

000 001 FAIRNESS VIA INDEPENDENCE: A GENERAL REGULAR- 002 IZATION FRAMEWORK FOR MACHINE LEARNING 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Fairness in machine learning has emerged as a central concern, as predictive models
012 frequently inherit or even amplify biases present in training data. Such biases
013 often manifest as unintended correlations between model outcomes and sensitive
014 attributes, leading to systematic disparities across demographic groups. Existing
015 approaches to fair learning largely fall into two directions: incorporating fairness
016 constraints tailored to specific definitions, which limits their generalizability, or
017 reducing the statistical dependence between predictions and sensitive attributes,
018 which is more flexible but highly sensitive to the choice of distance measure. The
019 latter strategy in particular raises the challenge of finding a principled and reliable
020 measure of dependence that can perform consistently across tasks. In this work,
021 we present a general and model-agnostic approach to address this challenge. The
022 method is based on encouraging independence between predictions and sensitive
023 features through an optimization framework that leverages the Cauchy–Schwarz
024 (CS) Divergence as a principled measure of dependence. Prior studies suggest
025 that CS Divergence provides a tighter theoretical bound compared to alternative
026 distance measures used in earlier fairness methods, offering a stronger foundation
027 for fairness-oriented optimization. Our framework, therefore, unifies prior efforts
028 under a simple yet effective principle and highlights the value of carefully chosen
029 statistical measures in fair learning. Through extensive empirical evaluation on
030 four tabular datasets and one image dataset, we show that our approach consistently
improves multiple fairness metrics while maintaining competitive accuracy.

031 032 1 INTRODUCTION 033

035 Fairness in machine learning has garnered growing concern, as machine learning (ML) models are
036 playing key roles in many high-stakes decision-making scenarios, such as credit scoring (Leal, 2022),
037 the job market (Hu & Chen, 2018), healthcare (Grote & Keeling, 2022), and education (Bøyum,
038 2014; Kizilcec & Lee, 2022). Among the various fairness notions, group fairness is one of the
039 most extensively studied ones as it addresses the prediction disparities across demographic groups,
040 including gender, age, skin color, and region (Mehrabi et al., 2021; Dwork et al., 2012; Barocas
041 et al., 2017). While many group fairness ML algorithms are proposed, they have challenges in their
042 applications, especially the *generalizability*, i.e., their adaptation to different fairness notions, and
043 *robustness*, i.e., the stability of the fairness when they encounter a slight change of model parameter.

044 Existing group fairness approaches can be intrinsically categorized into two main approaches based
045 on their debiasing objectives: *i*) directly integrate the fairness notion into the training objective,
046 *ii*) minimizing the correlation between predictions and sensitive attributes. Methods belongs to *i*)
047 such as a demographic parity (DP) regularizer, and an equality of opportunity (EO) regularizer. The
048 benefit of this approach is that the model trained by the target fairness objective can perform well on
049 specific fairness notions. For example, the machine learning model trained at demographic parity
050 has a high possibility of achieving good demographic parity in testing. However, such methods
051 limited their generalizability to other fairness notions (shown in Figure 1). Method *ii*) solves from a
052 more fundamental way that can deal with generalizability, including using information theory, or an
053 adversarial approach to minimize the correlation between the prediction and the sensitive attribute.
The most straightforward way is to use a distance measurement that assesses the relationship between
the sensitive attribute and the prediction, thus minimizing this distance during the training. This

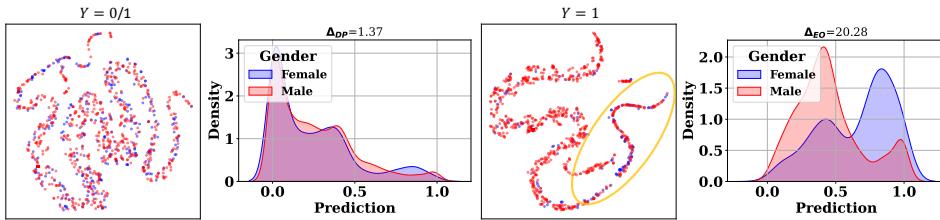
054
055
056
057
058
059
060
061

Figure 1: From left to right: (1) Prediction distribution of *all classes*; (2) T-SNE plot of embeddings for samples from *all classes*; (3) Prediction distribution of *class 1*; (4) T-SNE plot of embeddings for samples from *Adult*, and the sensitive attribute is gender. The blue points represent samples with sensitive attribute 0, while the red points represent samples with sensitive attribute 1.

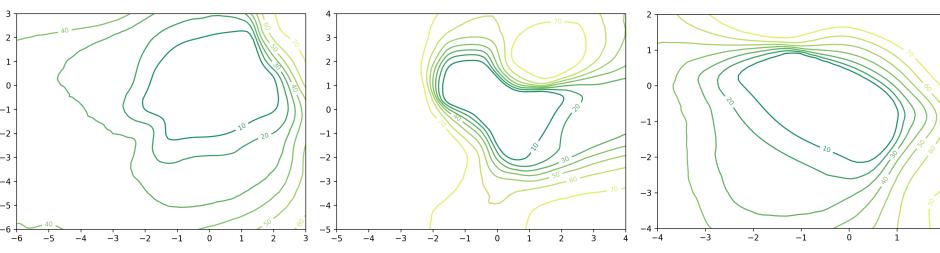
066
067
068
069
070
071
072
073
074

Figure 2: Fairness loss landscapes evaluated using three functions, presented from left to right: Kullback-Leibler (KL) divergence, Hilbert-Schmidt Independence Criterion (HSIC), and Cauchy-Schwarz (CS) divergence. A smaller inner circle indicates greater robustness. Among these methods, the CS divergence achieves the smallest inner circle, ranging from -2 to 1 , while the inner circles of KL and HSIC divergences both span from -2 to 2 .

080

enables the generalizability, but ascribes the pressure of the fairness performance to the quality of the distance measurement.

083
084
085
086
087
088
089
090
091
092
093
094
095
096

Existing fairness regularizers mainly assessed this correlation using gap parity, **Maximum Mean Discrepancy (MMD)** (Gretton et al., 2012), **Kullback-Leibler (KL) divergence**, and the **Hilbert-Schmidt Independence Criterion (HSIC)**. However, the current fairness regularizers are sensitive to the model parameter change, making them less robust in maintaining fairness, responding to a small change in the model parameters (shown in Figure 2). Theoretical studies have shown that the Cauchy-Schwarz divergence provides a tighter bound compared to the Kullback-Leibler divergence and gap parity, suggesting its potential to improve fairness in machine learning models. Motivated by this, we would like to know if using CS divergence can result in more generalizable and **more consistent utility-fairness trade-off across hyperparameter settings than standard regularizers** due to the benefit of a tighter upper bound of CS divergence. In light of this, we propose a new fairness regularizer based on the Cauchy-Schwarz divergence for fair machine learning. To evaluate the generalizability, we tested the fairness under a wide range of fairness notions proposed by previous studies, and to evaluate the robustness, we visualize whether a small change in the learned model parameters can influence the fairness. We summarize our contributions as follows:

097
098
100
101
102
103
104
105
106
107

- We introduce the Cauchy-Schwarz divergence to fair machine learning and present a novel regularization method.
- We elucidate the relationships between the Cauchy-Schwarz regularizer and other fairness regularizers, emphasizing its superior effectiveness in debiasing.
- Our experimental results, obtained from four tabular datasets and one image dataset, validate the efficacy of the proposed Cauchy-Schwarz regularizer in achieving fairness across multiple fairness notions simultaneously.

This develops CS divergence in a fairness-specific setting that is complementary to prior work on CS-based density estimation, clustering, and representation learning.

To Revvier
ESYX: Question 2

To Revvier JuMp:
W1

To Revvier LT9c:
W3

To Revvier JuMp:
W3

108 **2 PRELIMINARIES**

110 In this section, we establish the foundational concepts for our study. We start by exploring the notion
 111 of fairness in machine learning, including the relevant notations. Next, we provide an overview
 112 of general fairness-aware machine learning methods. Finally, we introduce the Cauchy-Schwarz
 113 divergence and discuss its benefits in reducing bias.

114 **Problem scope.** This paper focuses on in-process group fairness in the context of binary classification
 115 with binary-sensitive attributes. In contrast, group fairness seeks to ensure that machine learning
 116 models treat different demographic groups equitably, where groups are defined based on sensitive
 117 attributes such as gender, race, and age (Feldman et al., 2015; Zemel et al., 2013).

118 **Notations.** Under this setting, we consider a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i, s_i)\}_{i=1}^M$, where M is the number
 119 of samples, $\mathbf{x}_i \in \mathbb{R}^d$ represents the features excluding the sensitive attribute, $y_i \in \{0, 1\}$ is the label
 120 of the downstream task, and $s_i \in \{0, 1\}$ is the sensitive attribute of the i -th sample. The predicted
 121 probability for the i -th sample is denoted as $z_i \in [0, 1]$, computed by the machine learning model
 122 as $z_i = f(\mathbf{x}_i, s_i) : \mathbb{R}^d \rightarrow [0, 1]$. The binary prediction is represented as $\hat{y}_i \in \{0, 1\}$, defined by
 123 $\hat{y}_i = \mathbf{1}_{\{\cdot \geq t[z_i]\}}$, where $\mathbf{1}_{\{\cdot \geq t(\cdot)\}}$ is the indicator function that evaluates whether its input is greater
 124 than or equal to the threshold t . Finally, X, Y, S , and \hat{Y} denote random variables corresponding to
 125 \mathbf{x}_i, y_i, s_i , and \hat{y}_i , respectively.

126 **Problem Formulation.** Generally, the fairness objective can be summarized as follows:

$$\min_f \quad \mathcal{L}_{utility} + \lambda \mathcal{L}_{fairness}, \quad (1)$$

130 where the term $\mathcal{L}_{utility}$ denotes the loss function that measures the utility of the model, often a binary
 131 entropy loss, for our binary classification problem, while $\mathcal{L}_{fairness}$ (also shown in Figure 2) indicates
 132 the fairness constraint applied in the model. The parameter λ is used to control the trade-off between
 133 utility and fairness.

135 **2.1 GROUP FAIRNESS**

137 There are many ways to define and measure the group fairness. Each definition focuses on distinct
 138 statistical measures aimed at achieving balance among subgroups within the data. Among these, Demographic
 139 Parity and Equal Opportunity are the most popular ones, other popular fairness definitions
 140 are summarized in Appendix J.2.

141 **Demographic Parity (DP).** Demographic Parity (Zafar et al., 2017; Feldman et al., 2015; Dwork
 142 et al., 2012) mandates that the predicted outcome \hat{Y} be independent of the sensitive attribute S ,
 143 expressed mathematically as $\hat{Y} \perp S$. Most of the existing literature primarily addresses binary
 144 classification and binary attributes, where $Y \in \{0, 1\}$ and $S \in \{0, 1\}$. Similar to the concept of equal
 145 opportunity, the metric evaluating the DP fairness is defined by:

$$\Delta_{DP} = |P(\hat{Y}|S = 0) - P(\hat{Y}|S = 1)|. \quad (2)$$

148 A lower value of Δ_{DP} signifies a fairer classifier.

150 **Equal Opportunity (EO).** Equal Opportunity (Hardt et al., 2016) mandates that a classifier achieves
 151 equal true positive rates across various subgroups, striving towards the ideal of a perfect classifier.
 152 The corresponding fairness measurement for EO can be articulated as follows:

$$\Delta_{EO} = |P(\hat{Y}|Y = 1, S = 0) - P(\hat{Y}|Y = 1, S = 1)|. \quad (3)$$

156 A low Δ_{EO} indicates that the difference in the probability of an instance in the positive class being
 157 assigned a positive outcome is relatively small for both subgroup members. Both DP and EO can be
 158 effectively extended to problems involving multi-class classifications and multiple sensitive attribute
 159 categories. Note that in the binary classification task, Δ_{DP} and Δ_{EO} are sometimes calculated after
 160 binarizing $P(\hat{Y})$. Specifically, Δ_{DP} is defined as $\Delta_{DP} = |P(\hat{Y} = 1|S = 0) - P(\hat{Y} = 1|S = 1)|$,
 161 and Δ_{EO} is defined as $\Delta_{EO} = |P(\hat{Y} = 1|Y = 1, S = 0) - P(\hat{Y} = 1|Y = 1, S = 1)|$ (Beutel et al.,
 2017; Dai & Wang, 2021; Dong et al., 2022).

162 2.2 CAUCHY-SCHWARZ DIVERGENCE
163

164 Motivated by the well-known Cauchy-Schwarz (CS) inequality for square-integrable functions¹,
165 which holds with equality if and only if $p(\mathbf{x})$ and $q(\mathbf{x})$ are linearly dependent, we can define a measure
166 of the distance between $p(\mathbf{x})$ and $q(\mathbf{x})$. This measure is referred to as the CS divergence (Principe
167 et al., 2000; Yu et al., 2023), given by:

$$168 D_{\text{CS}}(p; q) = -\log \left(\frac{\left(\int p(\mathbf{x})q(\mathbf{x})dx \right)^2}{\int p(\mathbf{x})^2dx \int q(\mathbf{x})^2dx} \right). \quad (4)$$

172 The CS divergence, denoted as D_{CS} , is symmetric for any two probability density functions (PDFs) p
173 and q , satisfying $0 \leq D_{\text{CS}} < \infty$. The minimum divergence is achieved if and only if $p(\mathbf{x}) = q(\mathbf{x})$.

174 3 WHAT MAKES A GOOD FAIRNESS REGULARIZER?

176 Current fair machine learning algorithms
177 adopt a variety of approaches, prompting us to
178 explore the essential properties that make for
179 effective fairness regularizers. By categorizing
180 these algorithms into three types and con-
181 ducting preliminary experiments, we aim to
182 identify the key characteristics that contribute
183 to mitigating bias in machine learning models.

Reg.	Fairness Objective ($\mathcal{L}_{\text{fairness}}$)
DP	$ \mathbb{E}(\hat{Y} S=0) - \mathbb{E}(\hat{Y} S=1) $
EO	$ \mathbb{E}(\hat{Y} Y=1, S=0) - \mathbb{E}(\hat{Y} Y=1, S=1) $
MMD	$D_{\text{MMD}}(Z_{S=0}, Z_{S=1})$
HSIC	$D_{\text{HSIC}}(\hat{Y}, S)$
PR	$D_{\text{PR}}(\hat{Y}, S)$

184 Table 1: Fairness regularizers (Reg.) and objectives.

185 3.1 BALANCING THE PREDICTION ACROSS DIFFERENT SENSITIVE GROUPS

186 The first is to directly integrate the fairness notions, such as DP and EO, into the fairness objective.

$$187 \mathcal{L}_{\text{fairness}} = D(\mathbb{P}, \mathbb{Q}) \quad \text{where}$$

$$188 \begin{cases} \mathbb{P} = P(\hat{Y} | S=0), \quad \mathbb{Q} = P(\hat{Y} | S=1) \text{ for DP,} \\ 189 \mathbb{P} = P(\hat{Y} | Y=1, S=0), \quad \mathbb{Q} = P(\hat{Y} | Y=1, S=1) \text{ for EO.} \end{cases} \quad (5)$$

190 Calculating the distance between \mathbb{P} and \mathbb{Q} has many ways by using difference distance measurement
191 D , and the most used one is to calculate the absolute distance between the mean empirical estimations:

$$192 \mathcal{L}_{\text{fairness}} = |\mathbb{E}(\mathbb{P}) - \mathbb{E}(\mathbb{Q})|, \quad (6)$$

193 where the expected values are calculated as the mean of summation since \mathbb{P} and \mathbb{Q} are discrete
194 distributions.

195 Previous fair machine learning studies have shown that the fairness loss of the testing can be upper
196 bounded by the loss of training. Therefore, a *distance function having a tighter generalization error*
197 *bound used in training will lead to a better fairness guarantee for testing*.

$$201 \mathcal{L}_{\text{fairness}} = |\mathbb{E}(\hat{Y}|S=0) - \mathbb{E}(\hat{Y}|S=1)|. \quad (7)$$

202 We can see that the basic idea is to balance the prediction distribution between two sensitive groups.
203 Therefore, for this type of approach, we can also use other distance measurements: Therefore, except
204 for the absolute distance between the two prediction distributions.

205 3.2 BALANCING THE LATENT REPRESENTATION ACROSS DIFFERENT SENSITIVE GROUPS

206 Distance measures and minimization:

$$207 \mathcal{L}_{\text{fairness}} = D(Z_{S=0}, Z_{S=1}), \quad (8)$$

208 where Z is the latent representation from the neural networks, and $Z_{S=0}$ and $Z_{S=1}$ are the represen-
209 tation when the sensitive attribute is 1 or 0. The distance metric $D(\cdot)$ here can be a Mean Maximum
210 Discrepancy (Louizos et al., 2016).

211 ¹ $(\int p(\mathbf{x})q(\mathbf{x}) dx)^2 \leq \int p(\mathbf{x})^2 dx \int q(\mathbf{x})^2 dx$

216 **Extension to multiple sensitive attributes.** While our theoretical development is presented for
 217 a single sensitive attribute S for notational simplicity, the proposed CS-based regularizer naturally
 218 extends to multiple sensitive attributes. Let $S = (S_1, \dots, S_K)$ denote a vector of sensitive at-
 219 tributes. One can either (i) treat S as a joint variable and apply the same CS divergence to (\hat{Y}, S) ,
 220 i.e., penalize $\tilde{D}_{\text{CS}}(P_{\hat{Y}|S}, P_{\hat{Y}}P_S)$, or (ii) sum the CS divergences over individual attributes, e.g.,
 221 $\sum_{k=1}^K \tilde{D}_{\text{CS}}(P_{\hat{Y}|S_k}, P_{\hat{Y}}P_{S_k})$, depending on whether joint or per-attribute control is desired. Both
 222 variants use exactly the same empirical estimator as in (Equation (13)) and do not require any
 223 algorithmic changes; only the definition of S in the mini-batch needs to be updated.

To Rewvier Pqrj:
W3-2

225 3.3 MINIMIZING THE RELATIONSHIP BETWEEN PREDICTIONS AND SENSITIVE ATTRIBUTES

227 The goal of this category is to ensure that a fair machine learning algorithm’s predictions retain
 228 minimal sensitive information.

$$229 \mathcal{L}_{\text{fairness}} = D(\hat{Y}, S), \quad (9)$$

231 The HSIC and PR in Table 1 belong to this category. Specifically, the $D_{\text{PR}}(\hat{Y}, S)$ is defined as:
 232

$$233 D_{\text{PR}}(\hat{Y}, S) = \sum_{\hat{y} \in \hat{Y}} \sum_{s \in S} p(\hat{y}, s) \log \left(\frac{p(\hat{y}, s)}{p(\hat{y})p(s)} \right). \quad (10)$$

236 Note that, adversarial debiasing methods (Zhang et al., 2018) fall into this category because they em-
 237 ploy discriminators to predict sensitive group membership from the learned encoded representations.
 238 These methods aim to make sensitive attributes difficult to deduce from the encoded representations.

239 4 CAUCHY-SCHWARZ FAIRNESS REGULARIZER

241 In this section, we first introduce three prominent fairness regularizers that assess distribution distance
 242 using different metrics: Mean Maximum Discrepancy, Kullback-Leibler divergence, and Hilbert-
 243 Schmidt Independence Criterion (HSIC). For each metric, we explore its relationship with CS
 244 divergence. Subsequently, we explain how CS divergence can be utilized to achieve fairness.

245 4.1 HOW CAN THE CAUCHY-SCHWARZ DIVERGENCE BE APPLIED TO MITIGATE BIAS?

247 Given samples $\{\mathbf{x}_i^p\}_{i=1}^m$ and $\{\mathbf{x}_i^q\}_{i=1}^n$ drawn independently and identically distributed (i.i.d.) from
 248 $p(\mathbf{x})$ and $q(\mathbf{x})$ respectively, we can estimate the empirical CS divergence. This estimation can be
 249 performed using the kernel density estimator (KDE) as described in (Parzen, 1962) and follows the
 250 empirical estimator formula in (Jenssen et al., 2006).

251 **Proposition 4.1.** *Given two sets of observations $\{\mathbf{x}_i^p\}_{i=1}^{N_1}$ and $\{\mathbf{x}_j^q\}_{j=1}^{N_2}$, let p and q denote the
 252 distributions of two groups. The empirical estimator of the CS divergence $D_{\text{CS}}(p; q)$ is then given by:*

$$254 \tilde{D}_{\text{CS}}(p; q) = \log \left(\frac{1}{N_1^2} \sum_{i,j=1}^{N_1} \kappa(\mathbf{x}_i^p, \mathbf{x}_j^p) \right) + \log \left(\frac{1}{N_2^2} \sum_{i,j=1}^{N_2} \kappa(\mathbf{x}_i^q, \mathbf{x}_j^q) \right) \\ 255 - 2 \log \left(\frac{1}{N_1 N_2} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \kappa(\mathbf{x}_i^p, \mathbf{x}_j^q) \right). \quad (11)$$

261 The proof of this proposition is detailed in Appendix B.1, where κ represents a kernel function, such
 262 as the Gaussian kernel defined as $\kappa_\sigma(x, x') = \exp(-\|x - x'\|_2^2/2\sigma^2)$. In the following sections, we
 263 will explore the relationship between this kernel function and the existing fairness regularizer.

264 As mentioned earlier, the goal of fairness is to ensure an equal distribution of predictions across
 265 sensitive attributes. To achieve this, fairness-aware algorithms focus on minimizing the dependency
 266 of predictions on these sensitive attributes. Therefore, effectively modeling the relationship between
 267 the outcome variable Y and the sensitive attribute S becomes crucial. The prediction distribution
 268 over the sensitive attribute S is defined as follows:

$$269 \mathbb{P} = P(\hat{Y} | S = 0); \quad \mathbb{Q} = P(\hat{Y} | S = 1). \quad (12)$$

270 By substituting the distribution of predictions over the sensitive attribute into Equation (20), where
 271 $p = \mathbb{P}$ and $q = \mathbb{Q}$, we can define the objective we aim to solve as follows:
 272

$$273 \min_{\theta} \mathcal{L}_{\text{BCE}} + \alpha \tilde{D}_{\text{CS}}(\mathbb{P}, \mathbb{Q}) + \frac{\beta}{2} \|\theta\|_2^2, \quad (13)$$

275 where \mathcal{L}_{BCE} is the binary cross-entropy loss, which measures the classifier's accuracy. It is defined as:
 276

$$277 \mathcal{L}_{\text{BCE}} = \frac{1}{M} \sum_{i=1}^M -Y_i \log \hat{Y}_i, \quad (14)$$

278 where \hat{Y}_i is the predicted output obtained from the training model parameterized by θ . This model
 279 can be a Multi-Layer Perceptron for tabular data or a ResNet for image data. Additionally, $\|\theta\|_2^2$
 280 serves as an L_2 regularizer.
 281

282 **Computational complexity.** Naively computing the CS divergence on all n training examples
 283 has $O(n^2)$ cost, since it involves pairwise interactions between samples. In practice, as is standard
 284 for kernel-based regularizers such as MMD and HSIC, we evaluate the CS-based fairness loss on
 285 *mini-batches*. Given a batch of size B (typically $B \ll n$), the additional cost per optimization
 286 step is $O(B^2)$ and is implemented using vectorized matrix operations that reuse the same mini-
 287 batches as the prediction loss. Under typical batch sizes used in our experiments, this overhead is
 288 modest and comparable to that of existing kernel-based fairness methods, making the CS regularizer
 289 computationally practical for both tabular and image models.
 290

To Rewvier Pqrj:
 W3-1

292 4.2 WHY IS THE CS DIVERGENCE MORE EFFECTIVE FOR ENSURING FAIRNESS?

293 The CS Divergence is particularly well-suited for promoting fairness due to several key reasons:
 294

295 **(1) Closed-form solution for the mixture of Gaussians.** The CS divergence has several advantageous
 296 properties, one of which is that it provides a *closed-form solution for the mixture of Gaussians* (Kampa
 297 et al., 2011). This particular property has facilitated its successful application in various tasks,
 298 including deep clustering (Trosten et al., 2021), disentangled representation learning (Tran et al.,
 299 2022), and point-set registration (Sanchez Giraldo et al., 2017).

300 **(2) CS Divergence has a tighter error bound than the KL divergence.**

301 **Proposition 4.2.** For any d -variate Gaussian distributions $p \sim \mathcal{N}(\mu_p, \Sigma_p)$ and $q \sim \mathcal{N}(\mu_q, \Sigma_q)$,
 302 where Σ_p and Σ_q are positive definite, the following inequality holds:
 303

$$304 D_{\text{CS}}(p; q) \leq D_{\text{KL}}(p; q) \text{ and } D_{\text{CS}}(p; q) \leq D_{\text{KL}}(q; p). \quad (15)$$

305 The proof can be found in Appendix B.3. It is important to note that the divergences are being
 306 compared under the same model parameter θ .
 307

308 **Remark on the Gaussian setting.** Proposition 4.2 should be interpreted as a stylized comparison
 309 that is carried out under a Gaussian model solely for analytical convenience. The Gaussian assumption
 310 allows us to derive closed-form expressions for both the CS and KL divergences and to make their
 311 relationship explicit. It is *not* required by our training procedure, and it is not assumed anywhere in
 312 the empirical evaluation; in particular, the CS-based regularizer we optimize is defined for arbitrary
 313 distributions of predictions and sensitive attributes.
 314

To Rewvier
 ESYX: W1-1

315 **(3) CS divergence can provide tighter bounds than MMD and DP when the distributions are far
 316 apart or when the scale of the embeddings varies significantly.** Based on Remark A.1, we know
 317 that CS divergence employs cosine distance, while MMD relies on Euclidean distance. In addition,
 318 DP Equation (5) utilizes a mean disparity, which is a Manhattan distance for the mean estimations
 319 of two distributions. CS divergence measures the angle between two distributions in the feature
 320 space, focusing on the difference in direction rather than magnitude. In cases where the distributions
 321 have significantly different variances or scales, MMD and DP may yield a large distance even if the
 322 distributions are aligned in the feature space. In contrast, CS divergence normalizes this comparison,
 323 resulting in a more accurate measure of similarity and thereby providing a tighter generalization
 bound. This normalization enhances the robustness of CS divergence, preventing MMD and DP from
 overestimating the discrepancy due to their reliance on an unnormalized distance measure.

	Methods	Utility				Fairness			
		ACC (%)	↑	AUC (%)	↑	Δ_{DP} (%)	↓	Δ_{EO} (%)	↓
Adult	MLP	85.63 \pm 0.34	—	90.82 \pm 0.23	—	16.52 \pm 0.91	—	8.43 \pm 3.20	—
	Gender DP	82.42 \pm 0.39	-3.75%	86.91 \pm 0.80	-4.31%	1.29 \pm 0.95	92.19%	20.15 \pm 1.13	-139.03%
	MMD	81.90 \pm 0.68	-4.36%	85.27 \pm 0.52	-6.11%	2.47 \pm 0.52	85.05%	17.53 \pm 1.36	-107.95%
	HSIC	82.89 \pm 0.23	-3.20%	87.25 \pm 0.41	-3.93%	2.66 \pm 0.54	83.90%	18.47 \pm 1.22	-119.10%
	PR	81.81 \pm 0.52	-4.46%	85.38 \pm 0.82	-5.99%	0.71 \pm 0.40	95.70%	12.45 \pm 2.38	-47.69%
Race	CS	83.31 \pm 0.47	-2.71%	90.15 \pm 0.49	-0.74%	2.42 \pm 0.85	85.35%	2.27 \pm 1.04	73.07%
	MLP	84.42 \pm 0.31	—	90.15 \pm 0.36	—	13.47 \pm 0.83	—	9.25 \pm 3.86	—
	Race DP	83.64 \pm 0.78	-0.92%	88.45 \pm 0.32	-1.89%	2.45 \pm 0.67	81.81%	2.16 \pm 1.06	76.65%
	MMD	83.12 \pm 0.82	-1.54%	88.36 \pm 0.67	-1.99%	2.58 \pm 0.75	80.85%	3.33 \pm 0.93	64.00%
	HSIC	84.98 \pm 0.17	0.66%	90.90 \pm 0.19	0.83%	7.90 \pm 0.72	41.35%	2.11 \pm 0.18	77.19%
COMPAS	PR	82.13 \pm 1.16	-2.71%	87.44 \pm 0.33	-3.01%	1.53 \pm 0.83	88.64%	0.86 \pm 0.60	90.70%
	CS	83.53 \pm 0.53	-1.05%	90.26 \pm 0.47	0.12%	2.16 \pm 0.61	83.96%	0.44 \pm 0.12	95.24%
	MLP	66.85 \pm 0.72	—	72.10 \pm 0.94	—	13.22 \pm 3.32	—	11.41 \pm 5.83	—
	Gender DP	64.20 \pm 1.58	-3.96%	70.64 \pm 1.05	-2.02%	5.78 \pm 0.33	56.28%	6.78 \pm 1.61	40.58%
	MMD	64.82 \pm 1.62	-3.04%	70.72 \pm 0.92	-1.91%	3.09 \pm 0.92	76.63%	3.15 \pm 4.37	72.39%
Race	HSIC	63.17 \pm 3.46	-5.50%	71.17 \pm 0.84	-1.29%	1.84 \pm 0.43	86.08%	2.60 \pm 0.63	77.21%
	PR	64.95 \pm 0.15	-2.84%	72.12 \pm 0.75	0.03%	3.85 \pm 0.60	70.88%	3.91 \pm 1.02	65.73%
	CS	64.25 \pm 0.97	-3.89%	71.53 \pm 0.61	-0.79%	1.30 \pm 0.47	90.17%	0.44 \pm 0.13	96.14%
	MLP	66.99 \pm 1.05	—	72.46 \pm 0.88	—	17.24 \pm 4.15	—	19.44 \pm 4.63	—
	Race DP	64.98 \pm 3.72	-3.00%	72.09 \pm 1.03	0.51%	8.70 \pm 1.12	49.54%	7.04 \pm 2.13	63.79%
ACS-I	MMD	64.41 \pm 2.04	-3.85%	72.10 \pm 1.83	0.50%	4.42 \pm 2.11	74.36%	5.60 \pm 1.25	71.19%
	HSIC	64.52 \pm 2.20	-3.69%	72.16 \pm 0.94	0.41%	2.21 \pm 0.68	87.18%	2.72 \pm 0.87	86.01%
	PR	67.22 \pm 0.90	0.34%	72.86 \pm 0.87	-0.55%	5.60 \pm 1.12	67.52%	6.52 \pm 1.30	66.46%
	CS	65.62 \pm 1.24	-2.05%	72.70 \pm 1.06	0.33%	1.79 \pm 0.96	89.62%	1.48 \pm 1.64	92.39%
	MLP	82.04 \pm 0.27	—	90.16 \pm 0.18	—	10.26 \pm 4.68	—	2.13 \pm 3.64	—
Gender	Gender DP	81.32 \pm 0.17	-0.88%	89.33 \pm 0.15	-0.92%	0.96 \pm 0.22	90.64%	5.37 \pm 0.32	-152.11%
	MMD	80.93 \pm 0.55	-1.35%	88.44 \pm 1.71	-1.91%	2.45 \pm 0.65	76.12%	4.91 \pm 1.48	-130.52%
	HSIC	81.40 \pm 0.12	-0.78%	89.53 \pm 0.10	-0.70%	1.54 \pm 0.18	84.99%	4.95 \pm 0.39	-132.39%
	PR	80.03 \pm 0.30	-2.45%	88.10 \pm 0.26	-2.28%	0.35 \pm 0.20	96.59%	4.54 \pm 0.41	-113.15%
	CS	81.86 \pm 0.94	-0.22%	89.15 \pm 0.60	-1.12%	0.77 \pm 0.38	92.5%	0.90 \pm 0.46	57.75%
Race	MLP	81.23 \pm 0.14	—	90.16 \pm 0.18	—	10.06 \pm 1.84	—	7.42 \pm 0.66	—
	Race DP	81.25 \pm 0.13	0.02%	89.45 \pm 0.11	-0.79%	0.56 \pm 0.30	94.43%	4.53 \pm 0.48	38.95%
	MMD	80.22 \pm 1.22	-1.24%	88.42 \pm 1.63	-1.93%	1.45 \pm 0.89	85.59%	4.01 \pm 0.54	45.96%
	HSIC	81.41 \pm 0.15	0.22%	89.67 \pm 0.12	-0.54%	1.04 \pm 0.53	89.66%	2.77 \pm 0.35	62.67%
	PR	80.27 \pm 0.26	-1.18%	88.45 \pm 0.21	-1.90%	0.37 \pm 0.30	96.32%	4.25 \pm 0.49	42.72%
ACS-I	CS	80.78 \pm 0.85	-0.55%	89.14 \pm 0.94	-1.13%	0.81 \pm 0.28	91.95%	1.35 \pm 0.64	81.81%

Table 2: Fairness performance of existing fair models on the tabular datasets, considering race and gender as sensitive attributes. \uparrow indicates accuracy improvement **compared to MLP**, with higher accuracy reflecting better performance, and \downarrow denotes fairness improvement **compared to MLP**, where lower values indicate better fairness. **Green** values denotes better than MLP on the corresponding metric (ACC/AUC \uparrow ; $\Delta_{DP}/\Delta_{EO}\downarrow$), while **red** denote worse. rAll results are based on 10 runs for each method. The best results for each metric and dataset are highlighted in **bold** text.

Distribution-free applicability. We also emphasize that the optimization framework of Section 4 is distribution-free. The empirical CS divergence used as a fairness loss is implemented via a kernel-based estimator that only depends on samples from the joint distribution of predictions and sensitive attributes and does not impose any parametric form on this distribution. Consequently, the same regularizer can be applied to both tabular and image models. Our experiments in Section 5, which span four tabular datasets and one image dataset with clearly non-Gaussian distributions, show that the CS-based regularizer consistently improves group-fairness metrics while maintaining competitive utility, supporting the practical relevance of the theoretical insight obtained from the Gaussian comparison in Proposition 4.2.

Limitations and discussion. While CS offers a flexible and theoretically grounded way to penalize dependence between predictions and sensitive attributes, it also has several limitations. First, CS is implemented as a kernel-based dependence measure and thus is computationally more expensive than simple gap-based penalties such as DP or EO alone: evaluating the regularizer on a mini-batch of size B incurs $O(B^2)$ cost, similar to MMD and HSIC, whereas DP/EO-style losses can often be computed in $O(B)$ time. Second, as with other kernel methods, the performance of CS can depend

To Revwier
ESYX: W1-1

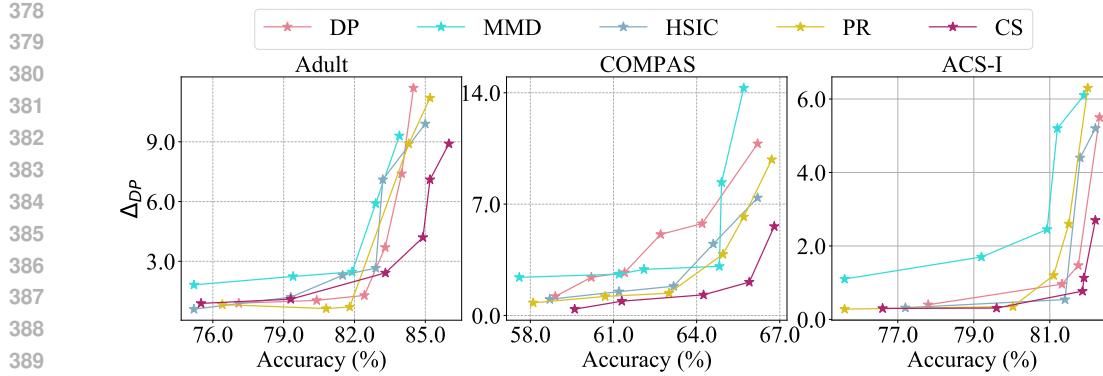


Figure 3: Fairness-accuracy trade-off curves on the test sets for (left) Adult, (middle) COMPAS, and (bottom) ACS-I. Ideally, results should be positioned in the bottom-right corner.

on the choice of kernel and bandwidth; in this work, we use an RBF kernel with the median heuristic, but exploring alternative kernels is an interesting direction for future work. Finally, our formulation specifically targets group fairness notions (e.g., DP and EO) and does not directly address other notions such as individual or causal fairness, which should be taken into account when deciding whether CS is suitable for a given application.

To Revvier LT9c:
W1

Choice of kernel. Throughout this work, we instantiate the empirical CS divergence with a Gaussian (RBF) kernel and choose its bandwidth using the median heuristic, following common practice in kernel-based dependence measures such as MMD and HSIC (Gretton et al., 2012; 2005). This choice is kept fixed across datasets and models to ensure a controlled comparison between regularizers. In principle, our framework is not restricted to a particular kernel family: alternative kernels (e.g., Laplacian or polynomial) can be plugged into the same estimator without changing the overall training objective or algorithm (Schölkopf & Smola, 2002; Shawe-Taylor & Cristianini, 2004). A more systematic investigation of how different kernel families and bandwidth-selection strategies affect the utility-fairness trade-off is an interesting direction for future work.

To Revvier Pqj:
W2-2

5 EXPERIMENTS

In this section, we evaluate the effectiveness of the CS fairness regularizer from several perspectives: (1) utility and fairness performance, (2) the tradeoff between utility and fairness, (3) prediction distributions across different sensitive groups, (4) T-SNE plots for these sensitive groups, and (5) the sensitivity of parameters in Equation (13). Our evaluation encompasses five datasets with diverse sensitive attributes, including four tabular datasets: Adult, COMPAS, ACS-I, and ACS-T, as well as one image dataset, CelebA-A. Utility performance is assessed based on accuracy and the area under the curve (AUC), while fairness performance is measured using Δ_{DP} Equation (2) and Δ_{EO} Equation (3). Detailed information about the datasets and baselines can be found in the Appendix. We denote an observation drawn from the results as **Obs..**

5.1 FAIRNESS AND UTILITY PERFORMANCE

We conducted experiments on five datasets along with their corresponding baselines, as previously mentioned. For each dataset, we performed 10 different splits to ensure robustness in our results. We calculated the mean and standard deviation for each metric across these splits. The accuracy and fairness performance of the downstream tasks is in Table 2. Our observations are as follows:

Obs. 1: CS achieves the best or near-best Δ_{EO} on most datasets and is competitive on Δ_{DP} , while maintaining high utility. Notably, CS demonstrates exceptional fairness performance on the image dataset, CelebA-A, where the disparity in the ‘Young’ and ‘Non-Young’ groups sees a Δ_{DP} reduction of 97.36% and a Δ_{EO} reduction of 98.58%. Furthermore, in the Adult and ACS-I datasets, which include gender groups, traditional methods such as DP, MMD, HSIC, and PR do not effectively optimize for EO fairness. In contrast, the proposed CS achieves significant reductions in Δ_{EO} by 72.12% and 63.85%, respectively, compared to MLP.

To Revvier Pqj:
W4

Obs. 2: CS achieves good fairness performance with a small sacrifice in utility. Specifically, CS exhibits a decrease of less than 3.1% in accuracy and less than 2.2% in AUC. The only exception is observed with COMPAS when gender is treated as a sensitive attribute, resulting in a slightly higher

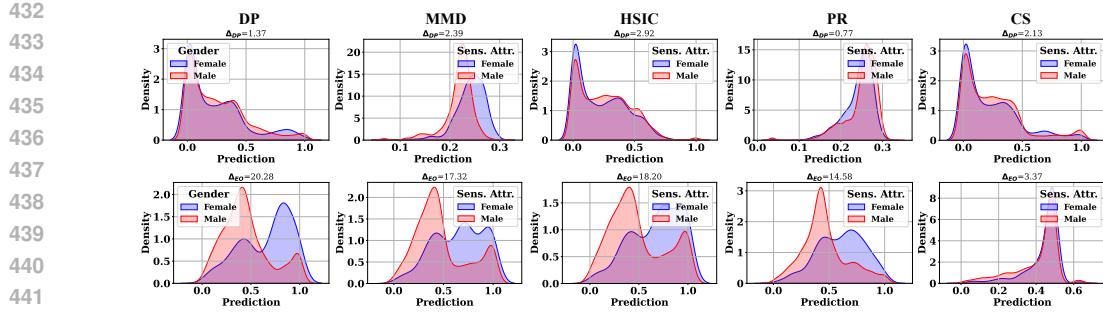


Figure 4: Prediction distributions for female and male groups in the Adult dataset. The top row shows kernel density estimates of the raw predictions \hat{Y} for all target labels, grouped by gender, while the bottom row shows the prediction densities for the positive class, $\hat{Y} = 1$, for the two gender groups. Each column corresponds to a different fairness regularizer. A larger overlap between the blue and red curves indicates better group fairness, and the reported values above each panel give the corresponding gaps in Δ_{DP} (top row) and Δ_{EO} (bottom row).

accuracy loss of 3.6%. Notably, CS demonstrates either equivalent or improved AUC performance, with increases of 0.02% and 0.58% on Adult for the gender and race groups, respectively, as well as a 0.35% increase on COMPAS for the race group. Among the baselines, HSIC ranks highest in utility, achieving the best performance on ACS-I for the race group and on ACS-T for both the gender and race groups. This is followed by PR, which shows the best utility on COMPAS for both the gender and race groups, as well as on CelebA-A for the gender group.

5.2 HOW DO ACCURACY AND FAIRNESS TRADE-OFF IN BASELINE MODELS AND CS?

We evaluate the trade-off between accuracy and Δ_{DP} for the baselines by varying the fairness hyperparameters (Yao et al., 2023; Deka & Sutherland, 2023). The results are presented in Figure 3, where the x-axis represents the target accuracy, while the y-axis shows the average Demographic Parity (DP) across both positive and negative target classes. It is important to note that the figure in the bottom right corner represents the optimal result.

Obs. 3: At the same utility level, CS is the most effective method in promoting fairness. Analyzing the results, we find that CS consistently achieves the lowest Δ_{DP} across most accuracy levels, with this effect becoming more pronounced at higher accuracy levels. This is evidenced by the significant gap in Δ_{DP} between CS and other baselines. It is important to note that while all baselines can demonstrate good fairness when the optimization prioritizes fairness over task objectives, the task objective remains critical for the practical application of these models. **Obs. 4: High accuracy can sometimes lead to worse fairness compared to MLP, as the fairness objective becomes more challenging to optimize when there is a stronger focus on task-specific objectives.** As shown in Table 2, the Δ_{DP} for MMD is over 14.0, which is greater than the average Δ_{DP} of 13.22 for MLP. However, these fairness regularizers generally prove effective in controlling bias in representations, especially when more emphasis is placed on the task-specific objective. Notably, some datasets with particularly sensitive attributes pose greater challenges for achieving fairness. For instance, the COMPAS dataset, which includes gender as a sensitive attribute, demonstrates this difficulty. One possible explanation is the relatively small sample size of COMPAS, which contains only 6,172 samples, significantly fewer than other datasets where fairness is easier to achieve. **Obs. 5: CS displays a significant increase in Δ_{DP} at a slower rate than other baselines as accuracy increases.** We analyze the slope of the lines representing the increase in Δ_{DP} with rising accuracy. Many methods, such as PR and DP, demonstrate strong fairness performance at low accuracy levels; however, they quickly lose control over fairness as accuracy begins to increase. This is evident from the abrupt rise in Δ_{DP} observed at around 82.0% on Adult, 63.0% on COMPAS, and 81.0% on ACS-I. In contrast, CS only exhibits a sudden increase at 85.0%, 65.5%, and 81.5% for the same datasets, respectively.

5.3 HOW CAN THE CS FAIRNESS REGULARIZER PERFORM WELL ON BOTH DP AND EO?

We visualize the kernel density estimate plot² of the predictions \hat{Y} across different sensitive groups to analyze how CS achieves a better balance of various fairness definitions compared to other baselines.

²<https://seaborn.pydata.org/generated/seaborn.kdeplot.html>

486 The *first row* displays the predictions for all target classes, specifically $Y = 0$ and $Y = 1$, grouped
 487 by sensitive attributes. In this row, the blue areas represent the prediction density for $S = 0$, while
 488 the red areas indicate the prediction density for $S = 1$. The *second row* illustrates the prediction
 489 density for the positive target class, $Y = 1$, across two different sensitive groups. Figure 4
 490 presents the results for Adult based on gender and race groups, with additional results for other datasets
 491 available in Appendix C.3.

492 **Obs. 6: CS effectively optimizes the prediction distributions for the two sensitive groups,
 493 specifically $\hat{Y}|S = 0$ and $\hat{Y}|S = 1$. Additionally, it optimizes the prediction distributions
 494 for these groups within the positive target group, i.e., $\hat{Y}|S = 0, Y = 1$ and $\hat{Y}|S = 1, Y = 1$.**
 495 Achieving DP and EO fairness requires different objectives. For instance, DP directly optimizes the
 496 Δ_{DP} , which results in reduced effectiveness for achieving EO fairness. This is evident across all
 497 datasets, as DP ranks among the worst, achieving 7/10 of the lowest EO fairness scores on Δ_{EO}
 498 when tested on five datasets with two types of sensitive attributes. The distribution plots for DP
 499 further illustrate this, showing a generally larger gap between the two sensitive groups in the EO
 500 plots compared to other methods. In contrast, CS consistently minimizes the prediction density gap
 501 between the two sensitive groups.

502 5.4 PARAMETER SENSITIVITY ANALYSIS

503 For all models, we tune the hyperparameters using cross-validation on the training set. The hyperparameters for these variants are determined through grid search during cross-validation.
 504 Specifically, we vary the parameters α and β in Equation (13) across the ranges $(1e-6, 150)$ and $(1e-3, 10)$, respectively. In this experiment, we specifically visualize the values of α in the range $(1e-4, 1e-1)$ for CS.

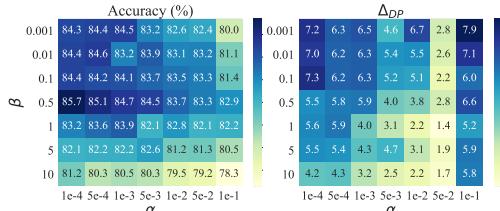
505 The heatmap in Figure 5 illustrates the accuracy and Δ_{DP} across various combinations of α and
 506 β values for the Adult. In the accuracy plots, darker colors indicate higher values, which are
 507 preferable, while lighter colors in the Δ_{DP} plots represent better fairness performance.

508 **Obs. 7: The highest accuracy is achieved when α is set to its smallest value, $1e-4$, while the
 509 best fairness is obtained with $\alpha = 5e-2$.** Notably, fairness drops significantly when α increases
 510 from $5e-2$ to $1e-1$. Generally, smaller values of α can still yield satisfactory fairness performance
 511 when paired with an appropriate range of β , specifically around $5 - 10$. **Obs. 8: The fairness
 512 performance is more sensitive to changes in α than in β .** For instance, adjusting β from $1e-3$ to
 513 10 , which represents a $10,000\times$ increase, results in only a slight decrease in Δ_{DP} from 7.2 to 4.2 .
 514 In contrast, increasing α from $1e-2$ to $5e-2$, a $5\times$ change, leads to a significant drop in Δ_{DP}
 515 from 6.7 to 2.8 , when keeping β fixed at $1e-3$.

526 6 CONCLUSION

527 In this paper, we introduce a novel fair machine learning method called the Cauchy-Schwarz (CS)
 528 fairness regularizer. Empirically, our approach achieves a more consistent utility-fairness trade-off
 529 across hyperparameter settings than standard regularizers, and yields more generalizable fairness by
 530 minimizing the Cauchy-Schwarz divergence between the prediction distribution and the sensitive
 531 attributes. We demonstrate that the CS divergence provides a tighter bound compared to both the
 532 Kullback-Leibler divergence and the Maximum Mean Discrepancy, as well as the mean disparity used
 533 in Demographic Parity regularization. This superiority is particularly evident when the distributions
 534 are significantly different or when there is substantial variation in the scale of the embeddings. As
 535 a result, our CS fairness regularizer delivers improved fairness performance in practical scenarios.
 536 While our work currently only evaluates on general machine learning tasks, and thus leave future
 537 work to other tasks such as graph learning.

To Rewvier JuMp:
W1



538 **Figure 5: Parameter sensitivity of the CS regularizer on ADULT:** heatmaps show test accuracy
 539 (left) and Δ_{DP} (right) as the fairness weight α and ℓ_2 weight β vary over the cross-validated ranges.
 Overall, CS exhibits a smooth utility-fairness trade-off, remaining stable over a broad range of
 β and becoming noticeably more sensitive only when α is very large.

540
541
ETHICS STATEMENT

542 Our work aims to improve fairness in machine learning by introducing a general and model-agnostic
 543 regularization method that reduces statistical dependence between predictions and sensitive attributes.
 544 By encouraging independence through Cauchy–Schwarz (CS) Divergence, our framework helps
 545 mitigate systematic biases that often lead to disparate treatment of demographic groups. We use
 546 publicly available datasets containing sensitive attributes (e.g., gender, race) solely for the purpose
 547 of evaluating fairness interventions. We acknowledge that fairness is a multidimensional concept
 548 and that no single method can fully eliminate all forms of bias. While our method improves several
 549 group fairness metrics, we caution against deploying fairness-enhancing techniques without thorough
 550 domain-specific evaluation and stakeholder engagement.

551
552
REPRODUCIBILITY STATEMENT
553

554 We ensure the reproducibility of our results by providing the following: (1) a comprehensive
 555 description of the CS Fairness Regularization method, including the formulation of the objective
 556 function and optimization procedure; (2) details on hyperparameters, model architectures, and training
 557 procedures used in all experiments; (3) evaluation on five benchmark datasets using multiple fairness
 558 metrics and baselines; and (4) open-sourcing our code and data preprocessing scripts upon publication.
 559 We follow standard experimental protocols and report average results across multiple runs to account
 560 for randomness. These practices facilitate independent verification and encourage future research on
 561 robust fairness methods.

562
563
REFERENCES

564 Tameem Adel, Isabel Valera, Zoubin Ghahramani, and Adrian Weller. One-network adversarial
 565 fairness. In *Proceedings of the AAAI*, pp. 2412–2420, 2019.

566 Chirag Agarwal, Himabindu Lakkaraju, and Marinka Zitnik. Towards a unified framework for fair
 567 and stable graph representation learning. In *UAI 2021: Uncertainty in Artificial Intelligence*, 2021.

568 Wael Alghamdi, Hsiang Hsu, Haewon Jeong, Hao Wang, Peter Michalak, Shahab Asoodeh, and
 569 Flavio Calmon. Beyond adult and compas: Fair multi-class prediction via information projection.
 570 *Advances in Neural Information Processing Systems*, 35:38747–38760, 2022.

571 Sina Baharlouei, Maher Nouiehed, Ahmad Beirami, and Meisam Razaviyayn. Rényi fair inference.
 572 In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HkgsUJrtDB>.

573 Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness in machine learning. *NeurIPS tutorial*,
 574 1:2017, 2017.

575 Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. Data decisions and theoretical implications when
 576 adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*, 2017.

577 Steinar Bøyum. Fairness in education—a normative analysis of oecd policy documents. *Journal of
 578 Education Policy*, 29(6):856–870, 2014.

579 Maarten Buyl and Tijl De Bie. Optimal transport of classifiers to fairness. In *Advances in Neural
 580 Information Processing Systems*, 2022.

581 Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R
 582 Varshney. Optimized pre-processing for discrimination prevention. In I. Guyon, U. Von Luxburg,
 583 S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural
 584 Information Processing Systems*, volume 30. Curran Associates, Inc., 2017a. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/9a49a25d845a483fae4be7e341368e36-Paper.pdf.

585 Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R
 586 Varshney. Optimized pre-processing for discrimination prevention. *NeurIPS*, 30, 2017b.

594 Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism
 595 prediction instruments. *Big data*, 5(2):153–163, 2017.
 596

597 Ching-Yao Chuang and Youssef Mroueh. Fair mixup: Fairness via interpolation. In *ICLR*, 2020.
 598

599 Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
 600

601 Enyan Dai and Suhang Wang. Say no to the discrimination: Learning fair graph neural networks
 602 with limited sensitive attribute information. In *Proceedings of the 14th ACM WSDM*, pp. 680–688,
 603 2021.

604 Namrata Deka and Danica J Sutherland. Mmd-b-fair: Learning fair representations with statistical
 605 testing. In *International Conference on Artificial Intelligence and Statistics*, pp. 9564–9576. PMLR,
 606 2023.

607 Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair
 608 machine learning. *NeurIPS*, 2021.
 609

610 Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. Edits: Modeling and mitigating data bias
 611 for graph neural networks. In *Proceedings of the ACM web conference 2022*, pp. 1259–1269, 2022.
 612

613 Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
 614

615 Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through
 616 awareness. In *Conference on Innovations in Theoretical Computer Science (ITCS)*, 2012.
 617

618 Harrison Edwards and Amos Storkey. Censoring representations with an adversary. *arXiv preprint
 619 arXiv:1511.05897*, 2015.
 620

621 Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubra-
 622 manian. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD
 623 international conference on knowledge discovery and data mining*, pp. 259–268, 2015.

624 Hans Gebelein. Das statistische problem der korrelation als variations-und eigenwertproblem und
 625 sein zusammenhang mit der ausgleichsrechnung. *ZAMM-Journal of Applied Mathematics and
 626 Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 21(6):364–379, 1941.
 627

628 Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical
 629 dependence with hilbert-schmidt norms. In *International conference on algorithmic learning
 630 theory*, pp. 63–77. Springer, 2005.

631 Arthur Gretton, Kenji Fukumizu, Choon Teo, Le Song, Bernhard Schölkopf, and Alex Smola. A
 632 kernel statistical test of independence. *Advances in neural information processing systems*, 20,
 633 2007.

634 Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A
 635 kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012.
 636

637 Thomas Grote and Geoff Keeling. Enabling fairness in healthcare through machine learning. *Ethics
 638 and Information Technology*, 24(3):39, 2022.
 639

640 Xiaotian Han, Zhimeng Jiang, Hongye Jin, Zirui Liu, Na Zou, Qifan Wang, and Xia Hu. Retiring
 641 $\$ \backslash \delta \backslash \text{text}\{DP\} \$$: New distribution-level metrics for demographic parity. *Transactions on
 642 Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=LjDFIWVVA>.
 643

644 Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances
 645 in neural information processing systems*, 29, 2016.
 646

647 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 648 recognition. In *CVPR*, June 2016.

648 Hermann O Hirschfeld. A connection between correlation and contingency. In *Mathematical pro-
649 ceedings of the cambridge philosophical society*, volume 31, pp. 520–524. Cambridge University
650 Press, 1935.

651

652 Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

653

654 Lily Hu and Yiling Chen. A short-term intervention for long-term fairness in the labor market. In
655 *WWW*, pp. 1389–1398, 2018.

656

657 Robert Jenssen, Jose C Principe, Deniz Erdogmus, and Torbjørn Eltoft. The cauchy–schwarz
658 divergence and parzen windowing: Connections to graph theory and mercer kernels. *Journal of
the Franklin Institute*, 343(6):614–629, 2006.

659

660 Ray Jiang, Aldo Pacchiano, Tom Stepleton, Heinrich Jiang, and Silvia Chiappa. Wasserstein fair
661 classification. In *Uncertainty in Artificial Intelligence*, pp. 862–872. PMLR, 2020.

662

663 Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimi-
664 nation. *Knowledge and information systems*, 33(1):1–33, 2012.

665

666 Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-aware classifier with
667 prejudice remover regularizer. In *ECML PKDD*. Springer, 2012.

668

669 Kittipat Kampa, Erion Hasanbelliu, and Jose C Principe. Closed-form cauchy–schwarz pdf divergence
670 for mixture of gaussians. In *IJCNN*, pp. 2578–2585. IEEE, 2011.

671

672 Jian Kang, Jingrui He, Ross Maciejewski, and Hanghang Tong. Inform: Individual fairness on
673 graph mining. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge
674 Discovery and Data Mining*, pp. 379–389, 2020.

675

676 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint
arXiv:1412.6980*, 2014.

677

678 René F Kizilcec and Hansol Lee. Algorithmic fairness in education. In *The ethics of artificial
679 intelligence in education*, pp. 174–202. Routledge, 2022.

680

681 Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determi-
682 nation of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.

683

684 Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *NeurIPS*,
685 30, 2017.

686

687 Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. Propublica compas analysis—data and
688 analysis for ‘machine bias’. <https://github.com/propublica/compas-analysis>,
689 2016. Accessed: 2023-03-13.

690

691 Tai Le Quy, Arjun Roy, Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsi. A survey on datasets for
692 fairness-aware machine learning. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge
693 Discovery*, 12(3):e1452, 2022.

694

695 Ana Alves Leal. Algorithms, creditworthiness, and lending decisions. In *International Conference
696 on Autonomous Systems and the Law*, pp. 321–353. Springer, 2022.

697

698 Peizhao Li, Yifei Wang, Han Zhao, Pengyu Hong, and Hongfu Liu. On dyadic fairness: Exploring and
699 mitigating bias in graph connections. In *International Conference on Learning Representations*,
700 2020.

701

Zhu Li, Adrian Perez-Suay, Gustau Camps-Valls, and Dino Sejdinovic. Kernel dependence
702 regularizers and gaussian processes with applications to algorithmic fairness. *arXiv preprint
arXiv:1911.04322*, 2019.

Hongyi Ling, Zhimeng Jiang, Youzhi Luo, Shuiwang Ji, and Na Zou. Learning fair graph representa-
703 tions via automated data augmentations. In *ICLR*, 2023.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In
704 *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

702 Christos Louizos, Kevin Swersky, Yujia Li, Max Welling, and Richard S Zemel. The variational fair
 703 autoencoder. In *International Conference on Learning Representations (ICLR)*, 2016.

704

705 Gilles Louppe, Michael Kagan, and Kyle Cranmer. Learning to pivot with adversarial networks.
 706 *NeurIPS*, 30, 2017.

707

708 David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and
 709 transferable representations. *International Conference on Machine Learning*, 2018.

710

711 Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey
 712 on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.

713

714 Anay Mehrotra and Nisheeth Vishnoi. Fair ranking with noisy protected attributes. *Advances in
 715 Neural Information Processing Systems*, 35:31711–31725, 2022.

716

717 Debarghya Mukherjee, Mikhail Yurochkin, Moulinath Banerjee, and Yuekai Sun. Two simple ways
 718 to learn individual fairness metrics from data. In *ICML*, pp. 7097–7107. PMLR, 2020.

719

720 Debarghya Mukherjee, Felix Petersen, Mikhail Yurochkin, and Yuekai Sun. Domain adaptation
 721 meets individual fairness. and they get along. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave,
 722 and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL
<https://openreview.net/forum?id=XSNfXG9HBAu>.

723

724 Emanuel Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3):1065–1076, 1962.

725

726 JC Principe, Dongxin Xu, J Fisher, and S Haykin. Information theoretic learning. unsupervised
 727 adaptive filtering. *Unsupervised Adapt Filter*, 1, 2000.

728

729 Alfréd Rényi. On measures of dependence. *Acta mathematica hungarica*, 10(3-4):441–451, 1959.

730

731 Luis G Sanchez Giraldo, Erion Hasanbelliu, Murali Rao, and Jose C Principe. Group-wise point-set
 732 registration based on renyi’s second order entropy. In *CVPR*, pp. 6693–6701, 2017.

733

734 Bernhard Schölkopf and Alexander J Smola. *Learning with kernels: support vector machines,
 735 regularization, optimization, and beyond*. MIT press, 2002.

736

737 John Shawe-Taylor and Nello Cristianini. *Kernel methods for pattern analysis*. Cambridge university
 738 press, 2004.

739

740 Changjian Shui, Gezheng Xu, Qi Chen, Jiaqi Li, Charles X Ling, Tal Arbel, Boyu Wang, and Christian
 741 Gagné. On learning fairness and accuracy on multiple subgroups. *Advances in Neural Information
 742 Processing Systems*, 35:34121–34135, 2022.

743

744 Gábor J Székely, Maria L Rizzo, and Nail K Bakirov. Measuring and testing dependence by
 745 correlation of distances. 2007.

746

747 Linh Tran, Maja Pantic, and Marc Peter Deisenroth. Cauchy–schwarz regularized autoencoder.
 748 *Journal of Machine Learning Research*, 23(115):1–37, 2022.

749

750 Daniel J Trosten, Sigurd Lokse, Robert Jenssen, and Michael Kampffmeyer. Reconsidering represen-
 751 tation alignment for multi-view clustering. In *CVPR*, pp. 1255–1265, 2021.

752

753 Ioannis Tsakousis and Mohammed H Alghamdi. Examining academic performance across gender
 754 differently: Measurement invariance and latent mean differences using bias-corrected bootstrap
 755 confidence intervals. *Frontiers in Psychology*, 13:896638, 2022.

756

757 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine
 758 learning research*, 9(11), 2008.

759

760 Sahil Verma and Julia Rubin. Fairness definitions explained. In *2018 ieee/acm international workshop
 761 on software fairness (fairware)*, pp. 1–7. IEEE, 2018.

756 Yao Yao, Qihang Lin, and Tianbao Yang. Stochastic methods for auc optimization subject to auc-
 757 based fairness constraints. In *International Conference on Artificial Intelligence and Statistics*, pp.
 758 10324–10342. PMLR, 2023.

760 Wenzhe Yin, Shujian Yu, Yicong Lin, Jie Liu, Jan-Jakob Sonke, and Efstratios Gavves. Domain
 761 adaptation with cauchy-schwarz divergence. *arXiv preprint arXiv:2405.19978*, 2024.

762 Shujian Yu, Hongming Li, Sigurd Løkse, Robert Jenssen, and José C Príncipe. The conditional
 763 cauchy-schwarz divergence with applications to time-series data and sequential decision making.
 764 *arXiv preprint arXiv:2301.08970*, 2023.

766 Shujian Yu, Xi Yu, Sigurd Løkse, Robert Jenssen, and Jose C Príncipe. Cauchy-schwarz divergence
 767 information bottleneck for regression. *arXiv preprint arXiv:2404.17951*, 2024.

769 Shujian Yu, Hongming Li, Sigurd Løkse, Robert Jenssen, and José C Príncipe. The conditional
 770 cauchy-schwarz divergence with applications to time-series data and sequential decision making.
 771 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

772 Mikhail Yurochkin and Yuekai Sun. Sensei: Sensitive set invariance for enforcing individual fairness.
 773 *arXiv preprint arXiv:2006.14168*, 2020.

775 Mikhail Yurochkin, Amanda Bower, and Yuekai Sun. Training individually fair ml models with
 776 sensitive subspace robustness. *arXiv preprint arXiv:1907.00020*, 2019.

778 Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. Fair-
 779 ness constraints: Mechanisms for fair classification. In *International Conference on Artificial
 780 Intelligence and Statistics (AISTATS)*, 2017.

781 Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations.
 782 In *International conference on machine learning*, pp. 325–333. PMLR, 2013.

784 Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial
 785 learning. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2018.

786 Guanhua Zhang, Yihua Zhang, Yang Zhang, Wenqi Fan, Qing Li, Sijia Liu, and Shiyu Chang.
 787 Fairness reprogramming. *arXiv preprint arXiv:2209.10222*, 2022.

789 Ji Zhao and Deyu Meng. Fastmmd: Ensemble of circular discrepancy for efficient two-sample test.
 790 *Neural computation*, 27(6):1345–1372, 2015.

792 Aoqi Zuo, Susan Wei, Tongliang Liu, Bo Han, Kun Zhang, and Mingming Gong. Counterfactual
 793 fairness with partially known causal graph. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave,
 794 and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL
 795 <https://openreview.net/forum?id=9aLbntHz1Uq>.

797 LLM DISCLOSURE

799 We used an LLM (OpenAI ChatGPT) only for copyediting—streamlining phrasing and correcting
 800 grammar, spelling, and stylistic inconsistencies. The model played no role in conceiving the research,
 801 designing methods, implementing code, running experiments, analyzing results, or shaping claims.
 802 All edits were reviewed by the authors, and only manuscript text was provided to the tool.

804 A WHAT IS THE RELATIONSHIP BETWEEN CS DIVERGENCE AND EXISTING 805 DISTRIBUTION DISTANCE MEASURES?

808 To illustrate the advantages of the CS fairness regularizer, we begin by summarizing the com-
 809 monly used distribution distance metrics: Maximum Mean Discrepancy (MMD), Kullback-Leibler
 divergence (KL), and Hilbert-Schmidt Independence Criterion (HSIC).

810
Demographic Parity Regularizer. The demographic parity regularizer is widely utilized in fairness-
811 focused machine learning studies (Chuang & Mroueh, 2020). It aims to optimize the mean disparity
812 between two *prediction distributions*. This regularizer can be formally expressed as:
813

$$814 \quad \text{DP}(p; q) = \left| \frac{1}{N_1} \sum_i^{N_1} p(\mathbf{x}_i) - \frac{1}{N_2} \sum_j^{N_2} q(\mathbf{x}_j) \right|, \quad (16)$$

$$815$$

$$816$$

817 where \mathbf{x}_i and \mathbf{x}_j are data points from $S = 0$ and $S = 1$, in the context of fairness. In the following,
818 we represent \mathbf{x}_i with distribution p and \mathbf{x}_j with distribution q as \mathbf{x}_i^p and \mathbf{x}_i^q for simplicity. However,
819 only optimizing on the mean disparity of two distributions cannot always generate an optimized DP
820 or EO, as the Equation (16) equals 0 is a necessary but not sufficient condition for achieving DP and
821 EO.

822 **Mean Maximum Discrepancy.** One of the most widely used distance metrics is the Mean Maximum
823 Discrepancy (MMD) (Gretton et al., 2012). In the context of fairness, previous studies have employed
824 MMD as a regularizer to enforce statistical parity among the *embeddings* of different sensitive groups
825 within a machine learning model (Deka & Sutherland, 2023; Louizos et al., 2016). This approach
826 aims to facilitate fair representation learning.

$$827 \quad \widetilde{\text{MMD}}^2(p; q) = \frac{1}{N_1^2} \sum_{i,j=1}^{N_1} \kappa(\mathbf{x}_i^p, \mathbf{x}_j^p) + \frac{1}{N_2^2} \sum_{i,j=1}^{N_2} \kappa(\mathbf{x}_i^q, \mathbf{x}_j^q) \\ 828 \quad - \frac{2}{N_1 N_2} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \kappa(\mathbf{x}_i^p, \mathbf{x}_j^q). \quad (17)$$

$$829$$

$$830$$

$$831$$

$$832$$

$$833$$

834 By comparing with Equation (20), we observe that the CS divergence introduces a logarithmic term
835 for each component of the MMD. Through simple transformations, we can deduce the following:

836 *Remark A.1.* CS divergence measures the cosine distance between empirical mean embedding
837 $\mu_p = \frac{1}{N_1} \sum_{i=1}^{N_1} f(\mathbf{x}_i^p)$ and $\mu_q = \frac{1}{N_2} \sum_{j=1}^{N_2} f(\mathbf{x}_j^q)$ in a Reproducing Kernel Hilbert Space, while
838 MMD utilizes Euclidean distance.

839 **Kullback-Leibler Divergence.** Kullback-Leibler (KL) Divergence is a key concept in information
840 bottleneck theory, where it is used to quantify the mutual information between two probability
841 distributions. This metric has gained popularity across various domains, including fair machine
842 learning (Kamishima et al., 2012).

$$843 \quad D_{\text{KL}} = \int p(\mathbf{x}) \log \left(\frac{p(\mathbf{x})}{q(\mathbf{x})} \right) \quad (18)$$

$$844$$

$$845$$

$$846$$

847 **Hilbert-Schmidt Independence Criterion (HSIC).** Let K and L denote the Gram matrices for the
848 variables x and y , respectively. Specifically, K is defined such that $K_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$, and L is defined
849 as $L_{ij} = \kappa(\mathbf{y}_i, \mathbf{y}_j)$, where κ is the Gaussian kernel function given by $\kappa = \exp \left(-\frac{\|\cdot\|^2}{2\sigma^2} \right)$. The Hilbert-
850 Schmidt Independence Criterion (HSIC) can be estimated using the following expression (Gretton
851 et al., 2007):
852

$$853 \quad \widetilde{\text{HSIC}}(p; q) = \frac{1}{N^2} \sum_{i,j}^N K_{ij} L_{ij} + \frac{1}{N^4} \sum_{i,j,q,r}^N K_{ij} L_{qr} \\ 854 \quad - \frac{2}{N^3} \sum_{i,j,q}^N K_{ij} L_{iq} = \frac{1}{N^2} \text{tr}(KHLH), \quad (19)$$

$$855$$

$$856$$

$$857$$

$$858$$

859 where $H = I - \frac{1}{N} \mathbf{1} \mathbf{1}^T$ represents a centering matrix of size $N \times N$. In this expression, I is the
860 identity matrix, $\mathbf{1}$ is a vector of ones, and $\frac{1}{N} \mathbf{1} \mathbf{1}^T$ computes the average across the columns, effectively
861 centering the data by subtracting the mean from each entry.

862 Compared to Equation (17), The HSIC can be interpreted as the MMD between the joint distribution
863 $p(\mathbf{x}, \mathbf{y})$ and the product of their marginal distributions $p(\mathbf{x})p(\mathbf{y})$.

864 **B DETAILS ON THE RELATION OF CS AND EXISTING FAIRNESS**
 865 **REGULARIZERS**
 866

867 **B.1 PROOF OF PROPOSITION 4.1**
 868

869 **Proposition 1.** *Given two sets of observations $\{\mathbf{x}_i^p\}_{i=1}^{N_1}$ and $\{\mathbf{x}_j^q\}_{j=1}^{N_2}$, let p and q denote the*
 870 *distributions of two groups. The empirical estimator of the CS divergence $D_{CS}(p; q)$ is given by:*
 871

$$872 \quad \tilde{D}_{CS}(p; q) = \log \left(\frac{1}{N_1^2} \sum_{i,j=1}^{N_1} \kappa(\mathbf{x}_i^p, \mathbf{x}_j^p) \right) + \log \left(\frac{1}{N_2^2} \sum_{i,j=1}^{N_2} \kappa(\mathbf{x}_i^q, \mathbf{x}_j^q) \right) \\ 873 \quad - 2 \log \left(\frac{1}{N_1 N_2} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \kappa(\mathbf{x}_i^p, \mathbf{x}_j^q) \right). \quad (20)$$

874

875 *Proof.* The CS divergence is defined as:
 876

$$877 \quad D_{CS}(p; q) = -\log \left(\frac{(\int p(\mathbf{x})q(\mathbf{x}) d\mathbf{x})^2}{\int p(\mathbf{x})^2 d\mathbf{x} \int q(\mathbf{x})^2 d\mathbf{x}} \right), \quad (21)$$

878

879 where $\hat{p}(\mathbf{x}) = \frac{1}{M} \sum_{i=1}^M \kappa_\sigma(\mathbf{x} - \mathbf{x}_i^p)$ and $\hat{q}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \kappa_\sigma(\mathbf{x} - \mathbf{x}_i^q)$ are kernel density estimation.
 880

881 Then we can obtain:
 882

$$883 \quad \int \hat{p}^2(\mathbf{x}) d\mathbf{x} = \frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M \kappa_{\sqrt{2}\sigma}(\mathbf{x}_i^p - \mathbf{x}_j^p). \quad (22)$$

884

885 By a similar approach,
 886

$$887 \quad \int \hat{q}^2(\mathbf{x}) d\mathbf{x} = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \kappa_{\sqrt{2}\sigma}(\mathbf{x}_i^q - \mathbf{x}_j^q), \quad (23)$$

888

889 and
 890

$$891 \quad \int \hat{p}(\mathbf{x})\hat{q}(\mathbf{x}) d\mathbf{x} = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N \kappa_{\sqrt{2}\sigma}(\mathbf{x}_i^p - \mathbf{x}_j^q). \quad (24)$$

892

893 Substituting (Equation (22))-(Equation (24)) into Eq. (21), we obtain:
 894

$$895 \quad \tilde{D}_{CS}(p; q) = \log \left(\frac{1}{M^2} \sum_{i,j=1}^M \kappa_{\sqrt{2}\sigma}(\mathbf{x}_i^p - \mathbf{x}_j^p) \right) + \log \left(\frac{1}{N^2} \sum_{i,j=1}^N \kappa_{\sqrt{2}\sigma}(\mathbf{x}_i^q - \mathbf{x}_j^q) \right) \\ 896 \quad - 2 \log \left(\frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N \kappa_{\sqrt{2}\sigma}(\mathbf{x}_i^p - \mathbf{x}_j^q) \right). \quad (25)$$

897

898 \square
 899

900 **B.2 PROOF OF REMARK A.1**
 901

902 **Remark 1.** *CS divergence measures the cosine distance between μ_p and μ_q in a Reproducing Kernel*
 903 *Hilbert Space, while MMD utilizes Euclidean distance.*
 904

905 *Proof.* Let \mathcal{H} be a Reproducing Kernel Hilbert Space (RKHS) associated with a kernel $\kappa(\mathbf{x}_i^p, \mathbf{x}_j^q) =$
 906 $\langle f(\mathbf{x}_i^p), f(\mathbf{x}_j^q) \rangle_{\mathcal{H}}$ (Yu et al., 2024). The mean embeddings of two distributions p and q in \mathcal{H} are
 907 denoted by $\mu_p = \frac{1}{N_1} \sum_{i=1}^{N_1} f(\mathbf{x}_i^p)$ and $\mu_q = \frac{1}{N_2} \sum_{j=1}^{N_2} f(\mathbf{x}_j^q)$ in \mathcal{H} , respectively. The CS divergence
 908 defined by Equation (20) can thus be written as:
 909

$$910 \quad \tilde{D}_{CS}(p; q) = -2 \log \frac{\langle \mu_p, \mu_q \rangle_{\mathcal{H}}}{\|\mu_p\|_{\mathcal{H}} \|\mu_q\|_{\mathcal{H}}} = -2 \log D_{\text{cos}}(\mu_p, \mu_q)$$

911

918 Here, $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ denotes the inner product in the RKHS, and $\|\cdot\|_{\mathcal{H}}$ represents the norm induced by the
 919 inner product. The mean embeddings $\boldsymbol{\mu}_p$ and $\boldsymbol{\mu}_q$ are elements of \mathcal{H} . Thus, the CS divergence is
 920 computed based on the cosine distance D_{COS} between $\boldsymbol{\mu}_p$ and $\boldsymbol{\mu}_q$.

921 Similarly, the Maximum Mean Discrepancy (MMD) between distributions p and q defined in Equation
 922 (17) can be written as:

$$924 \quad \text{MMD}^2(p, q) = \|\boldsymbol{\mu}_p - \boldsymbol{\mu}_q\|_{\mathcal{H}}^2 = D_{\text{EUC}}(\boldsymbol{\mu}_p, \boldsymbol{\mu}_q).$$

926 Thus, the MMD measures the Euclidean distance between the mean embeddings of p and q in the
 927 RKHS \mathcal{H} , i.e., the $\boldsymbol{\mu}_p$ and $\boldsymbol{\mu}_q$. \square

929 B.3 PROOF OF PROPOSITION 4.2

931 **Proposition 2.** For any d -variate Gaussian distributions $p \sim \mathcal{N}(\boldsymbol{\mu}_p, \Sigma_p)$ and $q \sim \mathcal{N}(\boldsymbol{\mu}_q, \Sigma_q)$ with
 932 positive definite Σ_p and Σ_q , the following inequality holds:

$$933 \quad D_{\text{CS}}(p; q) \leq D_{\text{KL}}(p; q) \text{ and } D_{\text{CS}}(p; q) \leq D_{\text{KL}}(q; p). \quad (26)$$

935 *Proof.* The KL divergence for p and q is given by:

$$937 \quad D_{\text{KL}}(p; q) = \frac{1}{2} \left(\text{tr}(\Sigma_q^{-1} \Sigma_p) - d + (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p)^\top \Sigma_q^{-1} (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p) + \log \left(\frac{|\Sigma_q|}{|\Sigma_p|} \right) \right). \quad (27)$$

940 The CS divergence is expressed as (Kampa et al., 2011):

$$942 \quad D_{\text{CS}}(p; q) = -\log(d_{xy}) + \frac{1}{2} \log(d_{xx}) + \frac{1}{2} \log(d_{yy}), \quad (28)$$

944 where:

$$945 \quad d_{pq} = \frac{\exp \left(-\frac{1}{2} (\boldsymbol{\mu}_p - \boldsymbol{\mu}_q)^\top (\Sigma_p + \Sigma_q)^{-1} (\boldsymbol{\mu}_p - \boldsymbol{\mu}_q) \right)}{\sqrt{(2\pi)^d |\Sigma_p + \Sigma_q|}}, \quad (30)$$

$$948 \quad d_{pp} = \frac{1}{\sqrt{(2\pi)^d |\Sigma_p|}}, \quad d_{qq} = \frac{1}{\sqrt{(2\pi)^d |\Sigma_q|}}. \quad (31)$$

950 We simplify:

$$953 \quad D_{\text{CS}}(p; q) = \frac{1}{2} (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p)^\top (\Sigma_p + \Sigma_q)^{-1} (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p) + \frac{1}{2} \log \left(\frac{|\Sigma_p + \Sigma_q|}{2^d \sqrt{|\Sigma_p| |\Sigma_q|}} \right). \quad (32)$$

955 When the mean vectors differ, based on the property (Horn & Johnson, 2012), $\Sigma_q^{-1} - (\Sigma_p + \Sigma_q)^{-1}$
 956 is positive semi-definite given $\Sigma_p = \Sigma_q$, we have:

$$958 \quad \begin{aligned} 2(D_{\text{CS}}(p; q) - D_{\text{KL}}(p; q)) \\ 959 = (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p)^\top (\Sigma_p + \Sigma_q)^{-1} (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p) \\ 960 - (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p)^\top \Sigma_q^{-1} (\boldsymbol{\mu}_q - \boldsymbol{\mu}_p) \leq 0. \end{aligned} \quad (33)$$

963 When the covariance matrices differ, let I be the d -dimensional identity matrix (Yin et al., 2024):

$$965 \quad \begin{aligned} 2(D_{\text{CS}}(p; q) - D_{\text{KL}}(p; q)) &= \log \left(\frac{|\Sigma_p + \Sigma_q|}{2^d \sqrt{|\Sigma_p| |\Sigma_q|}} \right) \\ 966 &- \log \left(\frac{|\Sigma_q|}{|\Sigma_p|} \right) - \text{tr}(\Sigma_q^{-1} \Sigma_p) + d \\ 967 &= -d \log 2 + \log (|\Sigma_q^{-1} \Sigma_p + I|) \\ 968 &+ \frac{1}{2} \log (|\Sigma_q^{-1} \Sigma_p|) - \text{tr}(\Sigma_q^{-1} \Sigma_p) + d. \end{aligned} \quad (34)$$

972 We have $|\Sigma_q^{-1}\Sigma_p| \leq \left(\frac{1}{d} \text{tr}(\Sigma_q^{-1}\Sigma_p)\right)^d$, and $|\Sigma_q^{-1}\Sigma_p + I| \leq \left(1 + \frac{1}{d} \text{tr}(\Sigma_q^{-1}\Sigma_p)\right)^d$. Thus, based
973 on Equation (34), we can obtain:
974

$$\begin{aligned} 975 \quad & 2(D_{\text{CS}}(p; q) - D_{\text{KL}}(p; q)) \\ 976 \quad & \leq -d \log 2 + d \log \left(1 + \frac{1}{d} \text{tr}(\Sigma_q^{-1}\Sigma_p)\right) \\ 977 \quad & + \frac{d}{2} \log \left(\frac{1}{d} \text{tr}(\Sigma_q^{-1}\Sigma_p)\right) - \text{tr}(\Sigma_q^{-1}\Sigma_p) + d. \end{aligned} \quad (35)$$

980 The combined Equation (33) and Equation (35), we can obtain:
981

$$982 \quad 2(D_{\text{CS}}(p; q) - D_{\text{KL}}(p; q)) \leq 0, \quad (36)$$

983 Similarly, we can obtain $2(D_{\text{CS}}(q; p) - D_{\text{KL}}(q; p)) \leq 0$. In conclusion, we conclude:
984

$$985 \quad D_{\text{CS}}(p; q) \leq D_{\text{KL}}(p; q) \text{ and } D_{\text{CS}}(p; q) \leq D_{\text{KL}}(q; p). \quad (37)$$

986 \square
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

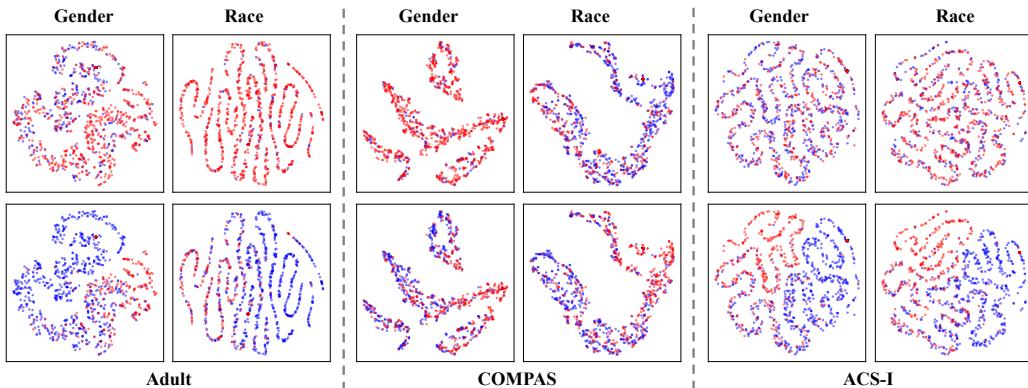
1026	1027	1028	Methods	Utility				Fairness				
				1029	1030	1031	1032	1033	1034	1035	1036	
1029	1030	1031	Gender	MLP	66.21 \pm 0.95	—	73.78 \pm 0.25	—	8.32 \pm 2.67	—	5.11 \pm 3.55	—
				DP	65.38 \pm 0.29	-1.25%	72.40 \pm 0.38	-1.87%	0.29 \pm 0.15	96.51%	1.83 \pm 0.26	64.19%
				MMD	64.48 \pm 0.27	-2.61%	72.92 \pm 0.31	-1.17%	1.22 \pm 0.36	85.34%	2.11 \pm 0.49	58.71%
				HSIC	66.01 \pm 0.29	-0.30%	73.16 \pm 0.32	-0.84%	0.98 \pm 0.26	88.22%	1.00 \pm 0.28	80.43%
				PR	62.72 \pm 1.01	-5.27%	69.36 \pm 0.85	-5.99%	0.78 \pm 0.50	90.63%	1.07 \pm 0.36	79.06%
				CS	65.70 \pm 0.42	-0.77%	72.83 \pm 0.58	-1.29%	0.17 \pm 0.08	97.96%	0.75 \pm 0.22	85.32%
1035	1036	1037	Race	MLP	66.38 \pm 0.42	—	73.69 \pm 0.63	—	9.28 \pm 1.63	—	6.21 \pm 1.63	—
				DP	64.96 \pm 0.23	-2.14%	71.86 \pm 0.23	-2.48%	0.82 \pm 0.33	91.16%	1.30 \pm 0.26	79.07%
				MMD	65.71 \pm 0.65	-1.01%	70.57 \pm 0.52	-4.23%	3.97 \pm 0.97	57.22%	1.55 \pm 0.79	75.04%
				HSIC	65.81 \pm 0.24	-0.86%	72.92 \pm 0.23	-1.04%	1.75 \pm 0.31	81.14%	0.43 \pm 0.23	93.08%
				PR	64.25 \pm 0.87	-3.21%	70.25 \pm 0.30	-4.67%	1.56 \pm 0.87	83.19%	1.21 \pm 0.74	80.52%
				CS	65.16 \pm 0.45	-1.84%	72.56 \pm 0.72	-1.41%	0.55 \pm 0.19	94.07%	1.38 \pm 0.46	77.78%
1040	1041	1042	Gender	RN	78.14 \pm 0.47	—	86.58 \pm 0.55	—	51.66 \pm 0.97	—	35.67 \pm 1.11	—
				DP	62.42 \pm 4.79	-20.12%	66.86 \pm 3.19	-22.78%	0.46 \pm 0.25	99.11%	4.84 \pm 2.37	86.43%
				MMD	62.54 \pm 4.26	-19.96%	66.47 \pm 3.85	-23.23%	1.39 \pm 0.64	97.31%	5.89 \pm 3.12	83.49%
				HSIC	63.39 \pm 3.63	-18.88%	69.33 \pm 3.25	-19.92%	2.24 \pm 0.36	95.66%	3.83 \pm 2.22	89.26%
				PR	65.51 \pm 3.52	-16.16%	71.70 \pm 2.88	-17.19%	4.00 \pm 0.52	92.26%	5.05 \pm 2.57	85.84%
				CS	65.05 \pm 3.80	-16.75%	71.42 \pm 2.46	-17.51%	0.98 \pm 0.62	98.10%	1.53 \pm 1.05	95.71%
1045	1046	1047	Young	RN	78.14 \pm 0.47	—	86.67 \pm 0.53	—	41.74 \pm 1.17	—	18.35 \pm 1.56	—
				DP	66.78 \pm 3.61	-14.54%	73.95 \pm 3.44	-14.68%	2.43 \pm 0.83	94.18%	0.91 \pm 1.77	95.04%
				MMD	65.82 \pm 4.87	-15.77%	72.84 \pm 3.61	-15.96%	3.49 \pm 0.83	91.64%	1.60 \pm 0.71	91.28%
				HSIC	66.04 \pm 3.01	-15.49%	73.08 \pm 2.69	-15.68%	1.99 \pm 0.55	95.23%	1.04 \pm 0.60	94.33%
				PR	62.98 \pm 4.69	-19.40%	69.63 \pm 4.02	-19.66%	1.32 \pm 0.49	96.84%	1.82 \pm 0.53	90.08%
				CS	65.33 \pm 4.26	-16.39%	73.15 \pm 3.84	-15.60%	1.28 \pm 0.40	96.93%	0.30 \pm 0.12	98.37%

Table 3: The fairness performance on the tabular dataset for existing fair models, and we consider race and gender as sensitive attributes. A higher accuracy metric indicates better performance. \uparrow represents the accuracy improvement compared to MLP. A lower fairness metric indicates better fairness. \downarrow represents the improvement of fairness compared to MLP. **Green** values denotes better than MLP on the corresponding metric (ACC/AUC \uparrow ; $\Delta_{DP}/\Delta_{EO}\downarrow$), while **red** values denote worse. The results are based on 10 runs for all methods.

C MORE EXPERIMENTAL RESULTS

C.1 EXPERIMENTS ON IMAGE DATASET

In this section, we present the experimental results on the CelebA-A image dataset. The CelebA-A face attributes dataset (Liu et al., 2015) contains over 200,000 face images, where each image has 40 human-labeled attributes. Among the attributes, we select ‘Attractive’ as a binary classification task and consider ‘Gender’ and ‘Young’ as sensitive attributes. The results are presented in Table 3. The results show a similar finding with the tabular dataset, demonstrating that 1) DP method always achieves a lower Δ_{DP} but a relatively high Δ_{EO} . 2) HSIC is a more promising fair model to achieve equal opportunity.

1080
1081 C.2 IS THE REPRESENTATION LEARNED BY APPLYING CS VIEWED AS FAIR?
1082
10831084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096 Figure 6: T-SNE visualizations of the latent representations on Adult, COMPAS, and ACS-I,
1097 colored by the target attribute (top) and the sensitive attribute (bottom).1098
1099 To further validate that CS can learn fair representations, we visualize the T-SNE embeddings of
1100 the latent space from the last layer before the prediction layer (Van der Maaten & Hinton, 2008)³.
1101 Figure 6 displays the representations learned from the last embedding layer on the Adult, COMPAS,
1102 and ACS-I datasets, while Figure 12 presents the results for ACS-T and CelebA-A. Based on
1103 these visualizations, we make the following observations:1104
1105 **Obs. 7: The CS can learn representations that are indistinguishable between sensitive groups.**
1106 This observation validates the effectiveness of CS in learning fair representations. Specifically, the
1107 plots in the first row of Figure 6 illustrate the embedding visualization of two sensitive groups: blue
1108 for $S=0$ and red for $S=1$. Overall, the points are uniformly dispersed, with no clear clusters of
1109 nodes sharing the same color. This indicates that the embeddings are learned independently of the
1110 sensitive attribute. Although some groups have a greater number of data points—such as in the
1111 Adult dataset with the sensitive attribute *race*, where the ratio of $S=0:S=1$ is 1:9.20, and in the
1112 COMPAS dataset with *gender*, where the ratio is 1:4.17 (as shown in Table 4)—the distribution of
1113 points in both colors remains even.1114
1115 **Obs. 8: The CS can learn distinguishable representations for different target attributes.**
1116 Observing the second row of Figure 6, we can identify a distinct pattern in the distribution of the
1117 blue and red points across different locations in the plot. Among these, the embedding for ACS-I
1118 exhibits the clearest pattern, followed by Adult. This observation is consistent with the utility
1119 results presented in Table 2, which show a decrease in accuracy and AUC as the degree of negativity
1120 increases, particularly evident in the \uparrow columns compared to the MLP. In contrast, COMPAS presents
1121 a greater challenge in ensuring utility while considering fairness, as indicated by the less distinct
1122 pattern in the learned embeddings, corroborated by the most significant utility drops in Table 2.1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133 ³<https://scikit-learn.org/stable/modules/generated/sklearn\protect\penalty\z@.manifold.TSNE.html>

1134
1135

C.3 MORE PREDICTION DISTRIBUTIONS OVER THE SENSITIVE GROUPS

1136
1137
1138
1139
1140
1141

As described in Section 5.3, the kernel density plots in this subsection visualize how the prediction distributions vary across sensitive groups, and how different fairness regularizers affect this alignment. In general, a **larger overlap** between the distributions of different sensitive groups indicates **better group fairness**. For each figure in this subsection, the top row shows the distributions of the raw predictions \hat{Y} for all target labels, grouped by a given sensitive attribute, while the bottom row focuses on the positive class.

1142
1143
1144
1145
1146

For example, in Figure 7: the first row plots the prediction densities for *Race*: the blue shaded area corresponds to $\hat{Y} | \text{Race} = \text{Black}$ and the red shaded area corresponds to $\hat{Y} | \text{Race} = \text{White}$; the second row then plots the prediction densities for the positive class, i.e., $\hat{Y} = 1$, again conditioned on the two race groups (blue for $\hat{Y} = 1 | \text{Race} = \text{Black}$ and red for $\hat{Y} = 1 | \text{Race} = \text{White}$).

1147
1148
1149
1150

The other figures are interpreted analogously for their respective sensitive attributes. Since the degree of overlap can sometimes be difficult to judge by eye, we also print the corresponding group-fairness metric in each subfigure: the first row reports Δ_{DP} (demographic parity gap) and the second row reports Δ_{EO} (equalized odds gap).

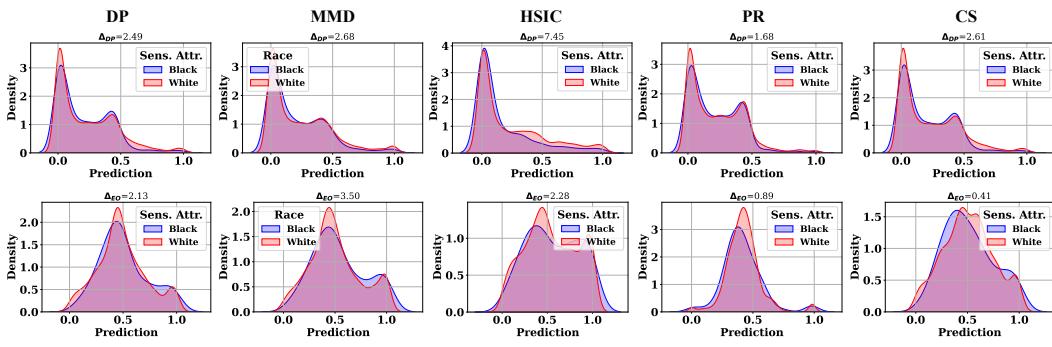
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162

Figure 7: Prediction distributions for black and white groups in the Adult dataset.

1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

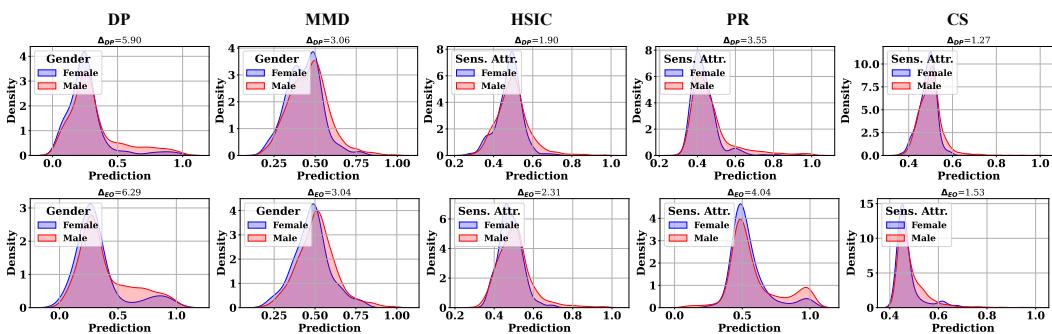
1199

1200

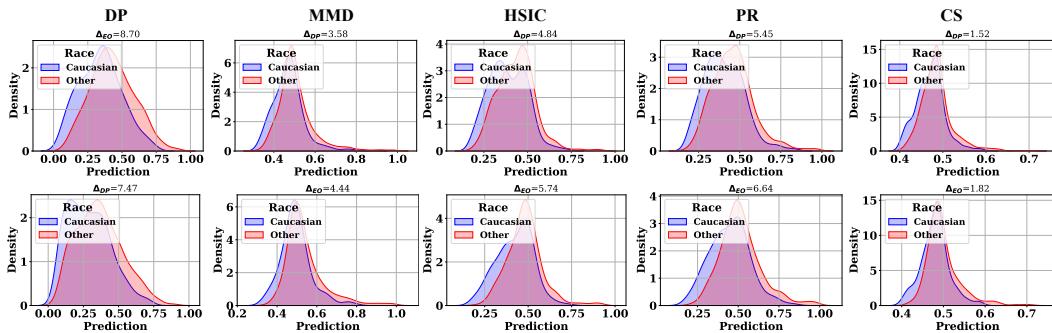
1201

1202

1203



(a) Prediction distributions for female and male groups in the COMPAS dataset.



(b) Prediction distributions for Caucasian and (all) other groups in the COMPAS dataset.

Figure 8: Prediction distribution over gender and race in the COMPAS dataset.

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

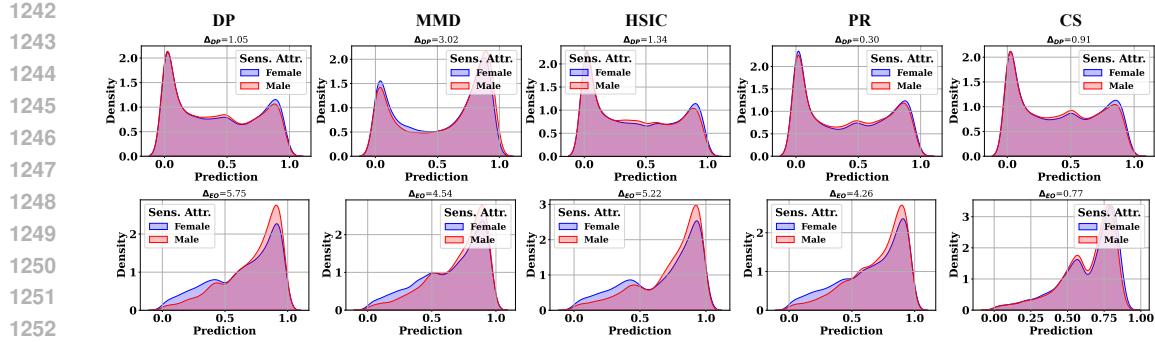
1237

1238

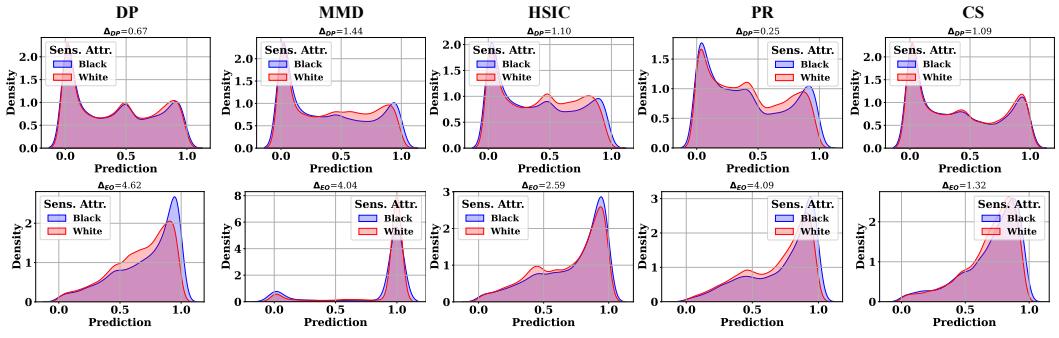
1239

1240

1241

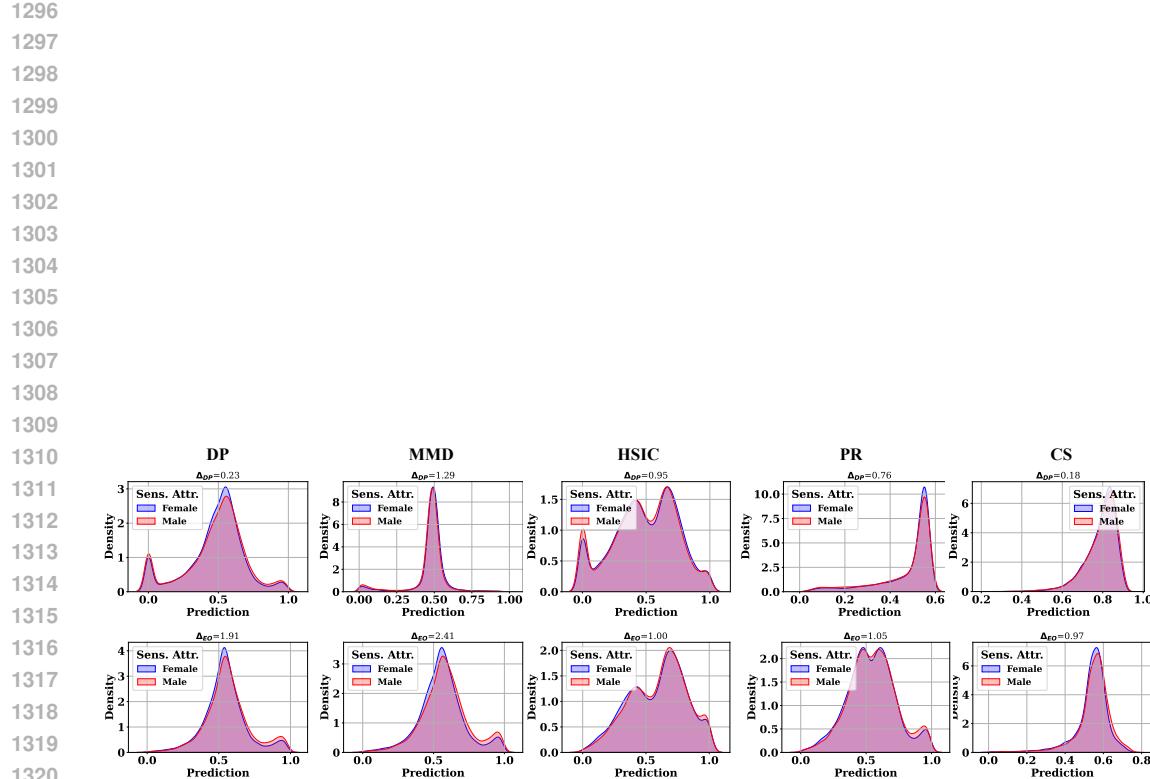


(a) Prediction distributions for female and male groups in the ACS-I dataset.



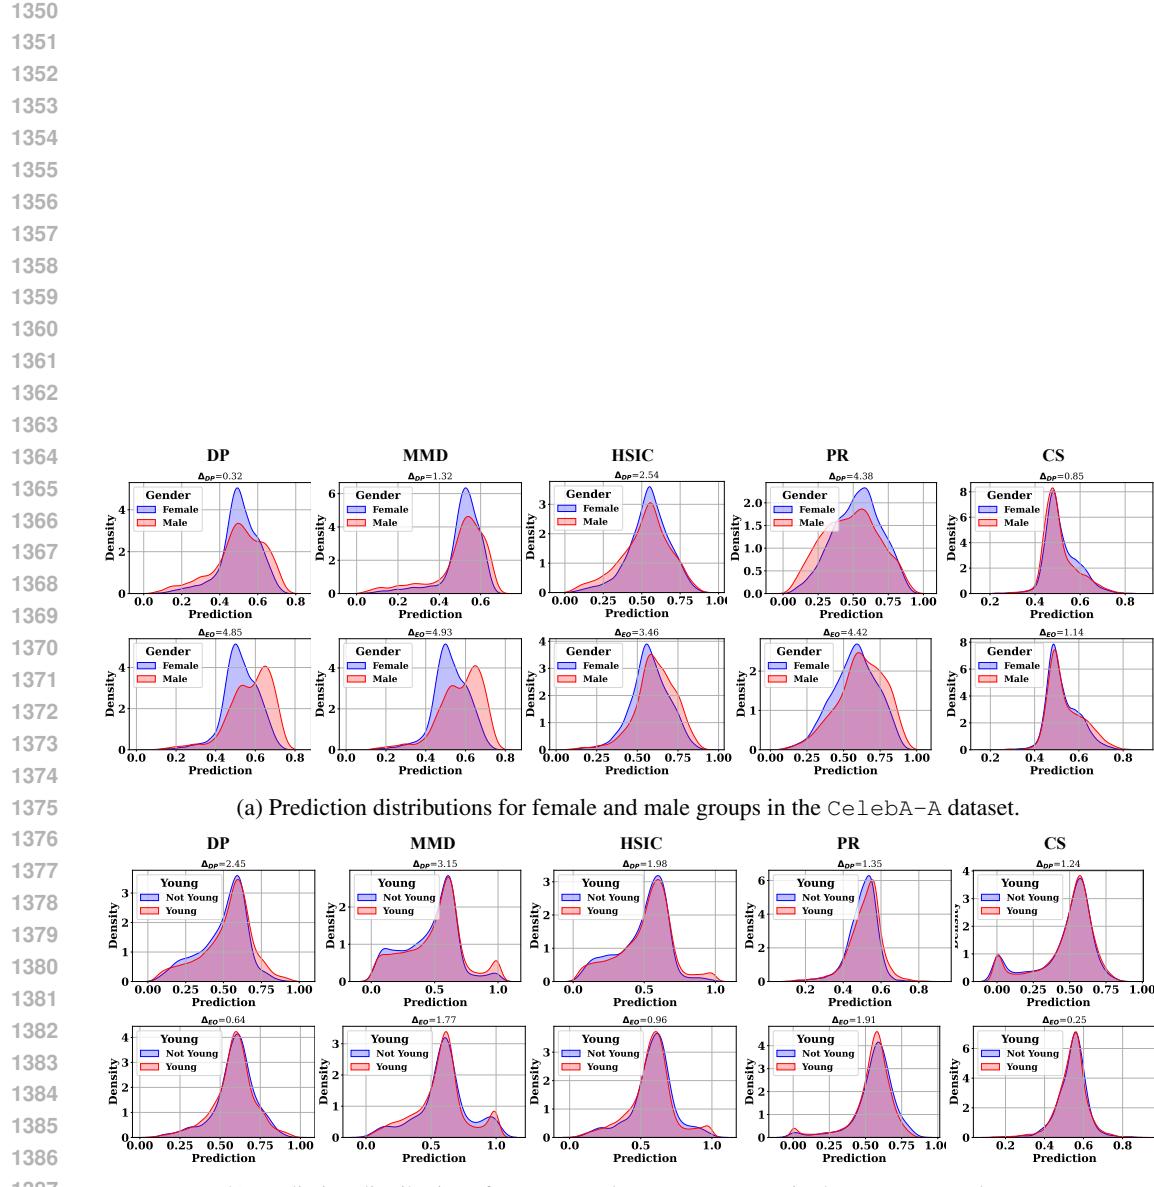
(b) Prediction distributions for black and white groups in the ACS-I dataset.

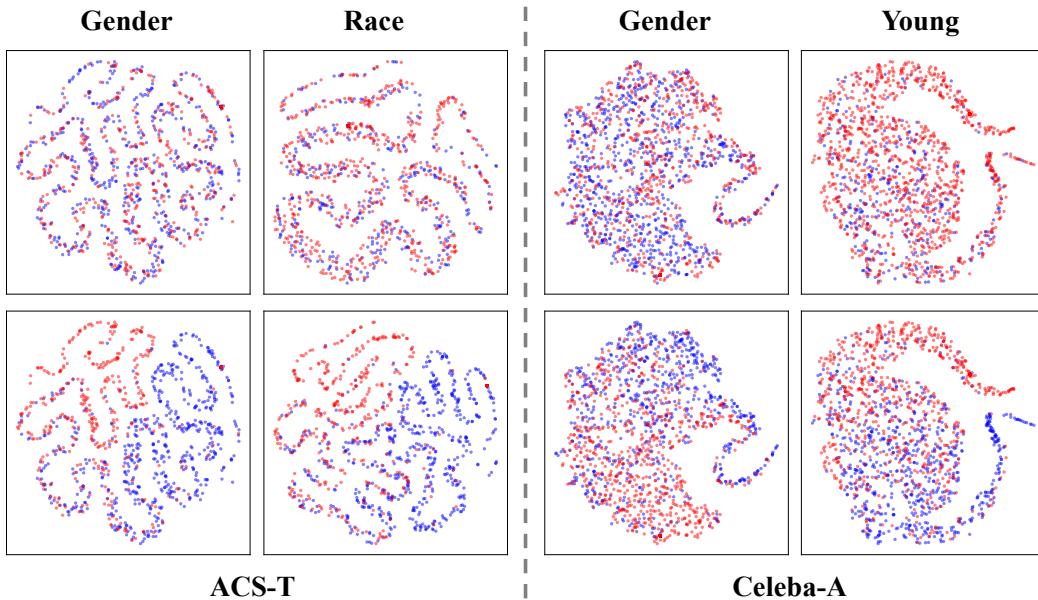
Figure 9: Accuracy and Δ_{DP} trade-off on ACS-I with sensitive attribute gender and race. Results located in the bottom-right corner are preferable.



1335 Figure 10: Accuracy and Δ_{DP} trade-off on ACS-T with sensitive attribute gender and race. Results
1336 located in the bottom-right corner are preferable.

1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349



1404
1405 C.4 MORE T-SNE PLOTS
14061427
1428 Figure 12: Accuracy and Δ_{DP} trade-off on ACS-T and CelebA-A. Results located in the bottom-right corner are preferable.
14291430 In addition to the T-SNE plots shown in Figure 6, which show the results on three datasets, we also
1431 include the T-SNE plots on the two remaining datasets ACS-T and CelebA-A in Figure 12.
14321433 D DATASET DESCRIPTIONS AND DETAILS
14341435 We conducted experiments on five datasets, including four tabular datasets and one image dataset.
1436 The introduction of these datasets is as follows:
1437

- **Adult**⁴ (Dua & Graff, 2017) The Adult dataset includes data from 45,222 individuals based on the 1994 US Census. The primary task is to predict whether an individual’s income exceeds \$50k USD, using various personal attributes. In this analysis, we focus on gender and race as sensitive attributes.
- **COMPAS**⁵ (Larson et al., 2016) The COMPAS dataset contains records of criminal defendants and is designed to predict the likelihood of recidivism within two years. It encompasses various attributes related to the defendants, including their criminal history, gender, and race.
- **ACS-I** and **ACS-T**⁶ (Ding et al., 2021) The ACS dataset is derived from the American Community Survey (ACS) Public Use Microdata Sample and encompasses several prediction tasks. These tasks include predicting whether an individual’s income exceeds \$50k and whether an individual is employed, with features such as race, gender, and other relevant characteristics tailored to each task.
- **CelebA-A**⁷ (Liu et al., 2015) The CelebFaces Attributes dataset comprises 20,000 face images of 10,000 distinct celebrities. Each image is annotated with 40 binary labels representing various facial attributes, including gender, hair color, and age. In this study, we focus on the ‘attractive’ label for a binary classification task, while considering ‘young’ and ‘gender’ as sensitive attributes.

1452 The detailed statistics for the aforementioned datasets are summarized as follows:
14531454
1455 ⁴<https://archive.ics.uci.edu/ml/datasets/adult>1456 ⁵<https://github.com/propublica/compas-analysis>1457 ⁶<https://github.com/zykls/folktables>1458 ⁷<https://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>

1458	Dataset	Task	Sen. Attr. (S)	#Samples	#Feat.	Class Y	0:1	1st S 0:1	2nd S 0:1
1459	Adult	Income	Gender, Race	45,222	101		1:0.33	1:2.08	1:9.20
1460	COMPAS	Credit	Gender, Race	6,172	405		1:0.83	1:4.17	1:0.52
1461	ACS-I	Income	Gender, Race	195,665	908		1:0.70	1:0.89	1:1.62
1462	ACS-T	Travel Time	Gender, Race	172,508	1,567		1:0.94	1:0.89	1:1.61
1463	CelebA-A	Attractive	Gender, Young	202,599	48 × 48		1:0.95	1:0.71	1:3.45

1464
 1465 Table 4: The table presents the statistics of the datasets. #Feat. refers to the total number of features
 1466 after preprocessing⁸. The ratio 0:1 represents the proportion between the two categories of the target
 1467 label or sensitive attributes.

1469 E BASELINES DETAILS

1471 We consider four widely used fairness methods: DP, MMD, HSIC, and PR. Specifically, DP and HSIC
 1472 minimize the demographic parity and Hilbert-Schmidt Independence Criterion, respectively. MMD
 1473 learns a classifier that optimizes the Mean Maximum Discrepancy. We also include base models MLP
 1474 and RN for tabular data and image data, respectively.

- 1476 • DP: It is a gap regularization method for demographic parity (Chuang & Mroueh, 2020). As these
 1477 fairness definitions cannot be optimized directly, gap regularization is differentiable and can be
 1478 optimized using gradient descent.
- 1479 • MMD: The Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) is a metric used to measure
 1480 the distance between probability distributions. Previous research has leveraged MMD to enhance
 1481 fairness in machine learning models, specifically in variational autoencoders (Louizos et al., 2016)
 1482 and MLPs (Deka & Sutherland, 2023). In this paper, we build on the methodologies from earlier
 1483 works (Zhao & Meng, 2015) to compute the MMD baseline.
- 1484 • HSIC: It minimizes the Hilbert-Schmidt Independence Criterion between the prediction accuracy
 1485 and the sensitive attributes (Gretton et al., 2005; Baharouei et al., 2020; Li et al., 2019).
- 1486 • Prejudice Remover (PR) (Kamishima et al., 2012) (Prejudice Remover) minimizes the prejudice
 1487 index, which is the mutual information between the prediction accuracy and the sensitive attributes.

1488 F MORE FAIRNESS DEFINITIONS

1490 G DETAILS OF THE GROUP FAIRNESS

1491 In this section, we provide the details of the group fairness. We first introduce the definition of group
 1492 fairness. Then, we introduce the existing group fairness metrics and algorithms.

- 1494 • DP (Demographic Parity or Statistical Parity) (Zemel et al., 2013). A classifier satisfies de-
 1495 mographic parity if the predicted outcome \hat{Y} is independent of the sensitive attribute S , i.e.,
 1496 $P(\hat{Y} | S = 0) = P(\hat{Y} | S = 1)$.
- 1498 • prule (Zafar et al., 2017). A classifier satisfies $p\%$ -rule if the ratio between the probability of
 1499 subjects having a certain sensitive attribute value assigned the positive decision outcome and the
 1500 probability of subjects not having that value also assigned the positive outcome should be no less
 1501 than $p/100$, i.e., $|P(\hat{Y} = 1 | S = 1)/P(\hat{Y} = 1 | S = 0)| \leq p/100$.
- 1502 • EOOpp (Equality of Opportunity) (Hardt et al., 2016). A classifier satisfies equalized opportunity
 1503 if the predicted outcome Y is independent of the sensitive attribute S when the label $Y = 1$, i.e.,
 1504 $P(\hat{Y} | S = 0, Y = 1) = P(\hat{Y} | S = 1, Y = 1)$.
- 1505 • EOOdd (Equalized Odds) (Hardt et al., 2016). A classifier satisfies equalized odds if the predicted
 1506 outcome Y is independent of the sensitive attribute S conditioned on the label Y , i.e., $P(\hat{Y} |$
 1507 $S = 0, Y = y) = P(\hat{Y} | S = 1, Y = y), y \in \{0, 1\}$.
- 1509 • ACC (Accuracy Parity). A classifier satisfies accuracy parity if the error rates of different sensitive
 1510 attribute values are the same, i.e., $P(\hat{Y} \neq Y | S = 0) = P(\hat{Y} \neq Y | S = 1), y \in \{0, 1\}$.
- 1511 • aucp (ROC AUC Parity). A classifier satisfies ROC AUC parity if its area under the receiver
 1512 operating characteristic curve with w.r.t. different sensitive attribute values is the same.

1512 • **ppv** (Predictive Parity Value Parity) A classifier satisfies predictive parity value parity if the
 1513 probability of a subject with a positive predictive value belonging to the positive class w.r.t.
 1514 different sensitive attribute values are the same, i.e., $P(Y = 1 | \hat{Y}, S = 0) = P(Y = 1 | \hat{Y}, S = 1)$.
 1515

1516 • **bnegc** (Balance for Negative Class). A classifier satisfies balance for the negative class if
 1517 the average predicted probability of a subject belonging to the negative class is the same w.r.t.
 1518 different sensitive attribute values, i.e., $\mathbb{E}[f(X) | Y = 0, S = 0] = \mathbb{E}[f(X) | Y = 0, S = 1]$.
 1519

1520 • **bposc** (Balance for Positive Class). A classifier satisfies balance for the positive class if the
 1521 average predicted probability of a subject belonging to the positive class is the same w.r.t.
 1522 different sensitive attribute values, i.e., $\mathbb{E}[f(X) | Y = 1, S = 0] = \mathbb{E}[f(X) | Y = 1, S = 1]$.
 1523

1524 • **abcc** (Area Between Cumulative density function Curves) (Han et al., 2023) is proposed to
 1525 precisely measure the violation of demographic parity at the distribution level. The new fairness
 1526 metrics directly measure the difference between the distributions of the prediction probability for
 1527 different demographic groups

H ADDITION EXPERIMENTS

H.1 ADDITION EXPERIMENTS ON MORE FAIRNESS METRICS

We provide additional results comparing our framework with baselines under the following fairness notions: Predictive Parity (PPV) (Chouldechova, 2017), p%-Rule (PRULE) (Zafar et al., 2017), Balance for Positive Class (BFP) (Kleinberg et al., 2016), and Balance for Negative Class (BFN) (Kleinberg et al., 2016). The dataset is Adult, using gender as the sensitive attribute. All other experimental settings are consistent with Table 1 in the paper.

Method	$\Delta_{PPV} (\downarrow)$	PRULE (\uparrow)	$\Delta_{BFP} (\downarrow)$	$\Delta_{BFN} (\downarrow)$
DP	27.35 ± 5.64	81.21 ± 9.04	11.25 ± 2.75	5.15 ± 0.44
MMD	35.19 ± 6.33	85.83 ± 7.15	18.32 ± 3.74	3.49 ± 0.25
HSIC	37.25 ± 3.19	96.18 ± 2.12	16.47 ± 1.21	4.04 ± 0.32
PR	25.46 ± 3.17	89.57 ± 7.39	21.45 ± 2.37	<u>3.46 ± 0.28</u>
CS	31.59 ± 4.35	97.75 ± 3.24	<u>15.25 ± 2.58</u>	3.18 ± 0.36

Table 5: Fairness performance comparison on the Adult dataset, with gender as the sensitive attribute under Δ_{PPV} , PRULE, Δ_{BFP} , and Δ_{BFN} .

We observe the following:

- CS generally achieves the best fairness trade-off performance across the four tested fairness notions.
- On the Adult, BFN is generally minimized more effectively than BFP.
- Since BFN is related to EO, the ranking of Δ_{BFN} aligns with Δ_{EO} in Table 1 of the paper. Note that, as stated in previous studies (Kleinberg et al., 2016), there is an inherent trade-off between BFP and BFN in practice.

H.2 ADDITIONAL EXPERIMENTS ON COMBINING MULTIPLE REGULARIZER TERMS SIMULTANEOUSLY

We conducted additional experiments where we combined both KL divergence and CS divergence as regularizers. The experiments were performed on the Adult, with gender as the sensitive attribute.

Actually, combining multiple fairness objectives has several drawbacks, which is why most existing studies avoid using multiple regularizers. Instead, they often choose to add simple constraint terms. The key drawbacks of combining fairness regularizers are summarized as follows:

⁸We adopt the preprocessing in previous studies (Le Quy et al., 2022; Mehrabi et al., 2021) involving identifying the target labels and sensitive attributes, and then selecting the relevant features for the analysis.

1566	Method	$\Delta_{DP} (\downarrow)$	$\Delta_{EO} (\downarrow)$
1567	DP	1.29 ± 0.95	20.15 ± 1.13
1568	CS	2.42 ± 0.85	2.27 ± 1.04
1569	KL	2.77 ± 0.86	10.42 ± 4.34
1570	CS+KL	2.46 ± 1.25	13.42 ± 6.12
1571	CS+0.5KL	2.25 ± 1.14	9.33 ± 6.36
1572			

Table 6: The fairness performance on Adult (gender).

- The CS divergence is upper-bounded by the KL divergence. Therefore, adding KL as an additional fairness objective is theoretically redundant and will not provide further benefits.
- Adding KL or other fairness metrics increases computational complexity, making the optimization process more challenging.

These experimental results further show the significance of our contribution: identifying a suitable, tighter-bounded fairness regularizer that balances effectiveness and computational efficiency.

H.3 ADDITIONAL EXPERIMENTAL ON EOoD AND EOoP

In this subsection, we provide additional experimental results on EOoD and EOoP, defined in Appendix G. Each line represents the results for one of the hyperparameter values $\alpha = 0.2/0.8/1.4$, denoted as 'lam' in the legend. We observe that both EOoD and EOoP regularizers perform well on the EO metrics but do not perform as well on the DP metrics. We summarize the results as follows:

Table 7: Fairness performance comparison on the ACS-I (gender) on additional metrics.

Method	$\Delta_{PPV} (\downarrow)$	$\Delta_{PRULE} (\uparrow)$	$\Delta_{BFP} (\downarrow)$	$\Delta_{BFN} (\downarrow)$	$\Delta_{DP} (\downarrow)$	$\Delta_{EO} (\downarrow)$
DP	11.86 ± 6.54	97.14 ± 12.41	4.83 ± 4.35	3.97 ± 4.02	0.96 ± 0.22	5.37 ± 0.32
EOoD	8.38 ± 1.28	84.75 ± 3.88	0.43 ± 0.63	0.31 ± 0.73	5.76 ± 1.42	0.64 ± 1.25
EOoP	8.00 ± 1.65	83.36 ± 4.41	0.34 ± 1.00	1.24 ± 1.43	6.32 ± 0.81	0.52 ± 1.29
CS	7.00 ± 8.35	96.90 ± 8.35	4.39 ± 1.32	2.77 ± 1.83	0.77 ± 0.38	0.90 ± 0.46

The observations from the results align with our claims, and the CS regularizer demonstrates significant effectiveness:

- As shown in Table 2 of our paper, the MLP achieves a Δ_{EO} of 2.13 ± 3.64 , whereas the DP regularizer gets a higher Δ_{EO} of 3.97 ± 4.02 . This indicates that the DP regularizer does not effectively optimize and may even negatively affects EO fairness.
- EO-based methods (EOoD and EOoP) show the worst performance in terms of DP fairness, even compared to other baselines such as MMD, HSIC, and PR (as reported in Table 2 in the paper). In particular, EOoP reaches approximately 15 in Δ_{DP} on the Adult dataset, as shown in the Appendix K. This high Δ_{DP} is consistently observed across different hyperparameter ($\alpha = 0.2, 0.8, 1.4$, referred to as 'lam' in the figure).

H.4 ADDITIONAL EXPERIMENTS ON PRE-PROCESSING AND POST-PROCESSING BASELINES

We have added a post-processing method, PostEO (Hardt et al., 2016) on the Adult dataset (with gender as the sensitive attribute).

The PostPro method is specifically designed to optimize for EO (Hardt et al., 2016), which explains its lower Δ_{EO} .

However, both pre-processing and post-processing methods share a common limitation: they result in lower utility (ACC or AUC). Considering the need for a balanced trade-off between fairness and utility, CS emerges as the most favorable option in our comparison.

To Revvier Pqrj:
W4

To Revvier LT9c:
W2

Method	ACC (\uparrow)	AUC (\uparrow)	Δ_{DP} (\downarrow)	Δ_{EO} (\downarrow)
DP	82.42 ± 0.39	86.91 ± 0.80	1.29 ± 0.95	20.15 ± 1.13
CS	83.31 ± 0.47	90.15 ± 0.49	2.42 ± 0.85	2.27 ± 1.04
PR	81.81 ± 0.52	85.38 ± 0.82	0.71 ± 0.40	12.45 ± 2.38
PostPro	80.25 ± 0.83	84.35 ± 0.98	5.75 ± 1.67	2.12 ± 1.44

Table 8: Comparison of methods on various metrics.

H.5 ADDITIONAL BASELINES WITH DEPENDENCE MEASURES

To further contextualize our CS-based fairness regularizer, we additionally compare against several classical dependence measures between the model prediction \hat{Y} and the sensitive attribute S : Hirschfeld–Gebelein–Rényi maximal correlation (HGR), mutual information (MI), and (distance) covariance (dCov) or empirical distance covariance. In all cases, we regularize the model by penalizing the corresponding dependence between \hat{Y} and S .

HGR maximal correlation (HGR). The Hirschfeld–Gebelein–Rényi maximal correlation (Hirschfeld, 1935; Gebelein, 1941; Rényi, 1959) between two random variables X and S is defined as:

$$\rho_{\text{HGR}}(X, S) = \sup_{f, g} \text{Corr}(f(X), g(S)) \text{ s.t. } \mathbb{E}[f(X)] = \mathbb{E}[g(S)] = 0, \mathbb{E}[f(X)^2] = \mathbb{E}[g(S)^2] = 1, \quad (38)$$

where the supremum is taken over square-integrable functions f and g . We instantiate an HGR-based fairness regularizer by penalizing $\rho_{\text{HGR}}(\hat{Y}, S)$ using a neural estimator.

Mutual information (MI). Mutual information (Cover, 1999) between X and S is

$$I(X; S) = \iint p_{X,S}(x, s) \log \frac{p_{X,S}(x, s)}{p_X(x) p_S(s)} dx ds, \quad (39)$$

or, for discrete variables,

$$I(X; S) = \sum_x \sum_s p_{X,S}(x, s) \log \frac{p_{X,S}(x, s)}{p_X(x) p_S(s)}. \quad (40)$$

Equivalently, MI can be written as a Kullback–Leibler divergence

$$I(X; S) = D_{\text{KL}}(p_{X,S} \parallel p_X p_S). \quad (41)$$

Our MI-based baseline regularizes the mutual information between the prediction and the sensitive attribute, $I(\hat{Y}; S)$, using a differentiable estimator.

Distance covariance (dCov) and empirical distance covariance. Distance covariance (Székely et al., 2007) between $X \in \mathbb{R}^p$ and $S \in \mathbb{R}^q$ is defined via their joint and marginal characteristic functions $\varphi_{X,S}$, φ_X , and φ_S as

$$\text{dCov}^2(X, S) = \int_{\mathbb{R}^{p+q}} |\varphi_{X,S}(t, u) - \varphi_X(t) \varphi_S(u)|^2 w(t, u) dt du, \quad (42)$$

for a suitable weight function $w(t, u)$. In practice, we use the standard empirical distance covariance estimator. Given samples $\{(x_i, s_i)\}_{i=1}^n$, define pairwise distances $a_{ij} = \|x_i - x_j\|$ and $b_{ij} = \|s_i - s_j\|$, and their double-centered versions:

$$A_{ij} = a_{ij} - \bar{a}_{i\cdot} - \bar{a}_{\cdot j} + \bar{a}_{\cdot\cdot}, \quad (43)$$

$$B_{ij} = b_{ij} - \bar{b}_{i\cdot} - \bar{b}_{\cdot j} + \bar{b}_{\cdot\cdot}, \quad (44)$$

where $\bar{a}_{i\cdot}$ and $\bar{a}_{\cdot j}$ denote row and column means, and $\bar{a}_{\cdot\cdot}$ is the grand mean (and analogously for b). The empirical distance covariance is then:

$$\widehat{\text{dCov}}^2(X, S) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n A_{ij} B_{ij}. \quad (45)$$

Our dCov-based regularizer penalizes $\widehat{\text{dCov}}^2(\hat{Y}, S)$.

1674
 1675 **Experimental setup.** To assess how our CS-based regularizer compares with other classical
 1676 dependence measures, we conduct an additional experiment on the ADULT dataset using *gender* as
 1677 the sensitive attribute. We keep **all settings identical** to the main experiment on ADULT in Section 5:
 1678 the same data split and preprocessing, the same MLP classifier architecture, optimizer, batch size,
 1679 learning rate, number of epochs, and the same protocol for tuning the fairness-regularization weight.
 1680 The only change is the choice of the fairness loss term $\mathcal{L}_{\text{fair}}$.

1681 Concretely, for each additional baseline, we replace the CS-based fairness loss in Equation (13) with
 1682 the corresponding dependence measure between the model prediction \hat{Y} and the sensitive attribute
 1683 S : (i) the **HGR** baseline penalizes the Hirschfeld–Gebelein–Rényi maximal correlation $\rho_{\text{HGR}}(\hat{Y}, S)$
 1684 using a neural estimator; (ii) the **MI** baseline penalizes the mutual information $I(\hat{Y}; S)$ between
 1685 prediction and sensitive attribute, estimated with a differentiable MI estimator; and (iii) the **dCov**
 1686 baseline penalizes the empirical distance covariance $\widehat{\text{dCov}}^2(\hat{Y}, S)$ defined before. In all cases, the
 1687 overall training objective retains the same form as Equation (1), and we evaluate the resulting models
 1688 on accuracy, AUC, Δ_{DP} , and Δ_{EO} using the same test split as in the main experiments. The results
 1689 are summarized in Table 9.

Method	ACC (\uparrow)	AUC (\uparrow)	$\Delta_{\text{DP}} (\downarrow)$	$\Delta_{\text{EO}} (\downarrow)$
HGR	80.13 ± 1.35	84.20 ± 1.27	3.82 ± 0.84	6.82 ± 3.77
dCov	82.31 ± 0.62	85.39 ± 0.89	4.75 ± 1.67	12.41 ± 1.44
PR	81.81 ± 0.52	85.38 ± 0.82	0.71 ± 0.40	12.45 ± 2.38
CS	83.31 ± 0.47	90.15 ± 0.49	2.42 ± 0.85	2.27 ± 1.04

1697 Table 9: Additional comparison of with HGR and dCov.

1698 From the results in Table 9, we observe that the proposed CS-based regularizer achieves the best
 1699 overall performance among all dependence-based baselines: it attains the highest ACC and AUC
 1700 while keeping both Δ_{DP} and Δ_{EO} low, confirming that CS offers a robust utility–fairness trade-off.
 1701 More specifically, (i) **HGR** is theoretically a very strong dependence measure, but in practice it
 1702 requires a neural estimator of the maximal correlation, which makes optimization noisy and sensitive
 1703 to hyperparameters; this is reflected in its relatively low ACC/AUC and larger standard deviations,
 1704 although its fairness metrics are still competitive, indicating that it can reduce dependence when the
 1705 optimization succeeds. (ii) **dCov** is a kernel-based or distance-based statistic with a closed-form
 1706 empirical estimator, so it is easier to optimize and leads to higher ACC/AUC and smaller variance
 1707 than HGR; however, its fairness performance is weaker, suggesting that penalizing average pairwise
 1708 distances between prediction and sensitive-feature embeddings is less aligned with group-rate gaps
 1709 than the density-ratio style CS divergence, which yields tighter control over the discrepancies that
 1710 drive Δ_{DP} and Δ_{EO} . (iii) **PR** (an MI-based regularizer) achieves very small Δ_{DP} , consistent
 1711 with its design of directly reducing mutual information between \hat{Y} and S and thereby aligning the
 1712 marginal prediction rates across groups, but its Δ_{EO} remains large and its utility is moderate, as
 1713 MI does not explicitly constrain the conditional error rates $P(\hat{Y} \mid Y, S)$ that underlie equalized
 1714 odds. Overall, these observations support CS as the most balanced choice among the considered
 1715 dependence measures.

1716 **Discussion: Relation between MI and our KL- and PR-based regularizers.** For completeness,
 1717 we briefly discuss how the generic MI baseline above relates to the KL-based fairness losses and the
 1718 Prejudice Remover (PR) baseline used in our main experiments. By definition, mutual information is
 1719 a Kullback–Leibler divergence between the joint distribution and the product of the marginals:

$$1720 \quad I(X; S) = D_{\text{KL}}(p_{X,S} \parallel p_X p_S). \quad (46)$$

1721 Thus, MI and KL belong to the same family of information-theoretic discrepancy measures: MI
 1722 uses KL to quantify *any* deviation of $p_{X,S}$ from independence, while many fairness regularizers
 1723 based on KL (including the KL term used in our fairness-loss landscape in Figure 2) penalize
 1724 Kullback–Leibler divergences between *conditional* distributions, such as $D_{\text{KL}}(p_{\hat{Y}|S=0} \parallel p_{\hat{Y}|S=1})$.
 1725 These conditional KL penalties are closely related to MI but not identical: if all $p_{\hat{Y}|S=s}$ coincide,
 1726 then both the conditional KL and $I(\hat{Y}; S)$ vanish, yet vanishing conditional KL for one pair of groups
 1727 does not necessarily minimize the full KL between $p_{\hat{Y},S}$ and $p_{\hat{Y}} p_S$.

1728 The Prejudice Remover (PR) (Kamishima et al., 2012) can be viewed as an explicit MI-based
 1729 regularizer. PR minimizes the *prejudice index*, which is defined as the mutual information between
 1730 the prediction and the sensitive attribute, $I(\hat{Y}; S)$, under a log-linear model of the conditional odds.
 1731 In this sense, PR instantiates the generic MI regularization principle with a particular parametric form
 1732 and optimization scheme.

1733 In summary, the generic MI baseline, the KL-based fairness penalties, and PR all enforce indepen-
 1734 dence between \hat{Y} and S using KL divergence in different guises: MI regularizes the KL divergence
 1735 between $p_{\hat{Y}, S}$ and $p_{\hat{Y}} p_S$; Our KL-based fairness loss penalizes KL divergences between group-
 1736 conditional predictions distributions; and PR minimizes a parametric approximation of $I(\hat{Y}; S)$. Our
 1737 CS-based regularizer complements this family by replacing the KL-based dependence measure with
 1738 the Cauchy–Schwarz divergence, which enjoys closed-form kernel estimators and the tighter bounds
 1739 analyzed in the main text.

1741 H.6 ADVERSARIAL METHODS EXPERIMENTS

1743 We conducted additional experiments using Adversarial Debiasing (Louppe et al., 2017), which we
 1744 refer to as ADV below.

1746 Method	1747 ACC (\uparrow)	1748 AUC (\uparrow)	1749 Δ_{DP} (\downarrow)	1750 Δ_{EO} (\downarrow)
1747 DP	1748 82.42 ± 0.39	1749 86.91 ± 0.80	1750 1.29 ± 0.95	1751 20.15 ± 1.13
1747 CS	1748 83.04 ± 0.51	1749 90.84 ± 0.35	1750 2.13 ± 0.89	1751 2.35 ± 1.15
1747 ADV	1748 81.58 ± 1.26	1749 83.08 ± 0.75	1750 16.3 ± 7.5	1751 14.2 ± 8.6

1752 Table 10: Additional experiment on the fairness performance of ADV on the Adult dataset (gender
 1753 attribute).

1754 From Table 10, we observe that:

- 1756 • The ADV method exhibits lower utility (in terms of ACC and AUC) and higher Δ_{DP} compared to
 1757 both the DP and CS fairness regularizers. It also performs worse than the CS regularizer in terms
 1758 of Δ_{EO} .
- 1759 • ADV also shows a higher variance in accuracy, likely due to the greater difficulty of optimizing
 1760 adversarial objectives compared to the DP and CS regularization approaches.

1762 I MORE EXPERIMENTAL DETAILS

1764 In this section, we describe the details of the experimental setup. In this work, we adopted a
 1765 straightforward stopping strategy. We employ a linear decay strategy for the learning rate, halving
 1766 it every 50 training step. The model training is stopped when the learning rate decreases to a value
 1767 below $1e^{-5}$. Across all datasets, we use a weight decay of 0.0, StepLR with a step size of 50 and
 1768 a gamma value of 0.1, and train for 150 epochs using the Adam Optimizer (Kingma & Ba, 2014).
 1769 The batch size and learning rate vary depending on the dataset, with specific values provided below.
 1770 Additionally, Table 11 lists the range of the control hyperparameter β for each fairness approach. The
 1771 experiments were executed using NVIDIA RTX A4000 GPUs with 16GB GDDR6 Memory.

1773 I.1 HYPERPARAMETER SETTINGS

1775 1. Training Hyperparameters:

- 1776 • Tabular data (Adult, COMPAS, ACS–I, and ACS–T):
 - 1777 – Learning rate: $1e^{-2}$
 - 1778 – Weight decay: 0.0
 - 1779 – StepLR_step: 50
 - 1780 – StepLR_gamma: 0.1
 - 1781 – Training epochs: 150

To Revvier
ESYX: R2-1

To Revvier Pqrj:
W4

To Revvier LT9c:
W2

To Revvier JuMp:
w2

To Revvier LT9c:
W2

1782 – Batch sizes: 1,024 on Adult, 32 on COMPAS, 4,096 on ACS-I, 4,096 on ACS-T
 1783 • Image data (CelebA-A):
 1784 – Learning rate: $1e^{-3}$
 1785 – Weight decay: 0.0
 1786 – StepLR_step: 50
 1787 – StepLR_gamma: 0.1
 1788 – Training epochs: 150
 1789 – Batch sizes: 256.
 1790

2. Architecture Hyperparameters:

1791 • Multilayer perceptron:
 1792 – Number of layers: 3
 1793 – Number of hidden neurons: {512, 256, 64}
 1794 • ResNet-18 (He et al., 2016):
 1795 – Model: <https://github.com/pytorch/vision/blob/main/torchvision/models/resnet.py>
 1796

I.2 HYPERPARAMETER SELECTION

1802 To implement CS and the baseline methods, we adjust the hyperparameter β by tuning it within a
 1803 specified range. The details of the hyperparameter selection process and the specific range for β are
 1804 provided below:

Method	Fairness Control Hyperparameter β
DP	0.5, 1.0, 1.2, 1.4, 1.6, 1.8, 2.0, 2.5, 3.0, 3.5, 4
HSIC	0.1, 1, 5, 10, 50, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1, 000
PR	0.05, 0.2, 0.3, 0.40, 0.50, 0.7, 0.9, 1.0
ADV	0.5, 1.0, 1.2, 1.4, 1.6, 1.8, 2.0, 2.5, 3.0, 3.5
CS	$1e^{-6}$, $1e^{-5}$, $1e^{-4}$, $1e^{-3}$, $1e^{-2}$, $2e^{-2}$, $5e^{-2}$, 0.1, 0.5, 1.0, 2.0, 3.0, 4.0, 50, 150

1813 Table 11: The selections of fairness control hyperparameter blue β .
 1814

J RELATED WORK

1819 In this section, we first review relevant prior studies, beginning with an overview of algorithmic
 1820 fairness in machine learning. We then narrow our focus to regularization-based in-processing methods,
 1821 which are central to our approach.

J.1 ALGORITHMIC FAIRNESS IN MACHINE LEARNING

1825 The importance of fairness in machine learning has grown significantly as the demand for unbiased
 1826 decision-making models for individuals and groups increases. This is especially critical in high-
 1827 stakes applications where the consequences of biased decisions can be severe. Fairness is commonly
 1828 categorized into three main types: *Individual fairness* (Yurochkin et al., 2019; Mukherjee et al.,
 1829 2020; Yurochkin & Sun, 2020; Kang et al., 2020; Mukherjee et al., 2022), which aims to ensure that
 1830 similar individuals are treated similarly; *Group fairness* (Hardt et al., 2016; Verma & Rubin, 2018;
 1831 Li et al., 2020; Ling et al., 2023), which focuses on achieving fairness across predefined subgroups,
 1832 often defined by sensitive attributes such as gender or race; *Counterfactual fairness* (Kusner et al.,
 1833 2017; Agarwal et al., 2021; Zuo et al., 2022), which seeks to ensure fairness by considering how
 1834 decisions would hold under alternative scenarios. Given the widespread adoption of group fairness
 1835 metrics in real-world applications and the increasing development of in-processing techniques for
 neural networks, particularly for tabular and image data.

1836 Various techniques for mitigating bias in machine learning models can be categorized into three
 1837 main approaches: *pre-processing*, *in-processing*, and *post-processing*. *Pre-processing* methods
 1838 focus on addressing biases present in the dataset itself to ensure that the trained model exhibits
 1839 fairness (Kamiran & Calders, 2012; Calmon et al., 2017a). For instance, these techniques may
 1840 involve rebalancing the dataset or modifying the data collection process (Calmon et al., 2017b).
 1841 *In-processing* methods, on the other hand, adjust the training objectives by incorporating fairness
 1842 constraints directly into the learning process (Kamishima et al., 2012; Zhang et al., 2018; Madras
 1843 et al., 2018; Zhang et al., 2022; Buyl & De Bie, 2022; Alghamdi et al., 2022; Shui et al., 2022;
 1844 Mehrrotra & Vishnoi, 2022). This approach aims to ensure that the model learns fair representations
 1845 during training. Finally, *post-processing* methods modify the predictions made by classifiers after
 1846 the model has been trained, with the goal of promoting fairness across different groups (Hardt et al.,
 1847 2016; Jiang et al., 2020; Tsaousis & Alghamdi, 2022). By categorizing these techniques, we can
 1848 better understand the different strategies available for mitigating bias in machine learning systems.
 1849

1850 J.2 REGULARIZATION-BASED IN-PROCESSING METHODS

1851 In this paper, we explore three types of regularization-based in-processing methods. First, *Gap*
 1852 *Regularization* (Chuang & Mroueh, 2020) streamlines the optimization process by offering a smooth
 1853 approximation of real-world loss functions, which are typically non-convex and difficult to optimize
 1854 directly. This category includes methods such as DP, EO, and EOD. Second, the *Independence*
 1855 approach integrates fairness constraints into the optimization, aiming to mitigate the influence of
 1856 protected attributes on model predictions while maintaining overall performance. Notable examples
 1857 of this approach include PR (Kamishima et al., 2012) and HSIC (Li et al., 2019). Lastly, *adversarial*
 1858 *debiasing* seeks to minimize utility loss while hindering an adversary’s ability to accurately predict
 1859 the protected attributes. This approach encompasses methods like ADV (Zhang et al., 2018; Louppe
 1860 et al., 2017; Beutel et al., 2017; Edwards & Storkey, 2015; Adel et al., 2019) and LAFTR (Madras
 1861 et al., 2018).

1862 J.3 CAUCHY-SCHWARZ DIVERGENCE IN OTHER ML SETTINGS

1864 Cauchy-Schwarz (CS) divergence has also been studied and applied in several machine-learning
 1865 problems outside of algorithmic fairness. Early work used CS divergence together with Parzen
 1866 window density estimates to build information-theoretic criteria for clustering and graph-based
 1867 learning, and to relate CS divergence to Mercer kernels and graph cuts (Jenssen et al., 2006). More
 1868 recent studies have employed CS divergence as a training objective for representation learning and
 1869 deep models, for example in information-bottleneck formulations for regression (Yu et al., 2024),
 1870 domain adaptation (Yin et al., 2024), and CS-regularized autoencoders that improve density estimation
 1871 and clustering performance (Tran et al., 2022). Conditional variants of CS divergence have further
 1872 been developed for time-series analysis and sequential decision making (Yu et al., 2025). These
 1873 works demonstrate that CS divergence is a versatile discrepancy measure for density estimation
 1874 and representation learning; our contribution is complementary, as we systematically develop and
 1875 evaluate CS divergence as an *in-processing fairness regularizer* with dedicated theoretical analysis
 1876 and extensive experiments in the algorithmic-fairness setting.

1876 In contrast to these applications, which primarily target density estimation, clustering, or represen-
 1877 tation learning objectives, our focus is on *algorithmic fairness*. To the best of our knowledge, our
 1878 work is the first to systematically develop Cauchy-Schwarz divergence as an *in-processing fairness*
 1879 *regularizer*, with theoretical analysis tailored to group-fairness notions and an extensive empirical
 1880 study of the resulting utility-fairness trade-offs.

To Revvier JuMp:
W3

To Revvier Pqrj:
W1

To Revvier JuMp:
W3

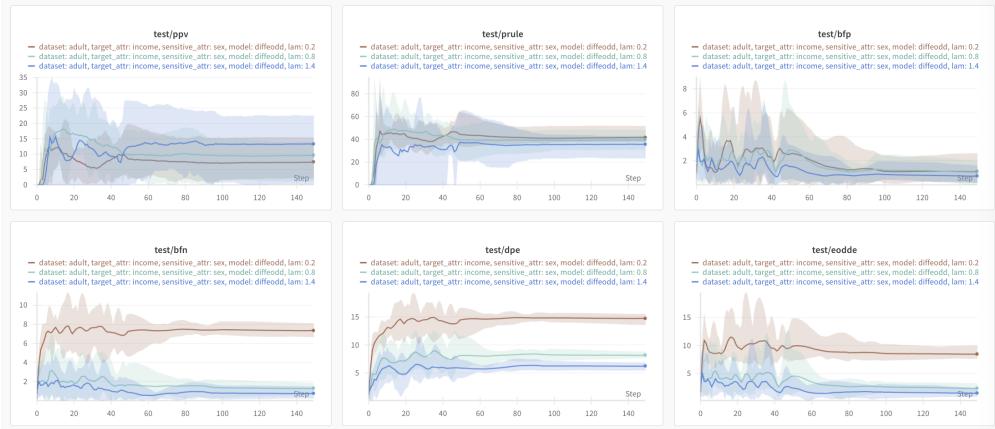
1890 K CURVES

1891 In this section, we show some important curves we recorded.

1892 *Note: The line represents the mean values, and the shaded area indicates the variation across all*
 1893 *runs.*

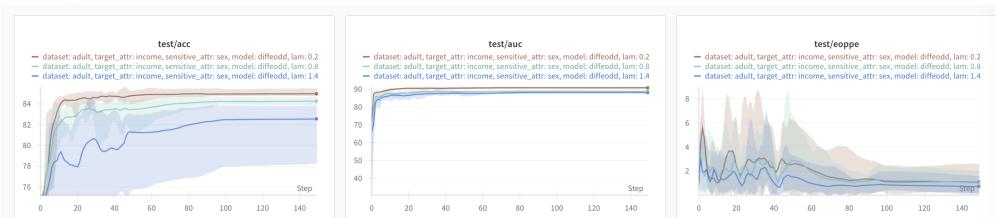
1894 K.1 EODD ADDITIONAL RESULTS

1895 [Adult (Gender)]



1914 Figure 13: Other metrics (Δ_{PPV} (\downarrow), $PRULE$ (\uparrow), Δ_{BFP} (\downarrow), Δ_{BFN} (\downarrow)), Δ_{DP} (Shown as
 1915 'dpe'), and eodde.

1916 Adult (Gender)



1926 Figure 14: ACC, AUC and Δ_{EO} (Shown as "eoppe")

1944

ACS-I (Gender)

1945

1946

1947

1948

1949

1950

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

1962

1963

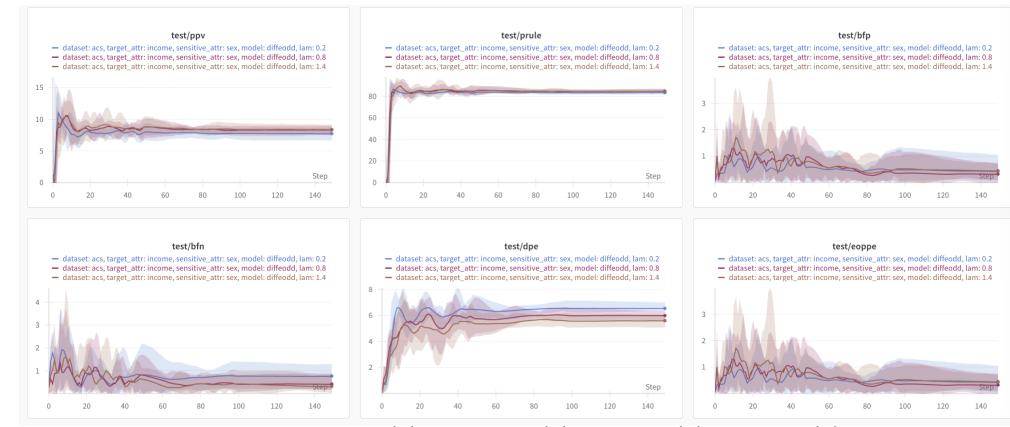


Figure 15: Other metrics (Δ_{PPV} (\downarrow), $PRULE$ (\uparrow), Δ_{BFP} (\downarrow), Δ_{BFN} (\downarrow), Δ_{DP} (Shown as 'dpe'), and eodde.

1964

1965

1966

1967

1968

1969

1970

1971

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997

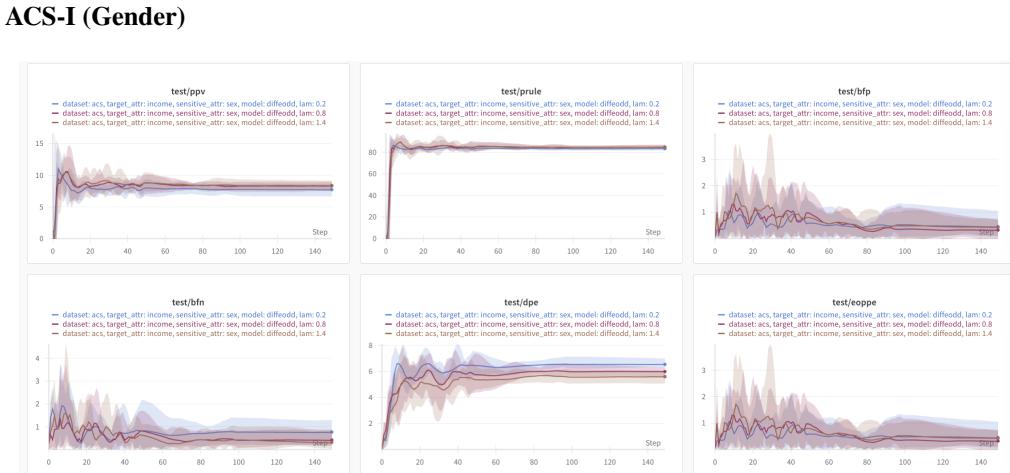


Figure 16: ACC, AUC and Δ_{EO} (Shown as “eoppe”).

1998
1999

K.2 EOPP ADDITIONAL RESULTS

2000

Adult (Gender)

2001

2002

2003

2004

2005

2006

2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

2026

2027

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045

2046

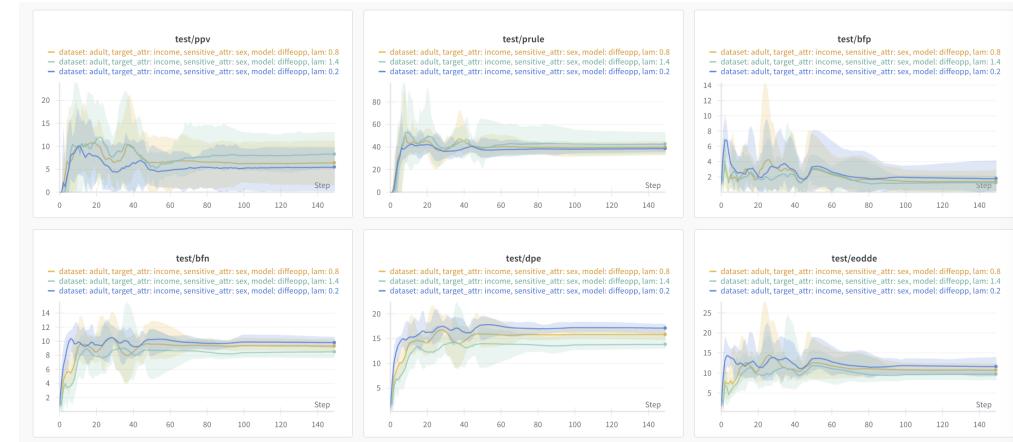
2047

2048

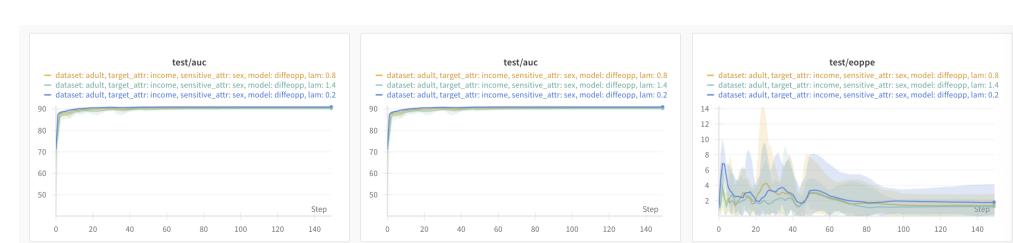
2049

2050

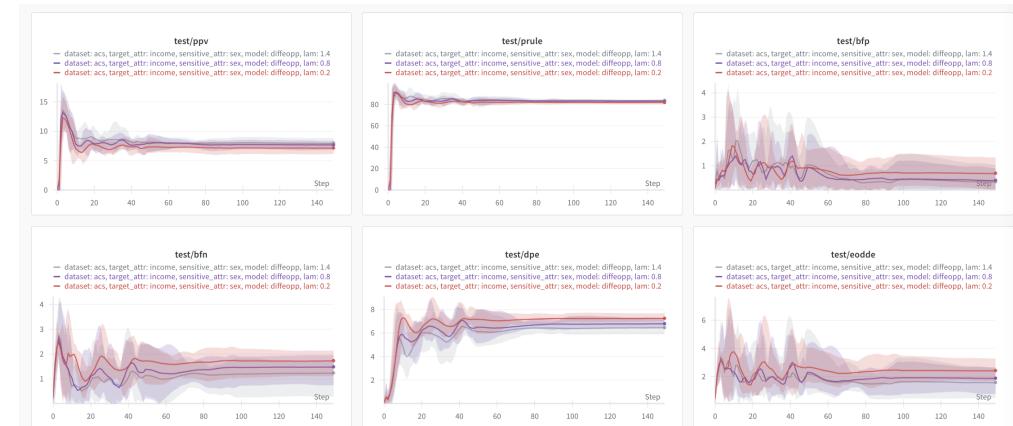
2051

Figure 17: Other metrics (\triangle_{PPV} (\downarrow), $PRULE$ (\uparrow), \triangle_{BFP} (\downarrow), \triangle_{BFN} (\downarrow), \triangle_{DP} (Shown as 'dpe'), and eodde).

Adult (Gender)

Figure 18: ACC, AUC and \triangle_{EO} (Shown as “eoppe”).

ACS-I (Gender)

Figure 19: Other metrics (\triangle_{PPV} (\downarrow), $PRULE$ (\uparrow), \triangle_{BFP} (\downarrow), \triangle_{BFN} (\downarrow), \triangle_{DP} (Shown as 'dpe'), and eodde).

2052 **ACS-I (Gender)**

2053

2054

2055

2056

2057

2058

2059

2060

2061

2062

2063

2064

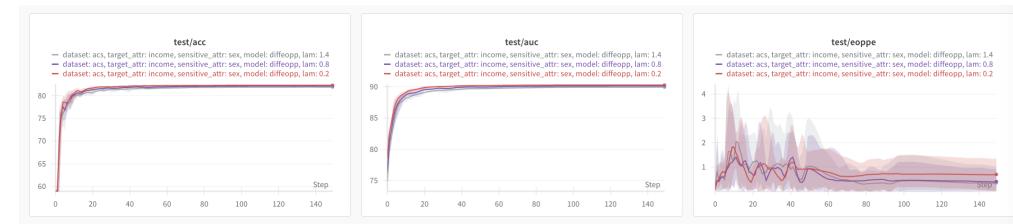
2065

2066

2067

2068

2069

Figure 20: ACC, AUC and Δ_{EO} (Shown as “eoppe”).

- **Dataset:** Adult

- **Sensitive Attribute:** Gender (represented as “Sex”)

- **Hyperparameters:** Learning rate = 1×10^{-2} , $\alpha = 5 \times 10^{-2}$, batch size = 1024, and $\beta = 1.0$ (kept consistent across all regularizers).

2070 **Training Fairness loss: HSIC, CS, and DP**

2071

2072

2073

2074

2075

2076

2077

2078

2079

2080

2081

2082

2083

2084

2085

2086

2087

2088

2089 **Training Fairness Loss: HSIC and CS**

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105

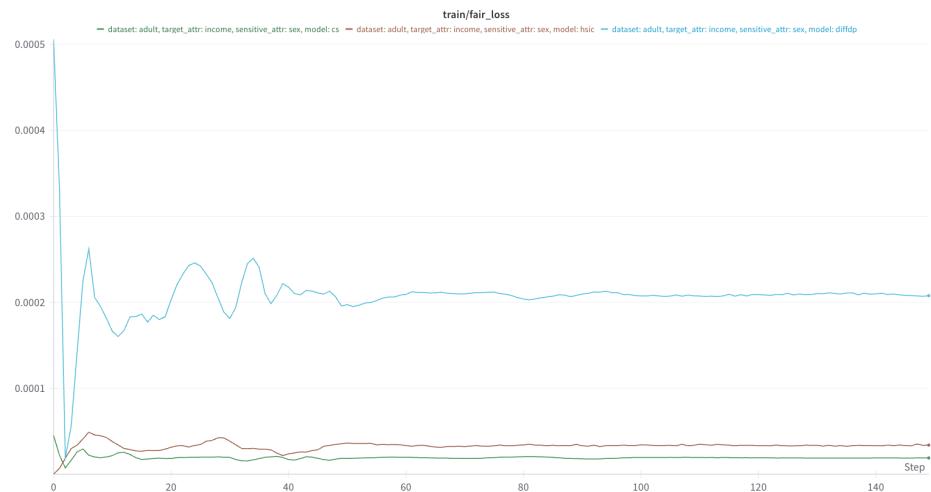


Figure 21: Training loss.

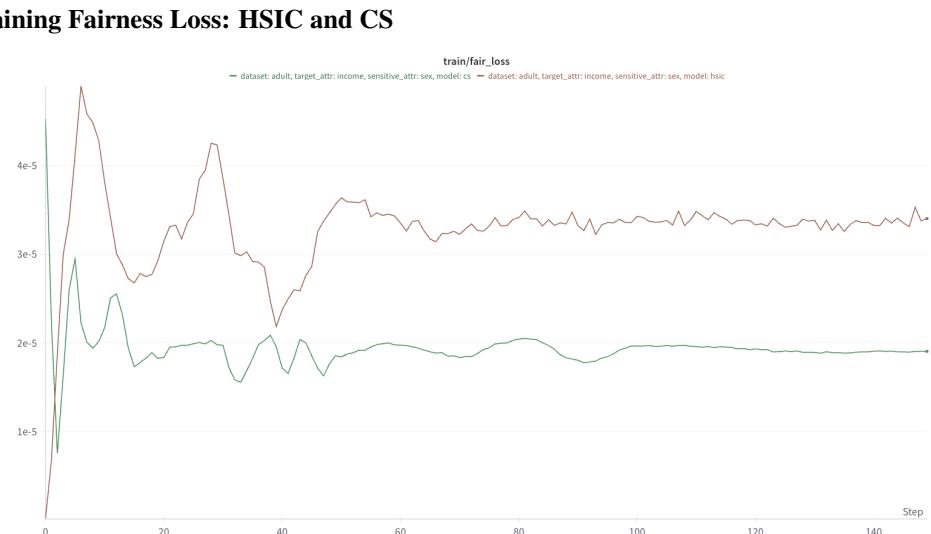
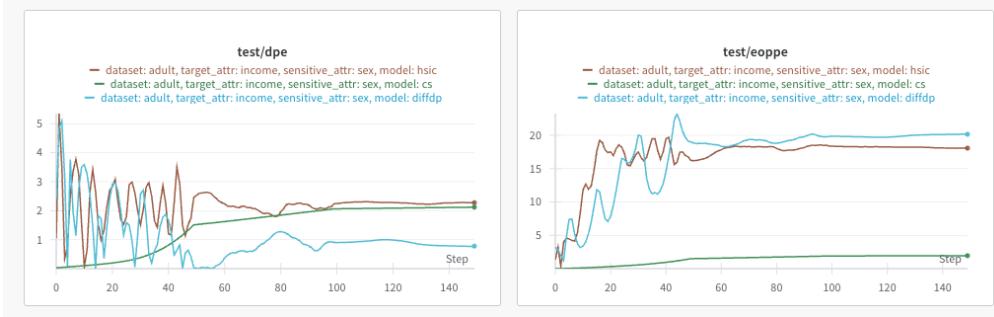


Figure 22: Training loss (excluding DP regularizer line) to more clearly observe the gap between HSIC and CS regularizers.

2106 **Test Δ_{DP} and Δ_{EO} :**

Figure 23: Test Δ_{DP} and Δ_{EO} .

L ADDITIONAL ABLATIONS AND EXPERIMENTS WITH MULTIPLE SENSITIVE ATTRIBUTES

To Rewvier Pqrj:
Follow up

L.1 ABLATION OVER KERNEL FAMILIES (GAUSSIAN/ LAPLACIAN/POLYNOMIAL)

Beyond the bandwidth study, we also compare different *kernel families* used inside the CS regularizer. We consider three standard choices:

- **Gaussian (RBF):** $k_{\text{rbf}}(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\|\mathbf{u}-\mathbf{v}\|_2^2}{2\sigma^2}\right)$.
- **Laplacian:** $k_{\text{lap}}(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\|\mathbf{u}-\mathbf{v}\|_2}{\sigma}\right)$.
- **Polynomial (degree 2):** $k_{\text{poly}}(\mathbf{u}, \mathbf{v}) = (\gamma \mathbf{u}^\top \mathbf{v} + 1)^2$, $\gamma = 1/\sigma$.

On the Adult–Income task (sensitive attribute Sex), we keep the model architecture, optimizer, training schedule, and regularization weight λ identical to the main experiment, and only change the kernel family used in the CS loss. For a fair comparison, we use the same bandwidth $\sigma_x = \sigma_y = \sigma_{\text{cross}} = 1$ for all three kernels. We summarize the final results from Figure 26 into Table 12.

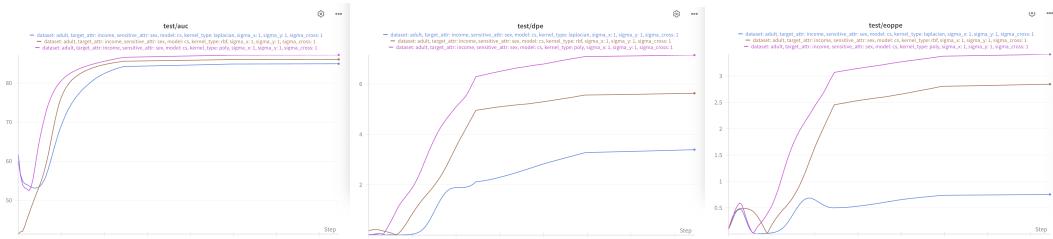


Figure 24: Ablation study: Accuracy and fairness of different kernel functions (Gaussian/Laplacian/polynomial) on Adult (sex).

Kernel type	AUC (%) \uparrow	Acc (%) \uparrow	$\Delta_{EO} \downarrow$	$\Delta_{DP} \downarrow$
Laplacian	84.2	83.5	0.75	3.3
Gaussian RBF	85.8	84.7	2.5	5.6
Polynomial	86.3	85.1	3.3	6.5

Table 12: Ablation study: CS performance on different kernel functions (Gaussian/ Laplacian/polynomial)

From Table 12, all three kernels achieve strong predictive performance (AUC \approx 84–86%), but they trace out different points on the accuracy–fairness trade-off curve:

2160 • The **Laplacian kernel** yields the **smallest EO and DP gaps**, i.e., the best group fairness, at the
 2161 cost of a small drop in AUC/accuracy.

2162 • The **polynomial kernel** attains slightly **higher AUC/accuracy**, but with noticeably larger EO/DP
 2163 gaps.

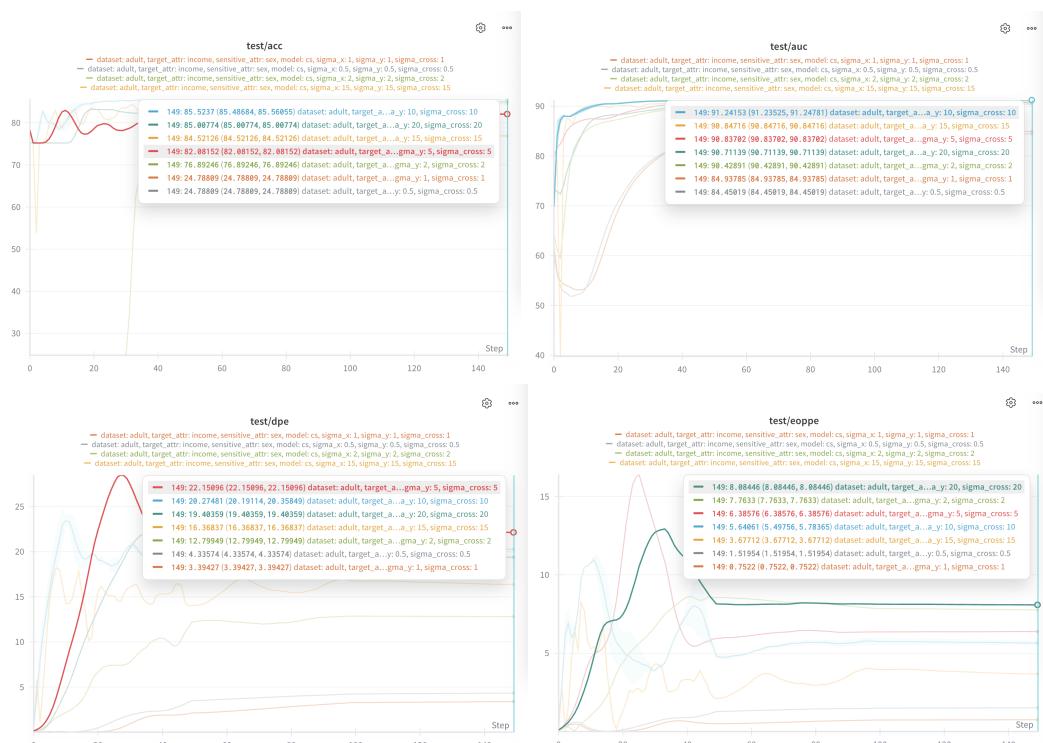
2164 • The **Gaussian (RBF) kernel** lies between the two, offering a **balanced trade-off**: it preserves
 2165 most of the performance benefits of the polynomial kernel while significantly improving fairness
 2166 compared to polynomial and remaining closer to Laplacian.

2168 These results confirm that CS is compatible with multiple kernel families and that the choice of kernel
 2169 can be used to tune the fairness–utility trade-off. In the main paper, we adopt the Gaussian kernel as
 2170 a default because it provides a **stable, middle-ground trade-off** and is widely used in dependence
 2171 measures (MMD/HSIC), making our comparison to existing divergences more direct.

2173 L.2 ABLATION OVER KERNEL BANDWIDTH σ

2174 To study the effect of the kernel bandwidth in the CS regularizer, we fix the model and training setup
 2175 used on Adult–Income (sensitive attribute Sex) and vary the kernel bandwidths while keeping all
 2176 other hyperparameters fixed (same optimizer, learning rate, batch size, and α as in the main Adult
 2177 experiments).

2178 We use an RBF kernel $k_\sigma(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\|\mathbf{u}-\mathbf{v}\|_2^2}{2\sigma^2}\right)$, and set $\sigma_x = 1$ for the prediction output, while
 2179 varying $\sigma_x = \sigma_y = \sigma_{\text{cross}} \in \{0.5, 1, 2, 5, 10, 15, 20\}$ for the sensitive attribute and cross terms in the
 2180 CS loss. For each configuration we record: (i) the Δ_{EO} ('test/eoppe'), (ii) the Δ_{DP} ('test/dpe'), (iii)
 2181 test accuracy, and (iv) test AUC at the final epoch.



2207 Figure 25: Ablation study: Ablation over kernel bandwidth of CS on Adult (sex). **How to read**
 2208 **these figures: in each subfigure, the lines are ordered in the *legend* from top to bottom according**
 2209 **to their values, from highest to lowest.**

2211 From Table 13, we see two observations:

2212 • **Extremely small bandwidths** ($\sigma \leq 1$) Here the RBF kernel becomes extremely peaked. The CS
 2213 loss forces almost pointwise independence, which makes optimization unstable: accuracy collapses

$\sigma_y = \sigma_x = \sigma_{\text{cross}}$	$\Delta_{\text{EO}} \downarrow$	$\Delta_{\text{DP}} \downarrow$	ACC (%) \uparrow	AUC (%) \uparrow
0.5	1.52	4.34	24.8	84.5
1	0.75	3.39	24.8	84.9
2	7.76	12.80	76.9	90.4
5	6.39	22.15	82.1	90.8
10	5.64	20.27	85.5	91.2
15	3.68	16.37	84.5	90.8
20	8.08	19.40	85.0	90.7

Table 13: Ablation study: Kernel bandwidth on the proposed CS fairness regularizer.

to $\approx 25\%$ and AUC drops to $\approx 84\text{--}85\%$. The very small DP/EO gaps in this regime are therefore misleading—they correspond to a nearly random classifier.

- **Moderate bandwidths** ($\sigma \in [2, 20]$) In this regime, the classifier maintains **high utility** (AUC $\approx 90\text{--}91\%$, accuracy between 77% and 85.5%). Fairness varies smoothly with σ :
- Δ_{EO} generally improves when moving from too-local kernels ($\sigma = 2, 5$) to more moderate ones; $\sigma = 15$ achieves the smallest Δ_{EO} among the stable runs.
- Δ_{DP} is best at $\sigma = 2$, but this comes with noticeably lower accuracy. For $\sigma \in \{10, 15, 20\}$, both DP and EO are within a similar, reasonable range while utility is highest.

Overall, the results show that CS is **not hypersensitive** to the exact kernel bandwidth: once σ is chosen in a reasonable range, the method achieves consistently high AUC with a stable fairness–utility trade-off. In our main experiments we therefore use a moderate bandwidth (e.g., $\sigma = 10$) that lies in this stable region, balancing strong accuracy (AUC $\approx 91\%$) with substantially reduced $\Delta_{\text{DP}}/\Delta_{\text{EO}}$.

L.3 EXPERIMENTS WITH MULTIPLE SENSITIVE ATTRIBUTES

To evaluate CS on multi-attribute fairness, we extend the Adult setting from a single sensitive attribute to an **intersectional attribute** combining sex and race. We construct four groups $S \in \{0, 1, 2, 3\}$ as White-Male, White-Female, Non-White-Male, and Non-White-Female. For each group g we define: $\text{DP}_g = \mathbb{P}(\hat{Y} = 1 \mid S = g)$, $\text{EO}_g = \mathbb{P}(\hat{Y} = 1 \mid Y = 1, S = g)$, $\text{Acc}_g = \mathbb{P}(\hat{Y} = Y \mid S = g)$. We then report the **intersectional demographic-parity gap** $\Delta\text{DP}_{\text{inter}} = \max_g \text{DP}_g - \min_g \text{DP}_g$, the **intersectional equal-opportunity gap** $\Delta\text{EO}_{\text{inter}} = \max_g \text{EO}_g - \min_g \text{EO}_g$, and the **worst-group accuracy** $\text{Acc}_{\min} = \min_g \text{Acc}_g$. The code snippet of these metrics are in Appendix N. Using the same MLP architecture and α ('lam' in the figure) = 0.5, we compare CS with three representative dependence-based regularizers (diffDP, HSIC, diffEOpp). The results are summarized below:

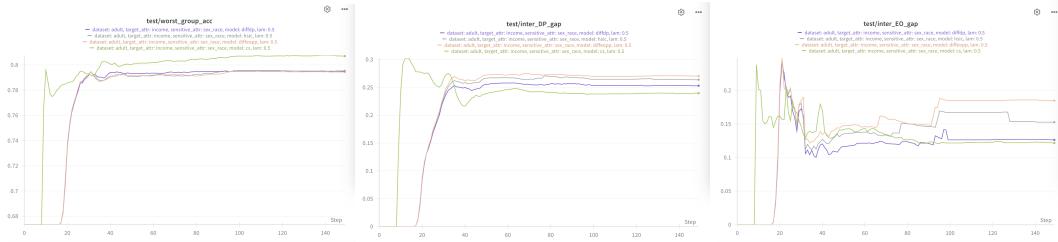


Figure 26: Evaluation of CS on the Adult (sex) as the sensitive attribute.

Method	$\Delta\text{DP}_{\text{inter}} \downarrow$	$\Delta\text{EO}_{\text{inter}} \downarrow$	$\text{Acc}_{\min} \uparrow$
diffDP	0.255	0.112	0.792
HSIC	0.262	0.152	0.791
diffEOpp	0.268	0.183	0.790
CS (ours)	0.241	0.108	0.803

Table 14: Evaluation of CS on the Adult (sex) as the sensitive attribute.

From Table 14, we observe that CS attains the **smallest** intersectional $\Delta\text{DP}^{\text{inter}}/\Delta\text{EO}^{\text{inter}}$ while also achieving the **highest** worst-group accuracy, indicating that the proposed Cauchy–Schwarz regularizer remains effective even when fairness is evaluated over four intersectional subgroups rather than a single sensitive attribute.

In particular, CS improves both $\Delta\text{DP}^{\text{inter}}$ and $\Delta\text{EO}^{\text{inter}}$ over HSIC and DP regularizer, and slightly improves Acc_{\min} compared to all baselines. These results suggest that the tighter dependence control provided by CS translates into more balanced treatment across intersectional groups without sacrificing worst-case utility.

M DISCUSSION: WHEN AND WHY IS CS EXPECTED TO OUTPERFORM OTHER DIVERGENCES?

To address the question of when CS is practically preferable to other dependence measures, we summarize the regimes where CS is theoretically and empirically advantageous:

- **Heavy-tailed or skewed group distributions.** When one group exhibits heavier prediction tails (e.g., more extreme probabilities), KL and DP penalties can be dominated by these tails. In contrast, the CS divergence, through its L_2 -normalization of group embeddings, limits the influence of such extremes and yields a more stable fairness penalty.
- **Scale-mismatched representations.** When latent embeddings for different groups differ markedly in variance or norm (a common scenario in deep models), Euclidean-based MMD can report large distances even when the group embeddings are well aligned in direction. CS compares *normalized* embeddings and therefore provides a tighter and more meaningful notion of “closeness” for fairness.
- **Imbalanced group sizes.** In highly imbalanced datasets, group-conditional densities are estimated with very different effective sample sizes. In such cases, KL and HSIC can fluctuate considerably with the minority group’s empirical variance, whereas the cosine-style normalization implicit in CS makes the fairness loss less sensitive to this sampling noise.

These regimes are not hypothetical: the datasets in Section 5 (Adult, COMPAS, ACS, CelebA-A) all exhibit at least one of these characteristics. This helps explain why CS often achieves lower Δ_{DP} and Δ_{EO} at comparable or better utility in our experiments.

N CODE SNIPPET

```

1 def calculate_intersectional_metrics(y_pred, y_target, sensitive):
2     if isinstance(y_pred, torch.Tensor):
3         y_pred = y_pred.detach().cpu().numpy()
4     if isinstance(y_target, torch.Tensor):
5         y_target = y_target.detach().cpu().numpy()
6     if isinstance(sensitive, torch.Tensor):
7         sensitive = sensitive.detach().cpu().numpy()
8
9     y_pred = y_pred.flatten()
10    y_target = y_target.flatten()
11    sensitive = sensitive.flatten()
12
13    y_pred_binary = (y_pred > 0.5).astype(int)
14
15    groups = [0, 1, 2, 3]
16
17    rates = {}
18    for g in groups:
19        mask = sensitive == g
20        if mask.sum() > 0:
21            rates[g] = y_pred_binary[mask].mean()
22        else:
23            rates[g] = 0.0

```

```

2322    24
2323    25     dp_gap = max(rates.values()) - min(rates.values())
2324    26
2325    27     tprs = {}
2326    28     for g in groups:
2327    29         mask = (sensitive == g) & (y_target == 1)
2328    30         if mask.sum() > 0:
2329    31             tprs[g] = y_pred_binary[mask].mean()
2330    32         else:
2331    33             tprs[g] = 0.0
2332    34
2333    35     eo_gap = max(tprs.values()) - min(tprs.values())
2334    36
2335    37     accs = {}
2336    38     for g in groups:
2337    39         mask = sensitive == g
2338    40         if mask.sum() > 0:
2339    41             accs[g] = (y_pred_binary[mask] == y_target[mask]).mean()
2340    42         else:
2341    43             accs[g] = 0.0
2342    44     worst_group_acc = min(accs.values())
2343    45
2344    46     return {
2345    47         "intersectional_DP_gap": dp_gap,
2346    48         "intersectional_EO_gap": eo_gap,
2347    49         "worst_group_acc": worst_group_acc,
2348    50     }

```

Listing 1: Calculation of Intersectional Fairness Metrics

```

2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375

```