EQUAL IMPROVABILITY: A NEW FAIRNESS NOTION CONSIDERING THE LONG-TERM IMPACT

Anonymous authors
Paper under double-blind review

ABSTRACT

Devising a fair classifier that does not discriminate against different groups is an important problem in machine learning. Although researchers have proposed various ways of defining group fairness, most of them only focused on the *immediate* fairness, ignoring the *long-term* impact of a fair classifier under the dynamic scenario where each individual can improve its feature over time. Such dynamic scenarios happen in real world, *e.g.*, college admission and credit loaning, where each rejected sample makes effort to change its features to get accepted afterwards. In this dynamic setting, the long-term fairness should equalize the samples' feature distribution across different groups after the rejected samples make some effort to improve. In order to promote long-term fairness, we propose a new fairness notion called *Equal Improvability* (EI), which equalizes the potential acceptance rate of the rejected samples across different groups assuming a bounded level of effort will be spent by each rejected sample. We analyze the properties of EI and its connections with existing fairness notions. To find a classifier that satisfies the EI requirement, we propose and study three different approaches that solve EI-regularized optimization problems. Through experiments on both synthetic and real datasets, we demonstrate that the proposed EI-regularized algorithms encourage us to find a fair classifier in terms of EI. Finally, we provide experimental results on dynamic scenarios which highlight the advantages of our EI metric in achieving the long-term fairness. Codes are available in anonymous GitHub repository [1].

1 INTRODUCTION

Over the past decade, machine learning has been used in a wide variety of applications. However, these machine learning approaches are observed to be unfair to individuals having different ethnicity, race, and gender. As the implicit bias in artificial intelligence tools raised concerns over potential discrimination and equity issues, various researchers suggested defining fairness notions and developing classifiers that achieve fairness. One popular fairness notion is *demographic parity* (DP), which requires the decision-making system to provide output such that the groups are equally likely to be assigned to the desired prediction classes, *e.g.*, acceptance in the admission procedure. DP and related fairness notions are largely employed to mitigate the bias in many realistic problems such as recruitment, credit lending, and university admissions (Zafar et al., 2017b; Hardt et al., 2016; Dwork et al., 2012; Zafar et al., 2017a).

However, most of the existing fairness notions are *zero-order*, *i.e.*, they only focus on immediate fairness, without taking potential follow-up inequity risk into consideration. In Fig. 1, we provide an example scenario when using a conventional zero-order fairness notion (DP in this example) has a long-term fairness issue, in a simple loan approval problem setting. Consider two groups (group 0 and group 1) with different distributions, where each individual has one label (approve the loan or not) and two features (credit score, income) that can be improved over time. Suppose each group consists of two clusters (with three samples each), and the distance between the clusters is different for two groups. Fig. 1 visualizes the distributions of two groups and the decision boundary of a classifier f which achieves DP among the groups. We observe that the rejected samples (left-hand-side of the decision boundary) in group 1 are located further away from the decision boundary than the rejected samples in group 0. As a result, the rejected applicants in group 1 need to make more effort to cross the decision boundary and get approval. This improvability gap between the two groups can

[1] https://anonymous.4open.science/r/ei_fairness-23CA/

make the rejected applicants in group 1 less motivated to improve their features, which may increase the gap between different groups in the future. Therefore, a new fairness notion that considers the improvability of the rejected samples is necessary to achieve the equality of true qualification status over time.

In this paper, we introduce a novel *first-order* fairness notion dubbed as *Equal Improvability* (EI) to address the problem of inequity of improvability. Let $\mathbf{x}$ be the feature of a sample and f be a score-based classifier, *e.g.*, predicting a sample as accepted if $f(\mathbf{x}) \geq 0.5$ holds and as rejected otherwise. We assume each rejected individual wants to get accepted in the future, thus improving its feature within a certain effort budget towards the direction that maximizes its score $f(\mathbf{x})$. Under this setting, we define EI fairness as the equity of the potential acceptance rate of the different rejected groups, once each individual makes the best effort within the predefined budget. This prevents the risk of exacerbating the gap between different groups in the long run.

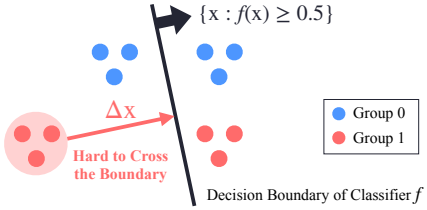

Figure 1: **Toy example showing the insufficiency of zero-order fairness notion.** We consider the binary classification (accept/reject) on 12 samples (dots), where $\mathbf{x}$ is the feature of the sample and the color of the dot represents the group. The given classifier f is fair in terms of a popular zero-order notion called demographic parity (DP), but does not have equal improvability of rejected samples ($f(\mathbf{x}) < 0.5$) in two groups; the rejected samples in group 1 needs more effort $\Delta\mathbf{x}$ to be accepted, *i.e.*, $f(\mathbf{x}+\Delta\mathbf{x}) \geq 0.5$, compared with the rejected samples in group 0.

Our key contributions are as follows:

- We propose a new group fairness notion called *Equal Improvability* (EI), which aims to equalize the probability of rejected samples being qualified after a certain amount of feature improvement, for different groups. EI encourages rejected individuals in different groups to have an equal amount of motivation to improve their feature to get accepted in the future. We analyze the properties of EI and the connections of EI with other existing fairness notions.
- We provide three methods to find a classifier that is fair in terms of EI, each of which uses a unique way of measuring the inequity in the improvability. Each method is solving a min-max problem where the inner maximization problem is finding the best effort to measure the EI unfairness, and the outer minimization problem is finding the classifier that has the smallest fairness-regularized loss. Experiments on synthetic/real datasets demonstrate that our algorithms find classifiers having low EI unfairness.
- We run experiments on dynamic scenarios where the data and the classifier evolve over multiple rounds, and show that training a classifier with EI constraints is beneficial for making the feature distributions of different groups identical in the long run, *i.e.*, achieving the long-term fairness.

2 Equal Improvability

Before defining our new fairness notion called *Equal Improvability* (EI), we first introduce necessary notations. For an integer n, let $[n] = \{0, \ldots, n-1\}$. We consider a binary classification setting where each data sample has an input feature vector $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$ and a label $\mathrm{y} \in \mathcal{Y} = \{0, 1\}$. In particular, we have a sensitive attribute $\mathrm{z} \in \mathcal{Z} = [Z]$, where Z is the number of sensitive groups. As suggested by Chen et al. (2021), we sort d features $\mathbf{x} \in \mathbb{R}^d$ into three categories: improvable features $\mathbf{x}_{\mathrm{I}} \in \mathbb{R}^{d_{\mathrm{I}}}$, mutable features $\mathbf{x}_{\mathrm{M}} \in \mathbb{R}^{d_{\mathrm{M}}}$, and immutable features $\mathbf{x}_{\mathrm{IM}} \in \mathbb{R}^{d_{\mathrm{IM}}}$, where $d_{\mathrm{I}} + d_{\mathrm{M}} + d_{\mathrm{IM}} = d$ holds. Here, improvable features $\mathbf{x}_{\mathrm{I}}$ refer to the features that can be improved and can directly affect the outcome, *e.g.*, salary in the credit lending problem, and GPA in the school's admission problem. In contrast, mutable features $\mathbf{x}_{\mathrm{M}}$ can be altered, but are not directly related to the outcome, *e.g.*, marital status in the admission problem, communication type in the credit lending problem. Although the individuals may manipulate these mutable features to get the desired outcome, we do not consider it as a way to make efforts as it does not affect the individual's true qualification status. Immutable features $\mathbf{x}_{\mathrm{IM}}$ are features that cannot be altered, such as race, age, or date of birth. Note that if sensitive attribute z is included in the feature vector, then it belongs to immutable features. For ease of notation, we write $\mathbf{x} = (\mathbf{x}_{\mathrm{I}}, \mathbf{x}_{\mathrm{M}}, \mathbf{x}_{\mathrm{IM}})$. Let $\mathcal{F} = \{f : \mathcal{X} \to [0, 1]\}$ be the set of classifiers, where each classifier is parameterized by $\boldsymbol{w}$, *i.e.*, $f = f_{\boldsymbol{w}}$. Given $f \in \mathcal{F}$, we consider the

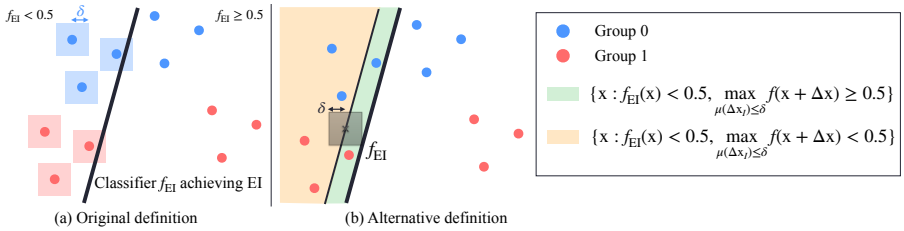

Figure 2: **Visualization of EI fairness.** For binary classification on 12 samples (dots) in two groups (red/blue), we visualize fairness notion defined in this paper: (a) shows the original definition in Def. 2.1, and (b) shows an alternative definition in Prop. 2.2. Here we assume two-dimensional features (both are improvable) and L_∞ norm for $\mu(\mathbf{x}) = \|\mathbf{x}\|_\infty$. The classifier f_{EI} achieves equal improvability (EI) since the same portion (1 out of 3) of unqualified samples in each group can be improved to qualified samples.

following deterministic prediction: $\hat{\mathrm{y}} \mid \mathbf{x} = \mathbf{1}\{f(\mathbf{x}) \geq 0.5\}$ where $\mathbf{1}\{A\} = 1$ if condition A holds, and $\mathbf{1}\{A\} = 0$ otherwise. We now introduce our new fairness notion.

Definition 2.1 (Equal Improvability)**.** Define a norm $\mu : \mathbb{R}^{d_{\text{I}}} \to [0, \infty)$. For a given constant $\delta > 0$, a classifier f is said to achieve *equal improvability with δ-effort* if

$$\mathbb{P}\left(\max_{\mu(\Delta \mathbf{x}_{\text{I}}) \leq \delta} f(\mathbf{x} + \Delta \mathbf{x}) \geq 0.5 \mid f(\mathbf{x}) < 0.5, \mathrm{z} = z\right) = \mathbb{P}\left(\max_{\mu(\Delta \mathbf{x}_{\text{I}}) \leq \delta} f(\mathbf{x} + \Delta \mathbf{x}) \geq 0.5 \mid f(\mathbf{x}) < 0.5\right)$$

holds for all $z \in \mathcal{Z}$, where $\Delta \mathbf{x}_{\text{I}}$ is the effort for improvable features and $\Delta \mathbf{x} = (\Delta \mathbf{x}_{\text{I}}, \mathbf{0}, \mathbf{0})$.

Note that the condition $f(\mathbf{x}) < 0.5$ represents that an individual is unqualified, and $f(\mathbf{x} + \Delta \mathbf{x}) \geq 0.5$ implies that the effort $\Delta \mathbf{x}$ allows the individual to become qualified. The above definition of fairness in *equal improvability* requires that unqualified individuals from different groups $z \in [Z]$ are equally likely to become qualified if appropriate effort is made. Note that μ can be defined on a case-by-case basis. For example, we can use $\|\mathbf{x}_{\text{I}}\| = \sqrt{\mathbf{x}_{\text{I}}{}^\top \boldsymbol{C} \mathbf{x}_{\text{I}}}$, where $\boldsymbol{C} \in \mathbb{R}^{d_{\text{I}} \times d_{\text{I}}}$ is a cost matrix that is diagonal and positive definite. Here, the diagonal terms of $\boldsymbol{C}$ characterize how difficult it is to improve each feature. For instance, consider the graduate school admission problem where $\mathbf{x}_{\text{I}}$ contains improvable features "number of publications" and "GPA". Since publishing more papers is harder than raising the GPA, the corresponding diagonal term for the number of publications feature in $\boldsymbol{C}$ should be greater than that for the GPA feature. The constant δ in Definition 2.1 can be selected depending on the classification task and the features. Appendix B.1 contains the interpretation of each term in Def. 2.1.

We introduce an alternative equivalent definition of EI fairness below.

Proposition 2.2. *The EI fairness notion defined in Def. 2.1 has an equivalent format: a classifier f achieves equal improvability with δ-effort if and only if*

$$\mathbb{P}(\mathbf{x} \in \mathcal{X}_-^{\text{imp}} \mid \mathbf{x} \in \mathcal{X}_-, \mathrm{z} = z) = \mathbb{P}(\mathbf{x} \in \mathcal{X}_-^{\text{imp}} \mid \mathbf{x} \in \mathcal{X}_-)$$

holds for all $z \in \mathcal{Z}$ where $\mathcal{X}_- = \{\mathbf{x} : f(\mathbf{x}) < 0.5\}$ is the set of features $\mathbf{x}$ for unqualified samples, and $\mathcal{X}_-^{\text{imp}} = \{\mathbf{x} : f(\mathbf{x}) < 0.5, \max_{\mu(\Delta \mathbf{x}_I) \leq \delta} f(\mathbf{x} + \Delta \mathbf{x}) \geq 0.5\}$ is the set of features $\mathbf{x}$ for unqualified samples that can be improved to qualified samples by adding $\Delta \mathbf{x}$ satisfying $\mu(\Delta \mathbf{x}_I) \leq \delta$.

The proof of this proposition is trivial from the definition of $\mathcal{X}_-$ and $\mathcal{X}_-^{\text{imp}}$. Note that $\mathbb{P}(\mathbf{x} \in \mathcal{X}_-^{\text{imp}} \mid \mathbf{x} \in \mathcal{X}_-)$ in the above equation indicates the probability that unqualified samples can be improved to qualified samples by changing the features within budget δ. This is how we define the "improvability" of unqualified samples, and the EI fairness notion is equalizing this improvability for all groups.

Visualization of EI. Fig. 2 shows the geometric interpretation of EI fairness notion in Def. 2.1 and Prop. 2.2, for a simple two-dimensional dataset having 12 samples in two groups $z \in \{\text{red}, \text{blue}\}$. Consider a linear classifier f_{EI} shown in the figure, where the samples at the right-hand-side of the decision boundary is classified as qualified samples ($f_{\text{EI}}(\mathbf{x}) \geq 0.5$). In Fig. 2a, we have L_∞ ball at each unqualified sample, representing that these samples have a chance to improve their feature in a way that the improved feature $\mathbf{x} + \Delta \mathbf{x}$ allows the sample to be classified as qualified, *i.e.*, $f_{\text{EI}}(\mathbf{x} + \Delta \mathbf{x}) \geq 0.5$. One can confirm that $\mathbb{P}\left(\max_{\mu(\Delta \mathbf{x}_{\text{I}}) \leq \delta} f(\mathbf{x} + \Delta \mathbf{x}) \geq 0.5 \mid f(\mathbf{x}) < 0.5, \mathrm{z} = z\right) = \frac{1}{3}$ holds for each group $z \in \{\text{red}, \text{blue}\}$, thus satisfying equal improvability according to Def. 2.1. In Fig. 2b, we check this in an alternative way by using the EI fairness definition in Prop. 2.2. Here, instead of making a set of improved features at each sample, we partition the feature domain $\mathcal{X}$ into

Table 1: Comparison of our EI fairness with existing fairness notions.

Name of fairness	Definition	Consider Dynamics?	Limitations
Equal Improvability (Ours)	$\mathbb{P}\left(\max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})>0.5 \mid f(\mathbf{x})<0.5, \mathrm{z}=z\right)=\mathbb{P}\left(\max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})>0.5 \mid f(\mathbf{x})<0.5\right)$	Yes	-
Demographic Parity	$\mathbb{P}\left(f(\mathbf{x})>0.5 \mid \mathrm{z}=z\right)=\mathbb{P}\left(f(\mathbf{x})>0.5\right)$	No	-
Equal Opportunity (Hardt et al., 2016)	$\mathbb{P}\left(f(\mathbf{x})>0.5 \mid \mathrm{y}=1, \mathrm{z}=z\right)=\mathbb{P}\left(f(\mathbf{x})>0.5 \mid \mathrm{y}=1\right)$	No	-
Equalized Odds (Hardt et al., 2016)	$\mathbb{P}\left(f(\mathbf{x})>0.5 \mid \mathrm{y}=y, \mathrm{z}=z\right)=\mathbb{P}\left(f(\mathbf{x})>0.5 \mid \mathrm{y}=y\right)$	No	-
Bounded Effort (Heidari et al., 2019)	$\mathbb{P}\left(\max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})>0.5, f(\mathbf{x})<0.5 \mid \mathrm{z}=z\right)=\mathbb{P}\left(\max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})>0.5, f(\mathbf{x})<0.5\right)$	Yes	Not able to handle imbalanced negative prediction rates
Equal Recourse (Gupta et al., 2019)	$\mathbb{E}\left[\min_{f(\mathbf{x}')>0.5} \mu(\mathbf{x}'-\mathbf{x}) \mid f(\mathbf{x})<0.5, \mathrm{z}=z\right]=\mathbb{E}\left[\min_{f(\mathbf{x}')>0.5} \mu(\mathbf{x}'-\mathbf{x}) \mid f(\mathbf{x})<0.5\right]$	Yes	Vulnerable to outliers

three parts: (i) the features for qualified samples $\mathcal{X}_{+}=\{\mathbf{x}: f_{\mathrm{EI}}(\mathbf{x})\geq 0.5\}$, (ii) the features for unqualified samples that can be improved $\mathcal{X}_{-}^{\mathrm{imp}}=\{\mathbf{x}: f_{\mathrm{EI}}(\mathbf{x})<0.5, \max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})\geq 0.5\}$ and (iii) the features for unqualified samples that cannot be improved $\mathcal{X}_{-}^{\mathrm{imp}}=\{\mathbf{x}: f_{\mathrm{EI}}(\mathbf{x})<0.5, \max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})<0.5\}$. In the figure, (ii) is represented as the green region and (iii) is shown as the yellow region. From Prop. 2.2, EI fairness means that $\frac{\text{\# samples in (ii)}}{\text{\# samples in (ii) + \# samples in (iii)}}$ is identical at each group $z\in\{\text{red}, \text{blue}\}$, which is true for the example in Fig. 2b.

Comparison of EI with related work. The suggested fairness notion *equal improvability* (EI) is in stark difference with existing popular zero-order fairness notions, *e.g.*, demographic parity, which can be "myopic" and focus only on achieving classification fairness in the current status. Our notion instead, uses classification fairness as a tool to equalize the true qualification status of different groups in the long run, thereby promoting social fairness. We here also note that EI has differences with existing first-order fairness notions that capture the dynamics of samples (Heidari et al., 2019; Huang et al., 2019; Gupta et al., 2019). Table 1 compares our fairness notion with the related existing notions. In particular, Bounded Effort (BE) fairness proposed by Heidari et al. (2019) equalizes 'the available reward after each individual making a bounded effort' for different groups, which is very similar to EI when we set a proper reward function. To be more specific, the BE fairness can be represented as in the Table 1. Comparing this BE expression with EI in Definition 2.1, one can confirm the difference: the inequality $f(\mathbf{x})<0.5$ is located at the conditional part for EI, which is not true for BE. EI and BE are identical if the negative prediction rates are equal across the groups, but in general they are different. The condition $f(\mathbf{x})<0.5$ here is very important since only looking into the unqualified members makes more sense when we consider *improvability*. More importantly, the BE definition is based on reward functions and we are presenting BE in a form that is closest to our EI fairness expression. Besides, Equal Recourse (ER) fairness proposed by Gupta et al. (2019) suggests to equalize the average effort of different groups without limiting the amount of effort that each sample can make. Note that ER is vulnerable to outliers. For example, when we have an unqualified outlier sample that is located far way from the decision boundary, ER disparity will be dominated by this outlier and fail to reflect the overall unfairness. Huang et al. (2019) proposed a causal-based fairness notion to equalize the minimum level of effort such that the expected prediction score of the groups are equal to each other. Note that, their definition is specific to causal-setting and it considers the sensitive groups not the rejected samples of the sensitive groups.

Compatibility of EI with other fairness notions. Here we prove the compatibility of three fairness notions (EI, DP, and BE), under two mild assumptions. Assumption 2.3 ensures that EI is well-defined, while Assumption 2.4 implies that the norm μ and the effort budget δ are chosen such that we have nonzero probability that unqualified individuals can become qualified after making efforts.

Assumption 2.3. For any classifier f, the probability of unqualified samples for each demographic group is not equal to 0, *i.e.*, $\mathbb{P}\left(f(\mathbf{x})<0.5, \mathrm{z}=z\right)\neq 0$ for all $z\in\mathcal{Z}$.

Assumption 2.4. For any classifier f, the probability of being qualified after the effort for unqualified samples is not equal to 0, *i.e.*, $\mathbb{P}\left(\max_{\mu(\Delta\mathbf{x}_{\mathrm{I}})\leq\delta} f(\mathbf{x}+\Delta\mathbf{x})\geq 0.5, f(\mathbf{x})<0.5\right)\neq 0$.

Under these assumptions, the following theorem reveals the relationship between DP, EI and BE.

Theorem 2.5. *If a classifier f achieves two of the following three fairness notions, DP, EI, and BE; then it has to achieve the remaining fairness notion as well.*

The proof of the Theorem 2.5 is provided in Appendix A. This theorem immediately implies the following corollary, which provides a condition such that EI and BE conflict with each other.

Corollary 2.6. *The above theorem says that if a classifier f achieves EI and BE, it has to achieve DP. Thus, by contraposition, if f does not achieve DP, then it cannot achieve EI and BE simultaneously.*

3 ACHIEVING EQUAL IMPROVABILITY

In this section, we discuss methods for finding a classifier that achieves EI fairness. Following existing in-processing techniques (Zafar et al., 2017c; Donini et al., 2018; Zafar et al., 2017a; Cho et al., 2020), we focus on finding a fair classifier by solving a fairness-regularized optimization problem. To be specific, we first derive a differentiable penalty term U_δ that approximates the unfairness with respect to EI, and then solve a regularized empirical minimization problem having the unfairness as the regularization term. This optimization problem can be represented as

$$\min_{f\in\mathcal{F}} \left\{ \frac{(1-\lambda)}{N} \sum_{i=1}^{N} \ell(\mathrm{y}_i, f(\mathbf{x}_i)) + \lambda U_\delta \right\}, \tag{1}$$

where $\{(\mathbf{x}_i, \mathrm{y}_i)\}_{i=1}^N$ is the given dataset, $\ell : \{0,1\} \times [0,1] \to \mathbb{R}$ is the loss function, $\mathcal{F}$ is the set of classifiers we are searching over, and $\lambda \in [0,1)$ is a hyperparameter that balances fairness and prediction loss. Here we consider three different ways of defining the penalty term U_δ, which are (a) covariance-based, (b) kernel density estimator (KDE)-based, and (c) loss-based methods. We first introduce how we define U_δ in each method, and then discuss how we solve equation 1.

Covariance-based EI Penalty. Our first method is inspired by Zafar et al. (2017c), which measures the unfairness of a score-based classifier f by the covariance of the sensitive attribute z and the score $f(\mathbf{x})$, when the demographic parity (DP) fairness condition ($\mathbb{P}(f(\mathbf{x}) > 0.5|\mathrm{z} = z) = \mathbb{P}(f(\mathbf{x}) > 0.5)$ holds for all z) is considered. The intuition behind this idea of measuring the covariance is that a perfect fair DP classifier should have zero correlation between z and $f(\mathbf{x})$. By applying similar approach to our fairness notion in Def. 2.1, the EI unfairness is measured by the covariance between the sensitive attribute z and the maximally improved score of rejected samples within the effort budget. In other words, $(\mathrm{Cov}(\mathrm{z}, \max_{\|\Delta\mathbf{x}_\mathrm{I}\|\le\delta} f(\mathbf{x}+\Delta\mathbf{x}) \mid f(\mathbf{x}) < 0.5))^2$ represents the EI unfairness of a classifier f where we took the square to penalize negative correlation case as well. Let $I_- = \{i : f(\mathbf{x}_i) < 0.5\}$ be the set of indices of unqualified samples, and $\bar{\mathrm{z}} = \sum_{i\in I_-} \mathrm{z}_i / |I_-|$. Then, EI unfairness can be approximated by the square of the empirical covariance, *i.e.*,

$$U_\delta \triangleq \left(\frac{1}{|I_-|} \sum_{i\in I_-} (\mathrm{z}_i - \bar{\mathrm{z}}) \left(\max_{\|\Delta\mathbf{x}_{\mathrm{I}_i}\|\le\delta} f(\mathbf{x}_i + \Delta\mathbf{x}_i) - \sum_{j\in I_-} \max_{\|\Delta\mathbf{x}_{\mathrm{I}_j}\|\le\delta} f(\mathbf{x}_j + \Delta\mathbf{x}_j)/|I_-| \right) \right)^2.$$

Since $\sum_{i\in I_-} (\mathrm{z}_i - \bar{\mathrm{z}}) \left(\sum_{j\in I_-} \max_{\|\Delta\mathbf{x}_{\mathrm{I}_j}\|\le\delta} f(\mathbf{x}_j + \Delta\mathbf{x}_j)/|I_-| \right) = 0$ from $\sum_{i\in I_-} (\mathrm{z}_i - \bar{\mathrm{z}}) = 0$, we have $U_\delta = \left(\frac{1}{|I_-|} \sum_{i\in I_-} (\mathrm{z}_i - \bar{\mathrm{z}}) \max_{\|\Delta\mathbf{x}_{\mathrm{I}_i}\|\le\delta} f(\mathbf{x}_i + \Delta\mathbf{x}_i) \right)^2$.

KDE-based EI Penalty. The second method is inspired by Cho et al. (2020), which suggests to first approximate the probability density function of the score $f(\boldsymbol{x})$ via kernel density estimator (KDE) and then put the estimated density formula into the probability term in the unfairness penalty. Recall that given m samples $\mathrm{y}_1, \ldots, \mathrm{y}_m$, the true density g_y on y is estimated by KDE as $\hat{g}_\mathrm{y}(\hat{\mathrm{y}}) \triangleq \frac{1}{mh} \sum_{i=1}^m g_k \left(\frac{\hat{\mathrm{y}} - \mathrm{y}_i}{h} \right)$, where g_k is a kernel function and h is a smoothing parameter.

Here we apply this KDE-based method for estimating the EI penalty term in Def. 2.1. Let $\mathrm{y}_i^{\max} = \max_{\|\Delta\mathbf{x}_{\mathrm{I}_i}\|\le\delta} f(\mathbf{x}_i + \Delta\mathbf{x}_i)$ be the maximum score achievable by improving feature $\mathbf{x}_i$ within budget δ, and $I_{-,z} = \{i : f(\mathbf{x}_i) < 0.5, \mathrm{z}_i = z\}$ be the set of indices of unqualified samples of group z. Then, the density of $\mathrm{y}_i^{\max}$ for the unqualified samples in group z can be approximated as[2]

$$\hat{g}_{\mathrm{y}^{\max}|f(\mathbf{x})<0.5,z}(\hat{\mathrm{y}}^{\max}) \triangleq \frac{1}{|I_{-,z}|h} \sum_{i\in I_{-,z}} g_k \left(\frac{\hat{\mathrm{y}}^{\max} - \mathrm{y}_i^{\max}}{h} \right).$$

Then, the estimate on the left-hand-side (LHS) probability term in Def. 2.1 is represented as $\hat{\mathbb{P}}\left(\max_{\mu(\Delta\mathbf{x}_\mathrm{I})\le\delta} f(\mathbf{x}+\Delta\mathbf{x}) \ge 0.5 \mid f(\mathbf{x}) < 0.5, \mathrm{z} = z\right) = \int_{0.5}^{\infty} \hat{g}_{\mathrm{y}^{\max}|f(\mathbf{x})<0.5,z}(\hat{\mathrm{y}}^{\max}) \mathrm{d}\hat{\mathrm{y}}^{\max} = \frac{1}{|I_{-,z}|h} \sum_{i\in I_{-,z}} G_k \left(\frac{0.5 - \mathrm{y}_i^{\max}}{h} \right)$ where $G_k(\tau) \triangleq \int_\tau^\infty g_k(y)\mathrm{d}y$. Similarly, we can estimate the right-hand-side (RHS) probability term in Def. 2.1, and the EI-penalty U_δ is computed as the summation of the absolute difference of the two probability values (LHS and RHS) among all groups z.

[2]This term is differentiable with respect to model parameters, since g_k is differentiable with respect to $\mathrm{y}_i^{\max}$, and $\mathrm{y}_i^{\max} = \max_{\|\Delta\mathbf{x}_{\mathrm{I}_i}\|\le\delta} f(\mathbf{x}_i + \Delta\mathbf{x}_i)$ is differentiable w.r.t. model parameters from (Danskin, 1967).

Loss-based EI Penalty. Another common way of approximating the fairness violation as a differentiable term is to compute the absolute difference of group-specific losses (Roh et al., 2021; Shen et al., 2022). Following the spirit of EI notion in Def. 2.1, we define EI loss of group z as $\tilde{L}_z \triangleq \frac{1}{|I_{-,z}|}\sum_{i\in I_{-,z}} \ell\left(1, \max_{\|\Delta\mathbf{x}_i\|\le\delta} f(\mathbf{x}_i+\Delta\mathbf{x}_i)\right)$. Here, $\tilde{L}_z$ measures how far the rejected samples in group z are away from being accepted after the feature improvement within budget δ. Similarly, EI loss for all groups is written as $\tilde{L} \triangleq \sum_{z\in\mathcal{Z}} \frac{I_{-,z}}{I_-}\tilde{L}_z$. Finally, the EI penalty term is defined as $U_\delta \triangleq \sum_{z\in\mathcal{Z}} \left|\tilde{L}_z - \tilde{L}\right|$.

Solving equation 1. For each approach defined above (covariance-based, KDE-based and loss-based), the penalty term U_δ is defined uniquely. Note that in all cases, we need to solve a maximization problem $\max_{\|\Delta\mathbf{x}_1\|\le\delta} f(\mathbf{x}+\Delta\mathbf{x})$ in order to get U_δ. Since equation 1 is a minimization problem containing U_δ in the cost function, it is essentially a minimax problem. We leverage adversarial training techniques to solve equation 1. The inner maximization problem is solved using one of two methods: (i) derive the closed-form solution for generalized linear models, (ii) use projected gradient descent for general settings. The details can be found in Appendix B.2.

4 Experiments

In this section, we provide empirical results on our EI fairness notion. First, we show that our methods suggested in Sec. 3 achieve EI fairness in various real/synthetic datasets. Second, focusing on the dynamic scenario where each individual can make effort to improve its outcome, we demonstrate that training an EI classifier at each time step promotes achieving the long-term fairness, *i.e.*, the feature distribution of two groups become identical in the long run. Codes are available in anonymous GitHub repository [3].

4.1 Suggested methods achieve EI fairness

Recall that Sec. 3 provided three approaches for achieving EI fairness. Here we check whether such methods successfully find a classifier with small EI disparity, compared with ERM which does not have fairness constraints. Due to the space limitation, the main body contains the results for logistic regression (LR) only, and in Appendix C contains the performance for multi-layer perceptron (MLP) as well as the effect of over-parameterization.

Experiment setting. For all experiments, we use the Adam optimizer and cross entropy loss. We perform cross-validation on the training set to find the best hyperparameter. We provide statistics for five trials having different random seeds. For KDE-based approach, we use Gaussian kernel.

Datasets. We perform the experiments on one synthetic dataset, and two real datasets: German Statlog Credit (Dua & Graff, 2017), and Adult (Dua & Graff, 2017). The synthetic dataset has two non-sensitive attributes $\mathbf{x} = (\mathrm{x}_1, \mathrm{x}_2)$, one binary sensitive attribute z, and a binary label y. Both features x_1 and x_2 are assumed to be improvable. We generate 20,000 samples where $(\mathbf{x}, \mathrm{y}, \mathrm{z})$ pair of each sample is generated independently as below. We define z and $(\mathrm{y}|\mathrm{z}=z)$ as Bernoulli random variables for all $z \in \{0,1\}$, and define $(\mathbf{x}|\mathrm{y}=y, \mathrm{z}=z)$ as multivariate Gaussian random variables for all $y, z \in \{0,1\}$. The numerical details are in Appendix C.1 The maximum effort δ for this dataset is set to 0.5. The ratio of the training versus test data is 4:1.

German Statlog Credit Data contains 1,000 samples and the ratio of the training versus test data is 4:1. The task is to predict the credit risk of an individual given its financial status. Following Jiang & Nachum (2020), we divide the samples into two groups using the age of thirty as the boundary, *i.e.*, $z = 1$ for samples with age above thirty. Four features $\mathbf{x}$ are considered as improvable: `checking account`, `saving account`, `housing` and `occupation`, all of which are ordered categorical features. For example, the `occupation` feature has four levels: (1) unemployed, (2) unskilled, (3) skilled, and (4) highly qualified. We set the maximum effort $\delta = 1$, meaning that an unskilled man (with level 2) can become a skilled man, but cannot be a highly qualified man.

The Adult dataset consists of the data for 48,842 people and is split into training set and test set in the ratio of 2:1. The task is predicting whether a person's income would exceed 50K USD per year. Following Jiang & Nachum (2020), we use gender as the sensitive attribute. We select education level (ordered categorical feature) to be the improvable feature. We set the maximum effort $\delta = 2$.

[3] https://anonymous.4open.science/r/ei_fairness-23CA/

Table 2: **Error rate and EI disparities of ERM and three proposed EI-regularized methods on logistic regression (LR).** For each dataset, the lowest EI disparity (disp.) value is in boldface. Classifiers obtained by our three methods have much smaller EI disparity values than the ERM solution, without having much additional error.

Dataset	Metric	Methods			
		ERM	Covariance-based	KDE-based	Loss-based
Synthetic	Error Rate(↓)	.221 ± .001	.253 ± .003	.250 ± .001	.246 ± .001
	EI Disp.(↓)	.117 ± .007	.003 ± .001	.005 ± .003	**.002 ± .001**
German Stat.	Error Rate(↓)	.220 ± .009	.262 ± .009	.243 ± .024	.237 ± .008
	EI Disp.(↓)	.041 ± .008	.021 ± .019	.035 ± .026	**.015 ± .009**
Adult	Error Rate(↓)	.152 ± .000	.173 ± .002	.181 ± .002	.183 ± .004
	EI Disp.(↓)	.051 ± .002	.016 ± .002	.009 ± .004	**.006 ± .005**

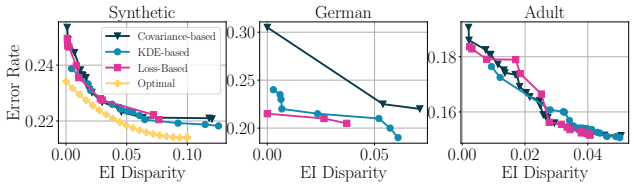

Figure 3: **Tradeoff between EI disparity and error rate.** We run three EI-regularized methods suggested in Sec. 3 for different regularizer coefficient λ and plot the frontier lines. For the synthetic dataset, the tradeoff curve for the ideal classifier is located at the bottom left corner, which is similar to the curves of proposed EI-regularized methods. This shows that our methods successfully find classifiers balancing EI disparity and error rate.

Results. Table 2 shows the test error rate and test EI disparity (disp.) for ERM and our three EI-regularized methods (covariance-based, KDE-based and loss-based) suggested in Sec. 3. For all three datasets, our EI-regularized methods successfully reduce the EI disparity without increasing the error rate too much, compared with ERM. Figure 3 shows the tradeoff between error rate and EI disparity of our EI-regularized methods. We marked the dots after running each method multiple times with different penalty coefficient λ, and plotted the frontier line. For the synthetic dataset with Gaussian feature, we numerically obtained the performance of the optimal EI classifier, which is added in the yellow line at the bottom left corner of the first column plot. The details of finding the optimal EI classifier is in Appendix B.3. One can confirm that our methods regularizing EI are having similar tradeoff curves for the synthetic dataset and Adult dataset. Especially, for synthetic dataset, the tradeoff curve of our methods nearly achieves that of the optimal EI classifier. For German dataset, the loss-based method is having a slightly better tradeoff curve than other methods.

4.2 EI promotes long-term fairness in dynamic scenarios

Recall that the motivation for proposing EI is to achieve *long-term fairness*, which is equalizing the feature distribution of samples in different groups in the long run, under the dynamic scenario where each individual can improve its feature. In this section, we provide simulation results on dynamic setting, which show that training a classifier with EI constraint encourages the long-term fairness.

4.2.1 Dynamic System Description

We consider a binary classification problem under the dynamic scenario with T rounds, where the improvable feature $\mathbf{x}_t \in \mathbb{R}$ and the label $\mathbf{y}_t \in \{0, 1\}$ of each sample as well as the classifier f_t evolve at each round $t \in \{0, \cdots, T-1\}$. We denote the sensitive attribute as $\mathrm{z} \in \{0, 1\}$, and the estimated label as $\hat{\mathrm{y}}_t$. We assume $\mathrm{z} \sim \mathrm{Bern}(0.5)$ and $(\mathrm{x}_t \mid \mathrm{z} = z) \sim \mathcal{P}_t^{(z)} = \mathcal{N}(\mu_t^{(z)}, \{\sigma_t^{(z)}\}^2)$. To mimic the admission problem, we only accept a fraction $\alpha \in (0, 1)$ of the population, *i.e.*, the true label is modeled as $\mathrm{y}_t = \mathbf{1}_{\mathrm{x}_t \geq \chi_\alpha^{(t)}}$, where $\chi_\alpha^{(t)}$ is the $(1-\alpha)$ percentile of the feature distribution at round t. We consider z-aware linear classifier outputting $\hat{\mathrm{y}}_t = \mathbf{1}_{\mathrm{x}_t \geq \tau_t^{(z)}}$, which is parameterized by the thresholds $(\tau_t^{(0)}, \tau_t^{(1)})$ for two groups. Note that this classification rule is equivalent to defining score function $f_t(x, z) = 1/(\exp(x - \tau_t^{(z)}) + 1)$ and $\hat{\mathrm{y}}_t = \mathbf{1}_{f(\mathbf{x}_t, \mathrm{z}) \geq 0.5}$.

Updating data parameters $(\mu_t^{(z)}, \sigma_t^{(z)})$. At each round t, we allow each sample can improve its feature from x to $x + \epsilon(x)$. Here we model $\epsilon(x) = \nu(x; z) = \frac{1}{(\tau_t^{(z)} - x + \beta)^2} \mathbf{1}_{x < \tau_t^{(z)}}$ for a constant $\beta > 0$. In this model, the rejected samples with larger gap $(\tau_t^{(z)} - x)$ with the decision boundary are making less effort Δx, which is inspired by the intuition that a rejected sample is less motivated to improve its feature if it needs to take a large amount of effort to get accepted in one scoop. After such effort is made, we compute the mean and standard deviation of each group: $\mu_{t+1}^{(z)} = \int_{-\infty}^{\infty} (x + \nu(x; z)) \phi(x; \mu_t^{(z)}, \sigma_t^{(z)}) \mathrm{d}x$ and $\sigma_{t+1}^{(z)} = \sqrt{\int_{-\infty}^{\infty} (x + \nu(x; z) - \mu_{t+1}^{(z)})^2 \phi(x; \mu_t^{(z)}, \sigma_t^{(z)}) \mathrm{d}x}$,

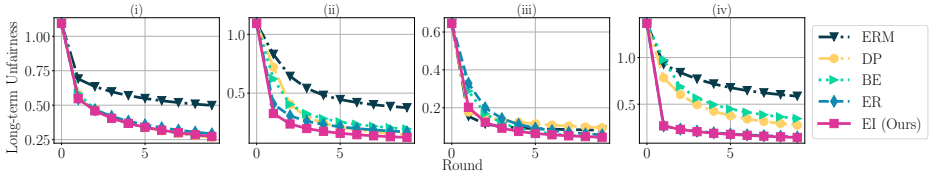

Figure 4: **Long-term unfairness $d_{\mathrm{TV}}(\mathcal{P}_t^{(0)}, \mathcal{P}_t^{(1)})$ at each round t for various algorithms.** Consider the binary classification problem over two groups, under the dynamic scenario where the data distribution and the classifier for each group evolves over multiple rounds. We plot how the long-term unfairness (measured by the total variation distance between two groups) changes as round t increases. Here, each column shows the result for different initial feature distribution, details of which are given in Sec. 4.2.2. The long-term unfairness of EI classifier reduces faster than other existing fairness notions, showing that EI proposed in this paper is helpful for achieving long-term fairness.

where $\phi(\cdot; \mu, \sigma)$ is the pdf of $\mathcal{N}(\mu, \sigma^2)$. We assume that the feature $\mathbf{x}_{t+1}$ in the next round follows a Gaussian distribution parameterized by $(\mu_{t+1}^{(z)}, \sigma_{t+1}^{(z)})$ for ease of simulation.

Updating classifier parameters $(\tau_t^{(0)}, \tau_t^{(1)})$. At each round t, we update the classifier depending on the current feature distribution $\mathbf{x}_t$. The EI classifier considered in this paper updates $(\tau_t^{(0)}, \tau_t^{(1)})$ as below. Note that the maximized score $\max_{\|\Delta \mathbf{x}_1\| \le \delta} f(\mathbf{x} + \Delta \mathbf{x})$ in Def. 2.1 can be written as $f_t(\mathrm{x}_t + \delta_t, \mathrm{z})$, and the equation $\max_{\|\Delta \mathbf{x}_1\| \le \delta} f(\mathbf{x} + \Delta \mathbf{x}) \ge 0.5$ is equivalent to $\mathrm{x}_t + \delta_t > \tau_t^{(z)}$. Consequently, EI classifier obtains $(\tau_t^{(0)}, \tau_t^{(1)})$ by solving

$$\min_{\tau_t^{(0)}, \tau_t^{(1)}} \left| \mathbb{P}(\mathrm{x}_t + \delta_t > \tau_t^{(0)} \,|\, \mathrm{z} = 0, \mathrm{x}_t < \tau_t^{(0)}) - \mathbb{P}(\mathrm{x}_t + \delta_t > \tau_t^{(1)} \,|\, \mathrm{z} = 1, \mathrm{x}_t < \tau_t^{(1)}) \right| \text{ s.t. } \mathbb{P}(\hat{\mathrm{y}}_t \neq \mathrm{y}_t) \le c,$$

where $c \in [0, 1)$ is the maximum classification error rate we allow, and δ_t is the effort level at iteration t. In our experiments, δ_t is chosen as the mean efforts the population makes, *i.e.*, $\delta_t = 0.5 \sum_{z=0}^{1} \int_{-\infty}^{\infty} \nu(x; z) \phi(x; \mu_t^{(z)}, \sigma_t^{(z)}) \mathrm{d}x$. We can similarly obtain the classifier for DP, BE and ER constraints, details of which are in Appendix C.4. In the experiments, we numerically obtain the solution of this optimization problem.

4.2.2 Experiments on Long-term Fairness

We first initialize the feature distribution in a way that both sensitive groups have either different mean (*i.e.*, $\mu_0^{(0)} \neq \mu_0^{(1)}$) or different variance (*i.e.*, $\sigma_0^{(0)} \neq \sigma_0^{(1)}$). At each round $t \in \{1, \cdots, T\}$, we update the data parameter $(\mu_t^{(z)}, \sigma_t^{(z)})$ for group $z \in \{0, 1\}$ and the classifier parameter $(\tau_t^{(0)}, \tau_t^{(1)})$, following the rule described in Sec. 4.2.1. At each round $t \in \{1, \cdots, T\}$, we measure the *long-term unfairness* defined as the total variation distance between the two group distributions: $d_{\mathrm{TV}}(\mathcal{P}^{(0)}, \mathcal{P}^{(1)}) = \int_{-\infty}^{\infty} |\phi(x; \mu_t^{(0)}, \sigma_t^{(0)}) - \phi(x; \mu_t^{(1)}, \sigma_t^{(1)})| \mathrm{d}x$. We run experiments on four different initial feature distributions: (i) $(\mu_0^{(0)}, \sigma_0^{(0)}, \mu_0^{(1)}, \sigma_0^{(1)}) = (0, 1, 1, 0.5)$, (ii) $(\mu_0^{(0)}, \sigma_0^{(0)}, \mu_0^{(1)}, \sigma_0^{(1)}) = (0, 0.5, 1, 1)$, (iii) $(\mu_0^{(0)}, \sigma_0^{(0)}, \mu_0^{(1)}, \sigma_0^{(1)}) = (0, 2, 0, 1)$, and (iv) $(\mu_0^{(0)}, \sigma_0^{(0)}, \mu_0^{(1)}, \sigma_0^{(1)}) = (0, 0.5, 1, 0.5)$, respectively. We set $\alpha = 0.2, c = 0.1, \beta = 0.25$.

Baselines. We compare our EI classifier with multiple baselines, including the empirical risk minimization (ERM) and algorithms with fairness constraints: demographic parity (DP), bounded effort (BE) (Heidari et al., 2019), and equal recourse (ER) (Gupta et al., 2019).

Results. Fig. 4 shows how the long-term unfairness $d_{\mathrm{TV}}(\mathcal{P}_t^{(0)}, \mathcal{P}_t^{(1)})$ changes as a function of round t, for cases (i) – (iv) having different initial feature distribution. Note that fairness-constrained algorithms (DP, BE, ER and EI) enjoys lower long-term unfairness compared with ERM, for case (i), (ii) and (iv). More importantly, EI accelerates the process of mitigating long-term unfairness, compared to other fairness notions. This observation highlights the benefit of EI in promoting true equity of groups in the long-term. Fig. 5 visualizes the initial distribution (at the leftmost column) and the evolved distribution at round $t = 3$ for multiple algorithms (at the rest of the columns). Each row represents different initial feature distribution, for cases (i) – (iv). One can confirm that EI brings the distribution of the two groups closer, compared with ERM, DP, BE, and ER.

5 Related Works

Fairness metric. Group fairness is a class of fairness notions requiring the classifier to treat different groups similarly. Most of the existing group fairness notions are myopic, measuring fairness by only comparing the positive rate of certain groups or subgroups at the current snapshot (Hardt et al., 2016;

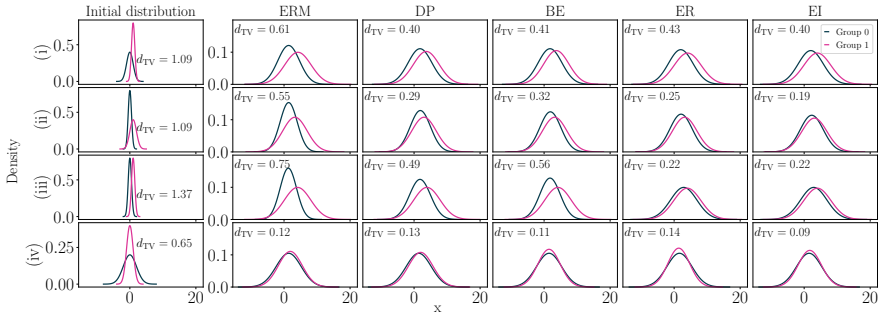

Figure 5: **Evolution of the feature distribution, when we apply each algorithm for $t = 3$ rounds.** At each row, the leftmost column shows the initial distribution and the rest of the columns show the evolved distribution for each algorithm, under the dynamic setting. Compared with existing fairness notions (DP, BE, and ER), our EI fairness achieves smaller feature distribution gap between groups.

Zafar et al., 2017a). In contrast, EI suggested in this paper is more farsighted by taking improvability into consideration. There exist a few works that propose effort-based fairness notions (Heidari et al., 2019; Huang et al., 2019; Gupta et al., 2019) that capture the potential follow-up impact. EI is closely related to, but also clearly distinct from these effort-based fairness notions, particularly in terms of motivation. Note that Heidari et al. (2019) aims to propose fairness notions that capture the dynamics, while Gupta et al. (2019) and Huang et al. (2019) aim to propose reasonable fairness notions that take effort into consideration. Our work takes one-step further, focusing on promoting the long-term fairness. Besides, Huang et al. (2019) only consider the causal setting and are restricted to certain models. Without the restricted assumptions introduced by Huang et al. (2019), our work applies to more general settings, and provides a fairness metric that is easy to measure. We provide a more systematic comparison between EI and related work in Sec. 2.

Fairness-aware algorithms. Most of the existing fair learning techniques fall into three categories: i) pre-processing approaches (Kamiran & Calders, 2012; 2010; Gordaliza et al., 2019; Jiang & Nachum, 2020), which primarily involves massaging the dataset to remove the bias; ii) in-processing approaches (Fukuchi et al., 2013; Kamishima et al., 2012; Calders & Verwer, 2010; Zafar et al., 2017c;a; Zhang et al., 2018; Cho et al., 2020; Roh et al., 2020; 2021; Shen et al., 2022), adjusting the model training for fairness; iii) post-processing approaches (Calders & Verwer, 2010; Alghamdi et al., 2020; Wei et al., 2020; Hardt et al., 2016) which achieve fairness by modifying a given unfair classifier. Prior work (Woodworth et al., 2017) showed that the in-processing approach generally outperforms other approaches due to its flexibility. Hence, we focus on the in-processing approach and propose three methods to achieve EI. These methods achieve EI by solving fairness-regularized optimization problems. In particular, our proposed fairness regularization terms are inspired by Zafar et al. (2017c); Cho et al. (2020); Roh et al. (2021); Shen et al. (2022).

Fairness dynamics. There are also a few attempts to study the long-term impact of different fairness policies (Zhang et al., 2020; Heidari et al., 2019; Hashimoto et al., 2018). In particular, Hashimoto et al. (2018) studies how ERM amplifies the unfairness of a classifier in the long run. The key idea is that if the classifier of the previous iteration favors a certain candidate group, then the candidate groups will be more unbalanced since fewer individuals from the unflavored group will apply for this position. Thus, the negative feedback leads to a more unfair classifier. In contrast, Heidari et al. (2019) and Zhang et al. (2020) focus more on long-term impact instead of classification fairness. To be specific, Heidari et al. (2019) studies how fairness intervention affects the different groups of individuals in terms of evenness, centralization, and clustering by simulating the population's response through effort. Zhang et al. (2020) investigates how different fairness policies affect the gap between the qualification rates of different groups under a partially observed Markov decision process. Besides, there are a few works which study how individuals may take strategic actions to improve their outcomes given a classifier (Chen et al., 2021). However, Chen et al. (2021) aims to address this problem by designing an optimization problem that is robust to strategic manipulation, which is orthogonal to our main focus.

6 Conclusion

In this paper, we proposed a novel group fairness notion called Equal Improvability (EI), which equalizes the potential acceptance of rejected samples in different groups, when appropriate effort is made by the rejected samples. We analyzed the properties of EI fairness, and provided three approaches for finding a classifier that achieves EI by using adversarial training. Experimental results showed that the proposed approaches reduce the EI disparity. Lastly, we formulated a dynamic model to showcase the benefit of EI in promoting the equity of feature distribution of different groups.

REFERENCES

Wael Alghamdi, Shahab Asoodeh, Hao Wang, Flavio P. Calmon, Dennis Wei, and Karthikeyan Natesan Ramamurthy. Model projection: Theory and applications to fair machine learning. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 2711–2716, 2020. doi: 10.1109/ISIT44484.2020.9173988.

Toon Calders and Sicco Verwer. Three naive bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2):277–292, 2010. doi: 10.1007/s10618-010-0190-x. URL https://doi.org/10.1007/s10618-010-0190-x.

Yatong Chen, Jialu Wang, and Yang Liu. Linear classifiers that encourage constructive adaptation, 2021.

Valeriia Cherepanova, Vedant Nanda, Micah Goldblum, John P Dickerson, and Tom Goldstein. Technical challenges for training fair neural networks. *arXiv preprint arXiv:2102.06764*, 2021.

Jaewoong Cho, Gyeongjo Hwang, and Changho Suh. A fair classifier using kernel density estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:15088–15099, 2020.

John Danskin. The theory of max-min and its application to weapons allocation problems. In *Springer*, 1967.

Michele Donini, Luca Oneto, Shai Ben-David, John Shawe-Taylor, and Massimiliano Pontil. Empirical risk minimization under fairness constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, pp. 2796–2806, 2018.

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

Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pp. 214–226, 2012. ISBN 9781450311151. doi: 10.1145/2090236.2090255. URL https://doi.org/10.1145/2090236.2090255.

Kazuto Fukuchi, Jun Sakuma, and Toshihiro Kamishima. Prediction with model-based neutrality. In *Machine Learning and Knowledge Discovery in Databases*, pp. 499–514, 2013. ISBN 978-3-642-40991-2.

Paula Gordaliza, Eustasio Del Barrio, Gamboa Fabrice, and Jean-Michel Loubes. Obtaining fairness using optimal transport theory. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2357–2365, 2019. URL http://proceedings.mlr.press/v97/gordaliza19a.html.

Vivek Gupta, Pegah Nokhiz, Chitradeep Dutta Roy, and Suresh Venkatasubramanian. Equalizing recourse across groups. *arXiv preprint arXiv:1909.03166*, 2019.

Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems (NIPS)*, volume 29, pp. 3315–3323, 2016.

Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1929–1938, 2018. URL http://proceedings.mlr.press/v80/hashimoto18a.html.

Hoda Heidari, Vedant Nanda, and Krishna P. Gummadi. On the long-term impact of algorithmic decision policies: Effort unfairness and feature segregation through social learning, 2019.

Wen Huang, Yongkai Wu, Lu Zhang, and Xintao Wu. Fairness through equality of effort, 2019.

Heinrich Jiang and Ofir Nachum. Identifying and correcting label bias in machine learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 702–712, 2020. URL http://proceedings.mlr.press/v108/jiang20a.html.

F. Kamiran and T.G.K. Calders. Classification with no discrimination by preferential sampling. In *Informal proceedings of the 19th Annual Machine Learning Conference of Belgium and The Netherlands*, pp. 1–6, 2010.

Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33:1–33, 2012.

Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Machine Learning and Knowledge Discovery in Databases*, pp. 35–50, 2012. ISBN 978-3-642-33486-3.

Yuji Roh, Kangwook Lee, Steven Whang, and Changho Suh. FR-train: A mutual information-based approach to fair and robust training. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8147–8157, 2020. URL http://proceedings.mlr.press/v119/roh20a.html.

Yuji Roh, Kangwook Lee, Steven Euijong Whang, and Changho Suh. Fairbatch: Batch selection for model fairness. In *International Conference on Learning Representations (ICLR)*, 2021. URL https://openreview.net/forum?id=YNnpaAKeCfx.

Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin, and Lea Frermann. Optimising equal opportunity fairness in model training. *arXiv preprint arXiv:2205.02393*, 2022.

Dennis Wei, Karthikeyan Natesan Ramamurthy, and Flavio Calmon. Optimized score transformation for fair classification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1673–1683, 2020. URL http://proceedings.mlr.press/v108/wei20a.html.

Blake Woodworth, Suriya Gunasekar, Mesrob I. Ohannessian, and Nathan Srebro. Learning non-discriminatory predictors. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1920–1953, 2017. URL http://proceedings.mlr.press/v65/woodworth17a.html.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 1171–1180, 2017a.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, Krishna P. Gummadi, and Adrian Weller. From parity to preference-based notions of fairness in classification. In *Advances in Neural Information Processing Systems (NIPS)*, NIPS'17, pp. 228–238, 2017b. ISBN 9781510860964.

Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P. Gummadi. Fairness Constraints: Mechanisms for Fair Classification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pp. 962–970, 2017c.

Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340, 2018.

Xueru Zhang, Ruibo Tu, Yang Liu, Mingyan Liu, Hedvig Kjellström, Kun Zhang, and Cheng Zhang. How do fair decisions fare in long-term qualification?, 2020.