdot)$  trying to learn the representation that can fool  $f_{\theta_Y}$ :

$$\min_{\theta_{h_2}} \max_{\theta_a} \mathcal{L}_a = \mathbb{E}_{X^2 \sim p(X^2 | \hat{A}^2=1)} [\log(f_{\theta_a}(h_2(X^2)))] + \mathbb{E}_{X^2 \sim p(X^2 | \hat{A}^2=0)} [\log(1 - f_{\theta_a}(h_2(X^2)))] \quad (5)$$

Where  $\theta_a$  are the parameters for the adversary to predict sensitive attributes. Finally, the overall objective function of adversarial debiasing for fair classification is a min-max function:

$$\min_{\theta_{h_2}, \theta_Y} \max_{\theta_a} \mathcal{L}_2 = \mathcal{L}_Y - \beta \mathcal{L}_a \quad (6)$$

Where  $\beta$  controls the importance of the bias predictor.

### 4.3. The proposed framework: FAIRDA

We have introduced how we can estimate sensitive attributes by adversarial domain adaptation, and how to ensure fair classification with adversarial debiasing. We integrate the

components together and the overall objective function of the dual adversarial learning is as follows:

$$\min_{\theta_{h_1}, \theta_{h_2}, \theta_Y, \theta_A} \max_{\theta_d, \theta_a} \mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 \quad (7)$$

The parameters in the objective are learned through RMSProp, which is an adaptive learning rate method that divides the learning rate by an exponentially decaying average of squared gradients. We choose RMSProp as the optimizer because it is a popular and effective method to determine the learning rate abortively which is widely used for training adversarial neural networks (Dou et al., 2019; Li et al., 2022; Zhou and Pan, 2022). Adam is another widely adopted optimizer that extends RMSProp with momentum terms, however, the momentum terms may make Adam unstable (Mao et al., 2017; Luo et al., 2018; Clavijo et al., 2021). We will also prioritize the training of  $f_{\theta_A}$  to ensure a good estimation of  $\hat{A}^2$ . Next, we will conduct a theoretical analysis of the dual adversarial learning for fairness guarantee.

## 5. A theoretical analysis on fairness guarantee

In this section, we perform a theoretical analysis of the fairness guarantee under the proposed framework FAIRDA with key assumptions. The model essentially contains two modules: (1) the adversarial domain adaptation for estimating sensitive



attributes ( $\mathcal{L}_1$  in Equation 3); and (2) the adversarial debiasing for learning a fair classifier ( $\mathcal{L}_2$  in Equation 6). The performance of the second module is relying on the output (i.e.,  $\hat{A}^2$ ) of the first module.

To understand the theoretical guarantee for the first module, we follow the recent work on analyzing the conventional unsupervised domain adaptation (Zhang et al., 2019; Liu et al., 2021). The majority of them consider the model effectiveness under different types of domain shifts such as label shift, covariate shift, conditional shift, etc. (Kouw and Loog, 2018; Liu et al., 2021). When analyzing domain adaptation models, researchers usually consider the existence of one shift while assuming other shifts are invariant across domains (Kouw and Loog, 2018). The existence of label shift is a common setting of unsupervised domain adaptation (Chen et al., 2018; Lipton et al., 2018; Azizzadenesheli et al., 2019; Wu et al., 2019). We follow the above papers to make the same assumption on label shift, which in our scenario is *sensitive attribute shift* across two domains, i.e., the prior distribution changes as  $p(A^1) \neq p(A^2)$ . In addition, we assume that the other shifts are invariant across domains, for example,  $p(X^1|A^1) = p(X^2|A^2)$ . Under the assumption of the sensitive attribute shift, we can derive the risk in the target domain as Kouw and Loog (2018):

$$R^2(h_1, \theta_A) = \sum_{A \in \mathcal{A}} \int_{\mathcal{X}} \ell(A, f_{\theta_A}(h_1(X))) \frac{p(X^2|A^2)p(A^2)}{p(X^1|A^1)p(A^1)} p(X^1, A^1) dX \tag{8}$$

$$= \sum_{A \in \mathcal{A}} \int_{\mathcal{X}} \ell(A, f_{\theta_A}(h_1(X))) \frac{p(A^2)}{p(A^1)} p(X^1, A^1) dX$$

Where the ratio  $p(A^2)/p(A^1)$  represents the change in the proportions in the sensitive attribute. Since we do not have samples with sensitive attributes from the target domain, we can use the samples drawn from the source distribution to estimate the target sensitive attributes distribution with mean matching (Gretton et al., 2009) by minimizing the following function  $\|M^1 p(A^1) - \mu^2\|_2^2$ : where  $M^1$  is the vector empirical sample means from the source domain, i.e.,  $[\mu^1(f_{\theta_A}(h_1(X^1))|A^1 = 0), \mu^1(f_{\theta_A}(h_1(X^1))|A^1 = 1)]$ .  $\mu^2$  is the encoded feature means for the target. We will incorporate the above strategy for estimating the target sensitive attribute proportions with gradient descent during the adversarial training (Li et al., 2019).

The noise induced for estimating the target sensitive attributes  $\hat{A}^2$  is nonnegligible, which may influence the adversarial debiasing in the second module. Next, we theoretically show that under mild conditions, we can satisfy fairness metrics such as demographic parity. First, we can prove that the global optimum of  $\mathcal{L}_2$  can be achieved if and only if  $p(Z^2|\hat{A}^2 = 1) = p(Z^2|\hat{A}^2 = 0)$ , where  $Z^2 = h_2(X^2)$  is the representation of  $X^2$  according to the Proposition 1.

in Goodfellow et al. (2014). Next, we introduce the following theorem under two reasonable assumptions:

Theorem 1. Let  $\hat{Y}^2$  denote the predicted label of the target domain, if

- (1) The estimated sensitive attributes  $\hat{A}^2$  and the representation of  $X^2$  are independent conditioned on the true sensitive attributes, i.e.,

$$p(\hat{A}^2, Z^2|A^2) = p(\hat{A}^2|A^2)p(Z^2|A^2);$$

- (2) The estimated sensitive attributes are not random, i.e.,

$$p(A^2 = 1|\hat{A}^2 = 1) \neq p(A^2 = 0|\hat{A}^2 = 1).$$

If  $\mathcal{L}_2$  reaches the global optimum, the label prediction  $f_{\theta_Y}$  will achieve demographic parity, i.e.,

$$p(\hat{Y}^2|A^2 = 0) = p(\hat{Y}^2|A^2 = 1).$$

We first explain the two assumptions:

- (1) Since we use two separate embedding layers  $h_1(\cdot)$  and  $h_2(\cdot)$  to predict the target sensitive attributes, and learn the latent presentation, it generally holds that  $\hat{A}^2$  is independent of the representation of  $X^2$ , i.e.,  $p(\hat{A}^2, Z^2|A^2) = p(\hat{A}^2|A^2)p(Z^2|A^2)$ ;
- (2) Since we are using adversarial learning to learn an effective estimator  $f_{\theta_A}$  for sensitive attributes, it is reasonable to assume that it does not produce random prediction results.

We prove Theorem 1 as follows:

Since  $p(\hat{A}^2, Z^2|A^2) = p(\hat{A}^2|A^2)p(Z^2|A^2)$ , we have  $p(Z^2|\hat{A}^2, A^2) = p(Z^2|A^2)$ . In addition, when the algorithm converges, we have  $p(Z^2|\hat{A}^2 = 1) = p(Z^2|\hat{A}^2 = 0)$ , which is equivalent with

$$\sum_{A^2} p(Z^2, A^2|\hat{A}^2 = 1) = \sum_{A^2} p(Z^2, A^2|\hat{A}^2 = 0). \text{ Therefore,}$$

$$\sum_{A^2} p(Z^2|A^2)p(A^2|\hat{A}^2 = 1) = \sum_{A^2} p(Z^2|A^2)p(A^2|\hat{A}^2 = 0) \tag{9}$$

Based on the above equation, we can get,

$$\frac{p(Z^2|A^2 = 1)}{p(Z^2|A^2 = 0)} = \frac{p(A^2 = 0|\hat{A}^2 = 1) - p(A^2 = 0|\hat{A}^2 = 0)}{p(A^2 = 1|\hat{A}^2 = 0) - p(A^2 = 1|\hat{A}^2 = 1)} = 1 \tag{10}$$

Which shows that a global minimum achieves, i.e.,  $p(Z^2|A^2 = 1) = p(Z^2|A^2 = 0)$ . Since  $\hat{Y}^2 = f_{\theta_Y}(Z^2)$ , we can get  $p(\hat{Y}^2|A^2 = 1) = p(\hat{Y}^2|A^2 = 0)$ , which is the demographic parity.

## 6. Experiments

In this section, we present the experiments to evaluate the effectiveness of FAIRDA. We aim to answer the following research questions (RQs):

- **RQ1:** Can FAIRDA obtain fair predictions without accessing sensitive attributes in the target domain?
- **RQ2:** How can we transfer fairness knowledge from the source domain while effectively using it to regularize the target's prediction?
- **RQ3:** How would different choices of the weights of the two adversarial components impact the performance of FAIRDA?

## 6.1. Datasets

We conduct experiments on four publicly available benchmark datasets for fair classification: COMPAS (Angwin et al., 2016), Adult (Dua and Graff, 2017), Toxicity (Dixon et al., 2018), and CelebA (Liu et al., 2015).

- **COMPAS**<sup>3</sup>: This dataset describes the task of predicting the recidivism of individuals in the U.S. Both “sex” and “race” could be the sensitive attributes of this dataset (Le Quy et al., 2022).
- **ADULT**<sup>4</sup>: This dataset contains records of personal yearly income, and the label is whether the income of a specific individual exceeds 50k or not. Following Zhao et al. (2022), we choose “sex” as the sensitive attribute.
- **Toxicity**<sup>5</sup>: This dataset collects comments with labels indicating whether each comment is toxic or not. Following Chuang and Mroueh (2021), we choose “race” as the sensitive attribute.
- **CelebA**<sup>6</sup>: This dataset contains celebrity images, where each image has 40 human-labeled binary attributes. Following Chuang and Mroueh (2021), we choose the attribute “attractive” as our prediction task and “sex” as the sensitive attribute.

The evaluation of all models is measured by the performance on the target dataset. We split train : eval : test set for the target dataset as 0.5 : 0.25 : 0.25. For the source dataset, because we do not evaluate on it, we split train : eval set as 0.6 : 0.4. It is worth noting that for the aforementioned datasets, we need to define a filter to derive source and target domains for empirical studies. Our general principle is to choose a feature that is reasonable in the specific scenario. For example, in COMPAS, we use age as the filter due to the prior knowledge that individuals in different age groups have different preferences to disclose their sex and race information. For ADULT, we consider the sex of each individual to be the sensitive attribute, and define two filters based on the `working class` and `country` to derive

<sup>3</sup> <https://github.com/propublica/compas-analysis>

<sup>4</sup> <https://archive.ics.uci.edu/ml/machine-learning-databases/adult/>

<sup>5</sup> <https://www.kaggle.com/c/jigsaw-unintended-bias-in-toxicity-classification>

<sup>6</sup> <http://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>

the source and target domains. The statistics of the datasets are shown in Table 1.

## 6.2. Experimental settings

### 6.2.1. Baselines

Since there is no existing work on cross-domain fair classification, we compare our proposed FAIRDA with the following representative methods in fair classification without sensitive attributes.

- **Vanilla:** This method directly trains a base classifier without explicitly regularizing on the sensitive attributes.
- **ARL** (Lahoti et al., 2020): This method utilizes reweighting on under-represented regions detected by adversarial model to alleviate bias. It focuses on improving Max-Min fairness rather than group fairness in our setting. However, it is an important work in fair machine learning. For the integrity of our experiment, we still involve it as one of our baselines.
- **KSMOTE** (Yan et al., 2020): It first derives pseudo groups, and then uses them to design regularization to ensure fairness.
- **FairRF** (Zhao et al., 2022): It optimizes the prediction fairness without sensitive attribute but with some available features that are known to be correlated with the sensitive attribute.

For KSMOTE<sup>7</sup> and FairRF<sup>8</sup>, we directly use the code provided by the authors. For all other approaches, we adopt a three-layer multi-layer-perceptron (MLP) as an example of the base classifier, and set the two hidden dimensions of MLP as 64 and 32, and use RMSProp optimizer with 0.001 as the initial learning rate. We have similar observations for other types of classifiers such as Logistic Regression and SVM.

### 6.2.2. Evaluation metrics

Following existing work on fair models, we measure the classification performance with Accuracy (ACC) and F1, and the fairness performance based on **Demographic Parity** and **Equal Opportunity** (Mehrabi et al., 2019). We consider the scenario when the sensitive attributes and labels are binary, which can be naturally extended to more general cases.

- **Demographic Parity:** A classifier is considered to be fair if the prediction  $\hat{Y}$  is independent of the sensitive attribute  $A$ . In other words, demographic parity requires each demographic group has the same chance for a positive outcome:  $\mathbb{E}(\hat{Y}|A = 1) = \mathbb{E}(\hat{Y}|A = 0)$ . We will

<sup>7</sup> <https://imbalanced-learn.org/stable/>

<sup>8</sup> <https://github.com/TianxiangZhao/fairlearn>

TABLE 1 Statistics of datasets.

Exp.	Ori. data	SA	Filter	Dataset	Domain	SA avail.	# Instance
1	COMPAS	Sex	age < 24	COMPAS 1	Source	Yes	881
			age ≥ 24	COMPAS 2	Target	No	4,397
2	COMPAS	Race	age < 24	COMPAS 3	Source	Yes	881
			age ≥ 24	COMPAS 4	Target	No	4,397
3	ADULT	Sex	work_class ≠ private	ADULT 1	Source	Yes	11,915
			work_class = private	ADULT 2	Target	No	33,307
4	ADULT	Sex	country = US	ADULT 3	Source	Yes	41,292
			country ≠ US	ADULT 4	Target	No	3,930
5	Toxicity	Race	Created date before 2017-7-1	Toxic 1	Source	Yes	11,010
			Created date after 2017-7-1	Toxic 2	Target	No	7,340
6	CelebA	Sex	Young	CelebA 1	Source	Yes	121,560
			Not young	CelebA 2	Target	No	81,040

SA refers to sensitive attributes. The filter defines the rule to group the original data to the datasets of the source and target domains.

report the difference of each sensitive group's demographic parity ( $\Delta_{DP}$ ):

$$\Delta_{DP} = |\mathbb{E}(\hat{Y}|A = 1) - \mathbb{E}(\hat{Y}|A = 0)| \quad (11)$$

- **Equal Opportunity:** Equal opportunity considers one more condition than demographic parity. A classifier is said to be fair if the prediction  $\hat{Y}$  of positive instances is independent of the sensitive attributes. Concretely, it requires the true positive rates of different sensitive groups to be equal:  $\mathbb{E}(\hat{Y}|A = 1, Y = 1) = \mathbb{E}(\hat{Y}|A = 0, Y = 1)$ . Similarly, in the experiments, we report the difference of each sensitive group's equal opportunity ( $\Delta_{EO}$ ):

$$\Delta_{EO} = |\mathbb{E}(\hat{Y}|A = 1, Y = 1) - \mathbb{E}(\hat{Y}|A = 0, Y = 1)| \quad (12)$$

Note that demographic parity and equal opportunity measure fairness performance in different ways, and the smaller the values are, the better the performance of fairness.

### 6.3. Fairness performances (RQ1)

To answer **RQ1**, we compare the proposed framework with the aforementioned baselines in terms of prediction and fairness. All experiments are conducted 5 times and the average performances and standard deviations are reported in Tables 2–4 in terms of ACC, F1,  $\Delta_{EO}$ ,  $\Delta_{DP}$ . We have the following observations:

- Adapting sensitive attribute information from a similar source domain could help achieve a better fairness performance in the target domain. For example, compared with ARL, KSMOTE and FairRF, which leverage the

implicit sensitive information within the target domain, overall FAIRDA performs better in terms of both fairness metrics  $\Delta_{DP}$  and  $\Delta_{EO}$  across all six experiments.

- In general, the proposed framework FAIRDA can achieve a better performance of fairness, while remains comparatively good prediction performance. For example, in Exp. 6 of CelebA dataset, the performance of  $\Delta_{DP}$  of FAIRDA increases 21.2% compared to Vanilla, on which no explicit fairness regularization applied, and the performance of ACC only drops 0.1%.
- We conduct experiments on various source and target datasets. We observe that no matter whether the number of instances in the source domain is greater or smaller than the target one, the proposed FAIRDA outperforms baselines consistently, which demonstrates the intuition of sensitive adaptation is effective and the proposed FAIRDA is robust.

### 6.4. Ablation study (RQ2)

In this section, we aim to analyse the effectiveness of each component in the proposed FAIRDA framework. As shown in Section (4), FAIRDA contains two adversarial components. The first one is designed for making better predictions on target domain's sensitive attribute with the knowledge learned from the source domain. The second one is designed for debiasing with the sensitive attribute predicted by the domain adaptive classifier. In order to answer **RQ2**, we investigate the effects of these components by defining two variants of FAIRDA:

- **FAIRDA w/o DA:** This is a variant of FAIRDA without domain adversary on the source domain. It first trains a classifier with the source domain's sensitive attribute and directly applies it to predict the target domain's sensitive attribute.



TABLE 2 Fairness performance comparison in the COMPAS dataset.

Datasets	Exp.	Methods	ACC	F1	$\Delta_{DP}$	$\Delta_{EO}$
COMPAS	1	Vanilla	0.677 ± 0.015	0.639 ± 0.018	0.128 ± 0.015	0.137 ± 0.005
		ARL	0.672 ± 0.008	0.609 ± 0.032	0.109 ± 0.038	0.119 ± 0.057
		KSMOTE	0.632 ± 0.012	0.525 ± 0.018	0.136 ± 0.017	0.146 ± 0.018
		FairRF	0.641 ± 0.013	0.527 ± 0.064	0.086 ± 0.020	0.108 ± 0.048
		FAIRDA	0.661 ± 0.023	0.639 ± 0.016	<b>0.050</b> ± 0.014	<b>0.042</b> ± 0.008
	2	Vanilla	0.675 ± 0.014	0.644 ± 0.023	0.141 ± 0.003	0.143 ± 0.007
		ARL	0.663 ± 0.001	0.598 ± 0.036	0.165 ± 0.015	0.171 ± 0.026
		KSMOTE	0.628 ± 0.010	0.615 ± 0.020	0.150 ± 0.012	0.148 ± 0.013
		FairRF	0.631 ± 0.007	0.603 ± 0.012	0.103 ± 0.019	0.105 ± 0.022
		FAIRDA	0.611 ± 0.021	0.658 ± 0.006	<b>0.085</b> ± 0.022	<b>0.096</b> ± 0.012

The numbers in bold indicate the best performances.

TABLE 3 Fairness performance comparison in the ADULT dataset.

Datasets	Exp.	Methods	ACC	F1	$\Delta_{DP}$	$\Delta_{EO}$
ADULT	3	Vanilla	0.857 ± 0.003	0.655 ± 0.009	0.100 ± 0.005	0.099 ± 0.014
		ARL	0.837 ± 0.006	0.668 ± 0.005	0.142 ± 0.019	0.122 ± 0.016
		KSMOTE	0.840 ± 0.002	0.511 ± 0.010	<b>0.072</b> ± 0.008	0.138 ± 0.055
		FairRF	0.848 ± 0.003	0.594 ± 0.005	0.083 ± 0.001	0.092 ± 0.015
		FAIRDA	0.845 ± 0.011	0.604 ± 0.008	0.080 ± 0.001	<b>0.073</b> ± 0.016
	4	Vanilla	0.839 ± 0.001	0.658 ± 0.001	0.107 ± 0.003	0.093 ± 0.009
		ARL	0.816 ± 0.015	0.673 ± 0.004	0.142 ± 0.025	0.079 ± 0.023
		KSMOTE	0.824 ± 0.004	0.589 ± 0.040	0.102 ± 0.016	0.140 ± 0.017
		FairRF	0.829 ± 0.003	0.611 ± 0.011	0.087 ± 0.006	<b>0.068</b> ± 0.013
		FAIRDA	0.825 ± 0.004	0.599 ± 0.017	<b>0.085</b> ± 0.004	0.070 ± 0.003

The numbers in bold indicate the best performances.

TABLE 4 Fairness performance comparison in the Toxicity and CelebA datasets.

Datasets	Exp.	Methods	ACC	F1	$\Delta_{DP}$	$\Delta_{EO}$
Toxicity	5	Vanilla	0.761 ± 0.002	0.550 ± 0.023	0.108 ± 0.017	0.118 ± 0.029
		ARL	0.759 ± 0.002	0.558 ± 0.022	0.113 ± 0.016	0.115 ± 0.025
		KSMOTE	0.619 ± 0.116	0.428 ± 0.040	0.104 ± 0.028	0.095 ± 0.013
		FairRF	0.705 ± 0.024	0.475 ± 0.097	0.099 ± 0.032	0.113 ± 0.006
		FAIRDA	0.739 ± 0.008	0.501 ± 0.011	<b>0.085</b> ± 0.012	<b>0.069</b> ± 0.048
CelebA	6	Vanilla	0.864 ± 0.001	0.443 ± 0.014	0.132 ± 0.010	0.237 ± 0.014
		ARL	0.801 ± 0.009	0.485 ± 0.009	0.291 ± 0.012	0.359 ± 0.020
		KSMOTE	0.859 ± 0.000	0.359 ± 0.038	0.124 ± 0.020	0.226 ± 0.030
		FairRF	0.864 ± 0.001	0.405 ± 0.024	0.115 ± 0.016	0.213 ± 0.026
		FAIRDA	0.863 ± 0.003	0.390 ± 0.030	<b>0.104</b> ± 0.007	<b>0.189</b> ± 0.017

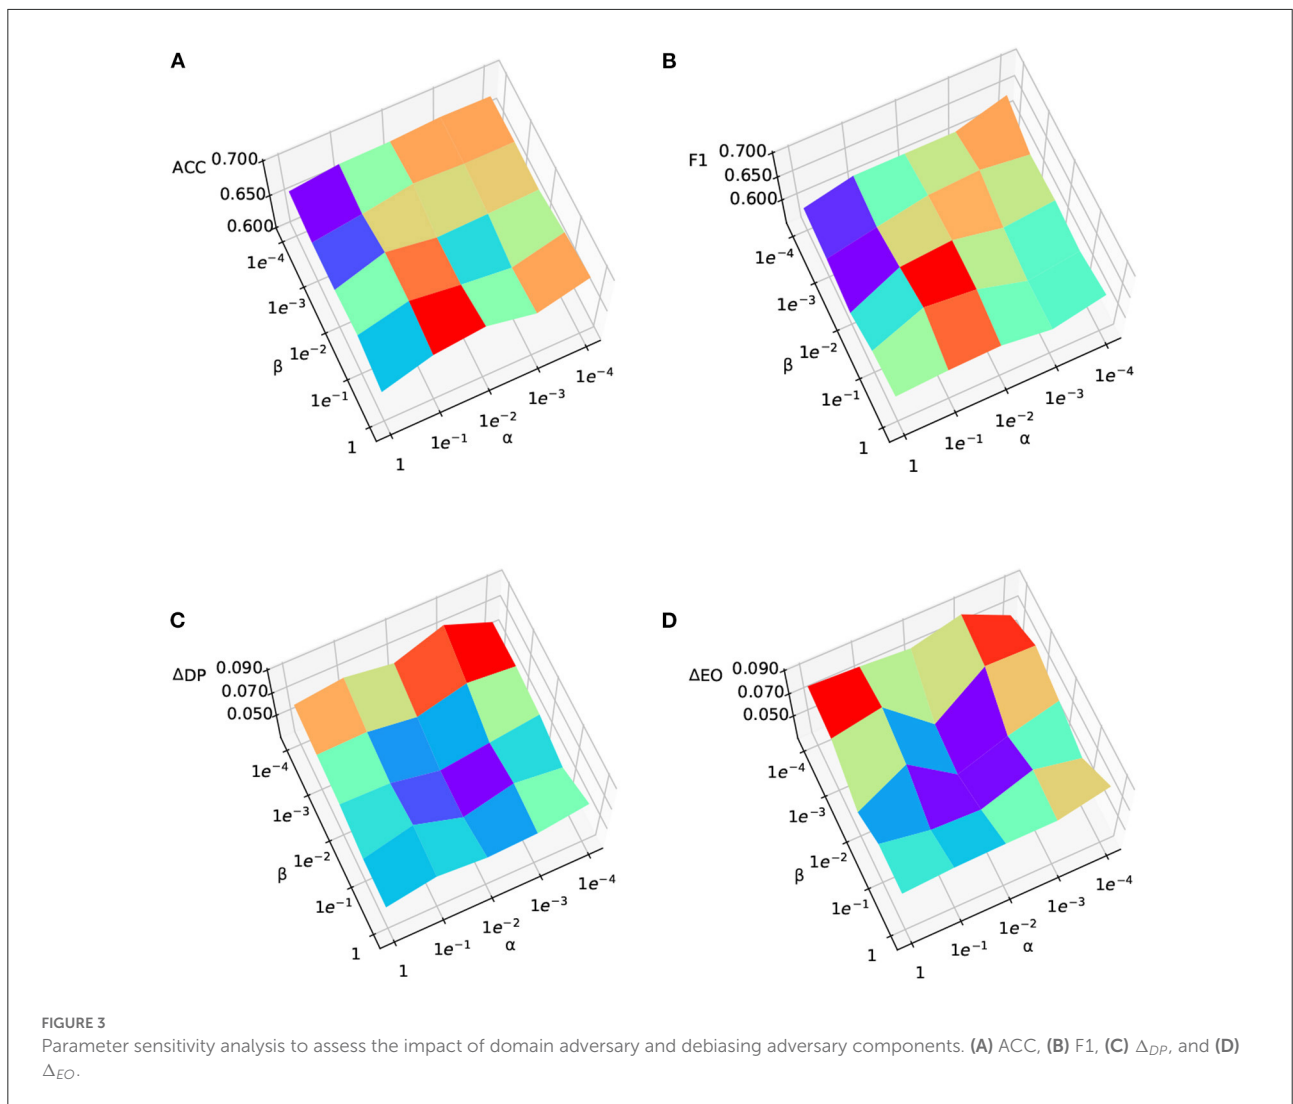
The numbers in bold indicate the best performances.

- FAIRDA *w/o* Debiasing: This is a variant of FAIRDA without debiasing adversary on the target domain. In short, it is equivalent to Vanilla.

Table 5 reports the average performances with standard deviations for each method. We could make the following observations:

TABLE 5 Assessing the effectiveness of each component in FAIRDA.

Variants	ACC	F1	$\Delta_{DP}$	$\Delta_{EO}$
FAIRDA w/o DA	$0.660 \pm 0.025$	$0.624 \pm 0.042$	$0.062 \pm 0.014$	$0.077 \pm 0.007$
FAIRDA w/o Debiasing	$0.677 \pm 0.015$	$0.639 \pm 0.018$	$0.128 \pm 0.015$	$0.137 \pm 0.005$
FAIRDA	$0.661 \pm 0.023$	$0.639 \pm 0.016$	$0.050 \pm 0.014$	$0.042 \pm 0.008$



- When we eliminate the domain adversary for adapting the sensitive attributes information (i.e., FAIRDA w/o DA), both  $\Delta_{DP}$  and  $\Delta_{EO}$  increase, which means the fairness performances are reduced. In addition, the classification performances also decline compared with FAIRDA. This suggests that the adaptation is able to make better estimations for the target sensitive attributes, which contributes to better fairness performances.

- When we eliminate the debiasing adversary (i.e., FAIRDA w/o Debiasing), there is no fairness regularization applied on the target domain. Therefore, it is equivalent to Vanilla. We could observe that the fairness performances drop significantly, which suggests that adding the debiasing adversary to FAIRDA is indispensable. In addition, even with the adversarial debiasing, FAIRDA can still achieve comparatively good F1 compared with FAIRDA w/o

Debiasing. This enhances the intuition of leveraging adversary to debias.

## 6.5. Parameter analysis (RQ3)

To answer **RQ3**, we explore the parameter sensitivity of the two important hyperparameters of our model, using Exp.1.  $\alpha$  controls the impact of the adversary to the sensitive attribute predictor, while  $\beta$  controls the influence of the adversary to the debiasing. We vary both  $\alpha$  and  $\beta$  from [0.0001, 0.001, 0.01, 0.1, 1], and other settings are the same as Exp.1. The results are shown in Figure 3. It is worth noting that, for Figures 3C, D, lower values are better fairness performances, while for Figures 3A, B, higher values indicate better prediction performance. From Figures 3C, D we could observe that: (1) generally, larger  $\alpha$  and  $\beta$  will achieve fairer predictions, while smaller  $\alpha$  and  $\beta$  result in worse fairness; (2) when we increase the value of  $\alpha$ , both  $\Delta_{DP}$  and  $\Delta_{EO}$  will first decrease and then increase when the value of  $\alpha$  is too large. This may be because the estimated target sensitive attributes are still noisy and strengthening the adversarial debiasing that optimizes the noisy sensitive attributes may not necessarily lead to better fairness performances; and (3) for  $\beta$ , the fairness performance is consistently better with a higher value. From Figure 3B, we observe that the classification performance is better when  $\alpha$  and  $\beta$  are balanced. Overall, we observe that when  $\alpha$  and  $\beta$  are in the range of [0.001, 0.1], FAIRDA can achieve relatively good performances of fairness and classification.

## 7. Conclusion and future work

In this paper, we study a novel and challenging problem of exploiting domain adaptation for fair and accurate classification for a target domain without the availability of sensitive attributes. We propose a new framework FAIRDA using a dual adversarial learning approach to achieve a fair and accurate classification. We provide a theoretical analysis to demonstrate that we can learn a fair model prediction under mild assumptions. Experiments on real-world datasets show that the proposed approach can achieve a more fair performance compared to existing approaches by exploiting the information from a source domain, even without knowing the sensitive attributes in the target domain. For future work, first, we can

consider multiple source domains available and explore how to exploit domain discrepancies across multiple domains to enhance the performance of the fair classifier in the target domain. Second, we will explore the fairness transfer under different types of domain shifts, such as the conditional shift of sensitive attributes. Third, we will explore other domain adaption approaches such as meta-transfer learning to achieve cross-domain fair classification.

## Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

## Author contributions

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

## Funding

This work was supported in part by a Cisco Research Award and XCMG American Research Corporation.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## References

Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., and Wallach, H. (2018). "A reductions approach to fair classification," in *International Conference on Machine Learning* (Stockholm), 60–69.

Angwin, J., Larson, J., Mattu, S., and Kirchner, L. (2016). "Machine bias," in *Ethics of Data and Analytics* (Auerbach Publications), 254–264.

- Azizzadenesheli, K., Liu, A., Yang, F., and Anandkumar, A. (2019). Regularized learning for domain adaptation under label shifts. *arXiv preprint arXiv:1903.09734*. doi: 10.48550/arXiv.1903.09734
- Barocas, S., Hardt, M., and Narayanan, A. (2017). "Fairness in machine learning." in *NeurIPS Tutorial, Vol. 1, 2*. Available online at: <http://www.fairmlbook.org>
- Behavod, Y., and Ligett, K. (2017). Learning fair classifiers: a regularization-inspired approach. *arXiv preprint arXiv:1707.00044*. doi: 10.48550/arXiv.1707.00044
- Caton, S., and Haas, C. (2020). Fairness in machine learning: a survey. *CoRR*, abs/2010.04053. doi: 10.48550/arXiv.2010.04053
- Chen, Q., Liu, Y., Wang, Z., Wassell, I., and Chetty, K. (2018). "Re-weighted adversarial adaptation network for unsupervised domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Salt Lake City, UT: IEEE), 7976–7985.
- Cheng, L., Varshney, K. R., and Liu, H. (2021). Socially responsible ai algorithms: Issues, purposes, and challenges. *arXiv preprint arXiv:2101.02032*. doi: 10.1613/jair.1.12814
- Chuang, C.-Y., and Mroueh, Y. (2021). Fair mixup: fairness via interpolation," in *International Conference on Learning Representations* (Vienna).
- Clavijo, J. M., Glaysher, P., Jitsev, J., and Katzy, J. (2021). Adversarial domain adaptation to reduce sample bias of a high energy physics event classifier. *Mach. Learn. Sci. Technol.* 3, 015014. doi: 10.1088/2632-2153/ac3dde
- Coston, A., Ramamurthy, K. N., Wei, D., Varshney, K. R., Speakman, S., Mustahsan, Z., et al. (2019). "Fair transfer learning with missing protected attributes," in *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (Honolulu, HI), 91–98.
- Dixon, L., Li, J., Sorensen, J., Thain, N., and Vasserman, L. (2018). "Measuring and mitigating unintended bias in text classification," in *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society* (New Orleans, LA), 67–73.
- Dou, Q., Ouyang, C., Chen, C., Chen, H., Glocker, B., Zhuang, X., et al. (2019). Pnp-adanet: plug-and-play adversarial domain adaptation network at unpaired cross-modality cardiac segmentation. *IEEE Access* 7:99065–99076. doi: 10.1109/ACCESS.2019.2929258
- Du, M., Yang, F., Zou, N., and Hu, X. (2020). Fairness in deep learning: a computational perspective. *IEEE Intell. Syst.* 36, 25–34. doi: 10.1109/MIS.2020.3000681
- Dua, D., and Graff, C. (2017). *UCI Machine Learning Repository*. Irvine, CA: University of California. Available online at: <http://archive.ics.uci.edu/ml>
- Dwork, C., Immorlica, N., Kalai, A. T., and Leiserson, M. (2018). "Decoupled classifiers for group-fair and efficient machine learning," in *Conference on Fairness, Accountability and Transparency* (New York, NY), 119–133.
- Ganin, Y., and Lempitsky, V. (2015). "Unsupervised domain adaptation by backpropagation," in *International Conference on Machine Learning* (Lille: PMLR), 1180–1189.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., et al. (2016). Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* 17, 1–35. doi: 10.1007/978-3-319-58347-1\_10
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). "Generative adversarial nets," in *NIPS* (Montreal, QC), 2672–2680.
- Gretton, A., Smola, A., Huang, J., Schmittfull, M., Borgwardt, K., and Schölkopf, B. (2009). Covariate shift by kernel mean matching. *Dataset Shift Mach. Learn.* 3, 5. doi: 10.7551/mitpress/9780262170055.003.0008
- Gupta, M., Cotter, A., Fard, M. M., and Wang, S. (2018). Proxy fairness. *arXiv preprint arXiv:1806.11212*. doi: 10.48550/arXiv.1806.11212
- Hamid, A. J., and Ahmed, T. M. (2016). Developing prediction model of loan risk in banks using data mining. *Mach. Learn. Appl.* 3, 3101. doi: 10.5121/mlajj.2016.3101
- Hardt, M., Price, E., and Srebro, N. (2016). Equality of opportunity in supervised learning. *arXiv preprint arXiv:1610.02413*. doi: 10.48550/arXiv.1610.02413
- Kamiran, F., and Calders, T. (2012). Data preprocessing techniques for classification without discrimination. *Knowl. Inf. Syst.* 33, 1–33. doi: 10.1007/s10115-011-0463-8
- Kang, J., He, J., Maciejewski, R., and Tong, H. (2020). "Inform: individual fairness on graph mining," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Diego, CA), 379–389.
- Kearns, M., Neel, S., Roth, A., and Wu, Z. S. (2018). "Preventing fairness gerrymandering: auditing and learning for subgroup fairness," in *International Conference on Machine Learning* (Stockholm), 2564–2572.
- Kouw, W. M., and Loog, M. (2018). An introduction to domain adaptation and transfer learning. *arXiv preprint arXiv:1812.11806*. doi: 10.48550/arXiv.1812.11806
- Lahoti, P., Beutel, A., Chen, J., Lee, K., Prost, F., Thain, N., et al. (2020). Fairness without demographics through adversarially reweighted learning. *arXiv preprint arXiv:2006.13114*.
- Le Quy, T., Roy, A., Iosifidis, V., Zhang, W., and Ntoutsis, E. (2022). "A survey on inspired datasets for fairness-aware machine learning," in *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, e1452.
- Li, Y., Murias, M., Major, S., Dawson, G., and Carlson, D. (2019). "On target shift in adversarial domain adaptation," in *The 22nd International Conference on Artificial Intelligence and Statistics* (Naha, Okinawa: PMLR), 616–625.
- Li, Y., Zhang, Y., and Yang, C. (2022). Unsupervised domain adaptation with joint adversarial variational autoencoder. *Knowl. Based Syst.* 17, 109065. doi: 10.1016/j.knsys.2022.109065
- Lipton, Z., Wang, Y.-X., and Smola, A. (2018). "Detecting and correcting for label shift with black box predictors," in *International Conference on Machine Learning* (Stockholm: PMLR), 3122–3130.
- Liu, X., Guo, Z., Li, S., Xing, F., You, J., Kuo, C.-C. J., et al. (2021). "Adversarial unsupervised domain adaptation with conditional and label shift: Infer, align and iterate," in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (Montreal, QC: IEEE), 10367–10376.
- Liu, Z., Luo, P., Wang, X., and Tang, X. (2015). "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*.
- Luo, Y., Zhang, S.-Y., Zheng, W.-L., and Lu, B.-L. (2018). "Wgan domain adaptation for eeg-based emotion recognition," in *International Conference on Neural Information Processing* (Montreal, QC: Springer), 275–286.
- Madras, D., Creager, E., Pitassi, T., and Zemel, R. (2018). "Learning adversarially fair and transferable representations," in *International Conference on Machine Learning* (Stockholm: PMLR), 3384–3393.
- Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., and Paul Smolley, S. (2017). "Least squares generative adversarial networks," in *Proceedings of the IEEE International Conference on Computer Vision* (Venice: IEEE), 2794–2802.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. (2019). A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*. doi: 10.48550/arXiv.1908.09635
- Mena, J. (2011). *Machine Learning Forensics for Law Enforcement, Security, and Intelligence*. CRC Press.
- Pan, S. J., and Yang, Q. (2009). A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* 22, 1345–1359. doi: 10.1109/TKDE.2009.191
- Redmond, M. (2011). *Communities and Crime Unnormalized Data set*. UCI Machine Learning Repository. Available online at: <http://www.ics.uci.edu/mllearn/MLRepository.html>
- Schumann, C., Wang, X., Beutel, A., Chen, J., Qian, H., and Chi, E. H. (2019). Transfer of machine learning fairness across domains. *arXiv preprint arXiv:1906.09688*. doi: 10.48550/arXiv.1906.09688
- Shu, Y., Cao, Z., Long, M., and Wang, J. (2019). Transferable curriculum for weakly-supervised domain adaptation. *Proc. AAAI Conf. Artif. Intell.* 33, 4951–4958. doi: 10.1609/aaai.v33i01.33014951
- Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., and Liu, C. (2018). "A survey on deep transfer learning," in *International Conference on Artificial Neural Networks* (Rhodes: Springer), 270–279.
- Tan, Y., Li, Y., and Huang, S.-L. (2021). "Otc: a transferability metric for cross-domain cross-task representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Nashville, TN: IEEE), 15779–15788.
- Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. (2017). "Adversarial discriminative domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Honolulu, HI: IEEE), 7167–7176.
- Wu, Y., Winston, E., Kaushik, D., and Lipton, Z. (2019). "Domain adaptation with asymmetrically-relaxed distribution alignment," in *International Conference on Machine Learning* (Long Beach, CA: PMLR), 6872–6881.
- Yan, S., Kao, H.-T., and Ferrara, E. (2020). "Fair class balancing: Enhancing model fairness without observing sensitive attributes," in *Proceedings of the 29th ACM International Conference on Information Knowledge Management*, 1715–1724.
- Zhang, L., Wu, Y., and Wu, X. (2017). "Achieving non-discrimination in data release," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Halifax, NS), 1335–1344.
- Zhang, Y., Liu, T., Long, M., and Jordan, M. (2019). "Bridging theory and algorithm for domain adaptation," in *International Conference on Machine Learning* (Long Beach, CA: PMLR), 7404–7413.

Zhao, T., Dai, E., Shu, K., and Wang, S. (2022). "You can still achieve fairness without sensitive attributes: exploring biases in non-sensitive features," in *WSDM '22, The Fourteenth ACM International Conference on Web Search and Data Mining, Phoenix, March 8-12, 2022* (Phoenix: ACM).

Zhou, R., and Pan, Y. (2022). "Floodan: unsupervised flood forecasting based on adversarial domain adaptation," in *2022 IEEE 5th International Conference on Big Data and Artificial Intelligence (BDAI)* (Fuzhou: IEEE), 6–12.

Zhuang, F., Cheng, X., Luo, P., Pan, S. J., and He, Q. (2015). "Supervised representation learning: Transfer learning with deep autoencoders," in *IJCAI* (Buenos Aires).

Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., et al. (2020). A comprehensive survey on transfer learning. *Proc. IEEE* 109, 43–76. doi: 10.1109/JPROC.2020.3004555