

Revisiting (Un)Fairness in Recourse by Minimizing Worst-Case Social Burden

Ainhize Barrainkua¹, Giovanni De Toni², Jose A. Lozano^{1,3}, Novi Quadrianto^{1,4}

¹Basque Center for Applied Mathematics, Bilbao, Spain

²Fondazione Bruno Kessler, Trento, Italy

³University of the Basque Country, Donostia - San Sebastian, Spain

⁴Predictive Analytics Lab, University of Sussex, Brighton, UK

abarrainkua@bcamath.org, gdetoni@fbk.eu, ja.lozano@ehu.eus, n.quadrianto@sussex.ac.uk

Abstract

Machine learning based predictions are increasingly used in sensitive decision-making applications that directly affect our lives. This has led to extensive research into ensuring the fairness of classifiers. Beyond just fair classification, emerging legislation now mandates that when a classifier delivers a negative decision, it must also offer actionable steps an individual can take to reverse that outcome. This concept is known as *algorithmic recourse*. Nevertheless, many researchers have expressed concerns about the fairness guarantees within the recourse process itself. In this work, we provide a holistic theoretical characterization of unfairness in algorithmic recourse, formally linking fairness guarantees in recourse and classification, and highlighting limitations of the standard equal cost paradigm. We then introduce a novel fairness framework based on *social burden*, along with a practical algorithm (MISOB), broadly applicable under real-world conditions. Empirical results on real-world datasets show that MISOB reduces the social burden across all groups without compromising overall classifier accuracy.

Code — <https://github.com/abarrainkua/MISOB>

Extended version — <https://arxiv.org/abs/2509.04128>

1 Introduction

The ubiquity of automated decision-making systems has considerably reshaped the landscape of how individuals access crucial resources and are granted important social and economic opportunities, ranging from securing credit to receiving essential public services (Chouldechova et al. 2018; Purificato et al. 2023; Perdomo et al. 2025). In recent years, growing concerns have emerged regarding the fairness of these systems, as well as their societal impact (Grgic-Hlaca et al. 2018; Jobin, Ienca, and Vayena 2019; Raji et al. 2022). Notably, instances have been documented in which algorithmic decision support systems produced discriminatory outcomes in domains such as welfare distribution (Heikkilä 2022) and criminal justice (Angwin et al. 2022). Moreover, in high-stakes scenarios, it is imperative that, whenever a model yields a negative decision, it also provides the affected individual with clear and actionable *recourse* (Karimi et al. 2022). Unfortunately, affected individuals usually lack

transparency and effective means to contest undesired decisions (e.g., bail denial) or take meaningful steps to improve their situation¹. For this reason, the “*right to recourse*” is slowly emerging as a legal requirement in several legislations, e.g., in the European GDPR and the Artificial Intelligence (AI) Act (EU Commission 2021).

Pursuing a given recourse often demands considerable effort from the affected individual. As noted by Ustun, Spangher, and Liu (2019), the average effort to achieve recourse can vary significantly across sensitive groups, raising important concerns about fairness in recourse. This disparity may arise when (i) the recourse itself *differs* depending on *group membership*, or when (ii) the recourse is formally identical, but the real-world *effort* required to carry it out *varies* in practice. For example, consider a method that offers suggestions to individuals denied a loan. It might recommend “*Increase your monthly income by \$500*” to someone from a privileged group, while advising “*Maintain stable employment for 24 more months*” to someone from an underprivileged background. The first recommendation is often more viable for individuals with access to high-paying jobs or stable employment opportunities. In contrast, maintaining long-term employment can be substantially more challenging for individuals in precarious jobs or members of low-income and otherwise underprivileged communities. Further, consider an algorithm suggesting that all rejected students retake standardized tests. It may impose minimal burden on those who can afford test preparation, tutoring, and multiple attempts, while students coming from low social and economic backgrounds often face barriers, including limited resources and time constraints, due to many socio-economic factors (Rodríguez-Hernández, Cascallar, and Kyndt 2020). Thus, what appears to be group-aware or uniform recommendations can translate into very different levels of effort, reinforcing structural inequalities.

In light of these concerns, it is paramount to develop metrics to reliably quantify and monitor the fairness of recourse recommendations, along with practical methods for improving these metrics. Existing research has primarily examined

¹For example, independent audits have shown fairness concerns and limited contestability options in nationwide-deployed systems, such as *COMPAS* for recidivism prediction in the United States (Angwin et al. 2022) and *VioGén* for gender-based violence risk assessment in Spain (Éticas Foundation 2022).

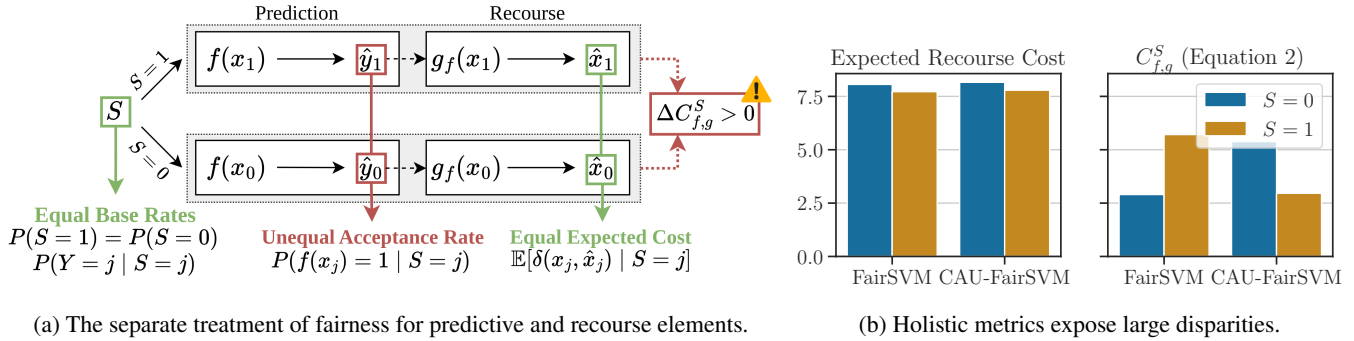


Figure 1: **Conventional metrics to measure fairness in recourse hide disparities.** (a) Given a population balanced by sensitive group membership and ground-truth outcomes (e.g., true ability to repay a loan), treating prediction and recourse fairness separately, as in the majority of prior work, can mask disparities: if positive classifications are unevenly distributed, one group bears recourse burdens more frequently even when per-instance recourse costs are equal; misclassification (e.g., false negatives) further amplifies this effect. (b) Conventional cost-parity metrics suggest balanced recourse costs across groups $S \in \{0, 1\}$ (left). Our metric (Eqn. 2) reveals substantial disparity once group differences in decision and error rates are taken into account, even with “fair” classifiers, FairSVM (Gupta et al. 2019) and CAU-FairSVM (Von Kügelgen et al. 2022) (right).

unfairness in recourse by analyzing the costs imposed on individuals receiving negative classifications (i.e., “reject”), either at the individual or group level (Von Kügelgen et al. 2022; Ehyaei et al. 2023; Kavouras et al. 2023; Bell et al. 2024; Yetukuri et al. 2024). However, we argue that fairness in recourse depends not only on *what* type of recourse is provided but also on *who* is more likely to receive the negative decision in the first place. As an illustration, consider two groups of people, $S = 0$ and $S = 1$, having the same number of individuals per group. If one of the groups (e.g., $S = 1$) has a higher acceptance rate than the other group, many more individuals from $S = 0$ receive a “reject” decision. Widely utilized fairness-in-recourse metric – *equal expected cost (of negatively predicted individuals)* – will not capture the inherent unfairness in the system as we can balance the cost of recourse per individual across two groups (e.g., $10/100 = 100/1000 = 0.1$). However, individuals in group $S = 0$ are disproportionately classified as negative (1000 individuals vs. 100) and therefore, a higher number of instances from $S = 0$ bear the cost of recourse.

We emphasize that the above illustrated problem is amplified when we consider the *ground-truth labels*, e.g., an individual’s latent (and unobserved) ability to repay a loan instead of just the classifier’s acceptance outcomes. Individuals may be wrongly classified and burdened with corrective actions, despite the error stemming from the decision support system itself. This misallocation of responsibility is troubling, particularly as error rates often differ systematically across sensitive groups, a well-documented challenge in algorithmic fairness (Mehrabi et al. 2021). All the issues mentioned above stem from the fact that current popular fairness-in-recourse paradigms do not correctly capture fairness issues in the whole system, which includes both the classifier’s components and the recourse’s components.

This paper presents a unified framework for assessing fairness in algorithmic recourse across the full decision pipeline, from initial classification through recourse recom-

mendation (see Fig. 1a for an overview). Our framework yields formal connections between *prediction fairness* and *recourse fairness*, showing how disparities in classification performance can propagate to recourse recommendations, and vice versa. Further, by leveraging these insights, we introduce a lightweight training method, MISOB (MINImax SOcial Burden), that empirically improves fairness throughout the pipeline, *without requiring access to sensitive attributes*, in real-world datasets. Lastly, we argue for moving beyond simple group-gap metrics and developing richer measures that more fully describe fairness guarantees.

Our contributions. Summarizing, we (i) theoretically characterize the sources of unfairness in algorithmic recourse, establishing formal connections between fairness guarantees in recourse and classification (Section 3), (ii) provide evidence of the limitations of approaches based on the standard equal cost paradigm (Section 3.1), (iii) introduce a novel framework for fair recourse (Algorithm 1) based on the concept of *social burden* (Section 4), and (iv) assess its performance on real-world datasets (Section 5).

2 Preliminaries and Related Work

Notation and Problem Formulation. We denote random variables with capital letters, their realizations with lowercase letters, and sets with capital letters in *italic*. Each decision subject (i.e., instance) is represented by the triplet (x, s, y) , where $x \in \mathcal{X} \subseteq \mathbb{R}^d$ represents the non-sensitive attributes, $s \in \mathcal{S} \subseteq \mathbb{R}^m$ the sensitive demographic attributes (e.g., race, gender, age), and $y \in \mathcal{Y}$ the class label. For simplicity, we assume a binary classification task with $\mathcal{Y} = \{0, 1\}$ (e.g., whether the decision subject is capable of repaying the loan). However, the proposed framework is general and can be extended to multiclass settings. The sensitive attributes s can be multidimensional, enabling the characterization of individuals by multiple protected attributes and the analysis of intersectional groups (e.g., *indigenous women*

	FAIRNESS METRIC			APPROACH AGNOSTIC TO		
	$f(x)$	$g_f(x)$	Y	$f(x)$	$g_f(x)$	S
GUPTA ET AL. (2019)	✗	✓	✗	✓	✗	✗
VON KÜGELGEN ET AL. (2022)	✗	✓	✗	✓	✗	✗
RAIMONDI, LAWRENCE, AND CHOCKLER (2022)	✓	✓	✗	-	-	-
KURATOMI ET AL. (2022)	✓	✓	✓	-	-	-
KAVOURAS ET AL. (2023)	✗	✓	✗	-	-	-
BELL ET AL. (2024)	✗	✓	✗	✓	✓	✗
THIS WORK (SECTION 3 AND 4)	✓	✓	✓	✓	✓	✓

Table 1: Summary of the existing works on fairness in recourse. We report the proposed fairness metric, whether it considers the classifier ($f(x)$) the recourse method ($g_f(x)$) or the ground truth outcomes (Y), and the approach used to enhance fairness; specifically, whether it is agnostic to the classifier, the recourse method, or the sensitive attributes (S). A “-” indicates works that propose only evaluation metrics without introducing any algorithm or methodology improving fairness.

under 30) that reflect more nuanced and compounded socio-economic disadvantages. Further, let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a binary classifier whose goal is to assign a label in a set \mathcal{Y} to instances in a set \mathcal{X} . Let $\delta((x, s), x')$ denote the cost for an individual with features (x, s) to change its non-sensitive profile to x' , e.g., ℓ_1 -norm (Karimi et al. 2022). Given an instance that received a negative classification (i.e., $f(x) = 0$), we want to find the “closest” instance x' achieving a positive classification by solving the following optimization problem:

$$\operatorname{argmin}_{x' \in \mathcal{X}} \delta((x, s), x') \quad \text{s.t. } f(x') \neq f(x) \quad (1)$$

In the literature, x' is also called *counterfactual explanation* (Wachter, Mittelstadt, and Russell 2017). Here, we assume the recommendation $x' = g_f(x)$ is generated by a recourse algorithm $g_f : \mathcal{X} \rightarrow \mathcal{X}$ tailored to the classifier f . A growing body of work has proposed diverse strategies $g_f(x)$ to offer recourse to individuals assigned unfavorable predictions. These methods vary in several key aspects, including the type of base classifier considered, and actionability constraints. Our discussion here is clearly limited; thus, please refer to Karimi et al. (2022) for a comprehensive survey on the topic. While providing recommendations is an important step toward accountability and transparency, significant fairness challenges remain, some stemming directly from the recourse process. Indeed, recourse can exacerbate social segregation (Gao and Lakkaraju 2023), and its cost may vary with user preferences (De Toni, Lepri, and Passerini 2023; De Toni et al. 2024) and sensitive attributes (Yetukuri et al. 2024), potentially deepening existing disparities.

Measuring Algorithmic Fairness. Research on algorithmic fairness has primarily examined disparities in model predictions across population subgroups, focusing on the *fairness of predictions themselves*. Fairness definitions typically fall into two categories: *individual-level*, which requires similar individuals to receive similar outcomes, and *group-level*, which evaluates whether performance metrics (e.g., accuracy) are comparable across demographic groups. Group metrics are widely used due to their computational simplicity, ease of integration into machine-learning pipelines, and ability to reveal structural disparities. Two common examples are *statistical parity* (Dwork et al. 2012), which measures differences in the probability of a positive prediction (i.e., acceptance rate) across groups regardless of

true labels, and *equality of opportunity* (Heidari et al. 2019), which measures such differences only among individuals in the positive class (i.e., true positive rate). However, *these metrics focus solely on the prediction stage and overlook the downstream consequences of model decisions*.

Measuring and Enhancing Fairness in Recourse. Fairness concerns in algorithmic recourse prompted the search for specific fairness criteria (Ustun, Spangher, and Liu 2019; Gupta et al. 2019). Table 1 summarizes the characteristics of the most relevant approaches that address fairness issues in recourse. Many works proposed both group- and individual-level fairness metrics in this direction (Von Kügelgen et al. 2022; Kavouras et al. 2023; Bell et al. 2024) and extended the analysis to incorporate classifier predictions and standard predictive performance measures (Kuratomi et al. 2022; Raimondi, Lawrence, and Chockler 2022). Several papers that introduced fairness metrics for evaluating algorithmic recourse have also proposed methods to improve fairness according to these metrics (Gupta et al. 2019; Von Kügelgen et al. 2022; Bell et al. 2024). For instance, Von Kügelgen et al. (2022) proposed a causality-based method that requires access to a structural causal model (SCM). Gupta et al. (2019) introduced an approach applicable only to recourse methods that map instances directly to the decision boundary, while Bell et al. (2024) addressed fairness under resource constraints, a setting closer to a ranking rather than a classification problem.

Most fairness metrics proposed in this context focus *exclusively on the recourse stage*, considering only negatively classified instances and ignoring the classifier’s overall behavior, such as group-specific acceptance and rejection rates. Furthermore, they generally lack *practical strategies* for improving fairness. Lastly, none of these studies adopt an *intersectional perspective*, and all assume access to sensitive attributes for all individuals *at training time*, an assumption rarely satisfied in real-world applications. Our work is conceptually aligned with Kuratomi et al. (2022) and Raimondi, Lawrence, and Chockler (2022), but their formulation lacks a clear connection to existing fairness metrics in prediction.

3 Revisiting (Un)Fairness in Recourse

We now outline three key limitations of common fairness metrics in pipelines with recourse: (1) they ignore the clas-

sifier’s decision behavior, (2) they omit ground truth information, and (3) they focus solely on metric equalization. We propose an alternative formulation that addresses these issues, offers a broader view of unfairness sources in recourse, and establishes explicit links between fairness in prediction and recourse. Proofs are provided in Appendix A.

3.1 The Impact of the Classifier’s Decisions

Let us exemplify the first issue with an example of a binary classification task. Consider a population X with a binary target variable $Y \in \{0, 1\}$ and two equally represented sensitive groups $S \in \{0, 1\}$, such that $P(S = 0) = P(S = 1)$ and $P(Y = 1 | S = 0) = P(Y = 1 | S = 1)$. Let $f(x)$ be a classifier approximating the posterior $P(Y | X)$, and assume that under f one group has a higher *acceptance rate* (AR), $P(f(x) = 1 | S = 1) > P(f(x) = 1 | S = 0)$. After recourse is provided to negatively classified instances, let us assume that the expected cost of the required changes is equal across groups $\mathbb{E}[\delta((x, s), g_f(x)) | S = 1] = \mathbb{E}[\delta((x, s), g_f(x)) | S = 0]$. According to standard fairness metrics, this situation would be considered fair. *But is it truly fair?* While the expected recourse cost is equal once offered, individuals in the group with lower AR are more likely to require recourse in the first place. We argue that *this disparity reflects an unfair situation, which existing popular metrics fail to capture*. We now introduce an alternative fairness metric that evaluates recourse costs across the entire population, rather than only among negatively classified instances.

Definition 3.1. Let f be a classifier and g_f a recourse algorithm. Given a sensitive group $s \in \mathcal{S}$, let $\delta((x, S = s), g_f(x))$ be the cost of applying the recourse provided by g_f . Then, the expected recourse cost $C_{f,g}^s$ can be defined as:

$$\underbrace{\mathbb{E}[\delta((X, S = s), g_f(X))]}_{\text{Expected cost for instances with } f(x)=0} \left(1 - \underbrace{P(f(X) = 1 | S = s)}_{\text{Acceptance Rate (AR)}}\right) \quad (2)$$

where the expectation is over $P(X | S = s, f(X) = 0)$, and $P(f(X) = 1 | S = s)$ indicates the *acceptance rate* of the classifier for the given sensitive group.

Namely, Eqn. 2 arises since instances receiving a positive prediction, $\{x : f(x) = 1, x \in X\}$, do not modify their non-sensitive features, implying no changes in recourse cost, $\delta((x, s), g_f(x)) = 0$. Fig. 1b shows how, when accounting for AR across sensitive groups, disparities in expected recourse costs become substantially larger. Given the CAULIN dataset from Von Kügelgen et al. (2022), we computed the empirical expected recourse cost by simply averaging only over negatively predicted instances (left), and by using Eqn. 2 (right), which also accounts for the AR of the base classifiers. Indeed, Fig. 1b shows that both groups have similar costs under the original metric with FairSVM, but $S = 1$ has the highest under the revised one, as a significantly larger proportion of this group receives negative predictions. Therefore, to avoid masking underlying unfairness, we argue that the cost metric must account for *acceptance rates* to accurately capture true recourse costs across groups.

3.2 The Burden of False Negatives

In Eqn. 2, we measure the expected recourse cost while accounting for the classifier’s predictions, providing a more precise notion of fairness. However, it focuses exclusively on the model’s predictions $f(x)$, without accounting for the true class label. As a result, it fails to distinguish between individuals who were incorrectly denied a positive outcome (*i.e.*, truly positive but misclassified) and those who were correctly assigned a negative outcome. In this section, we take a step further by focusing on minimizing the expected cost associated with positive instances that the classifier incorrectly predicts as negative. This next step naturally leads to the concept of social burden.² *Social burden* measures the expected cost incurred by individuals who are already qualified (true positives) but receive an unfavorable prediction by a classifier f (Milli et al. 2019). The social burden is zero if the classifier positively classifies all qualified individuals without requiring changes to their non-sensitive attributes. Conversely, the social burden is greater than zero when qualified individuals are assigned a negative classification and forced to alter their features to be recognized as valid. Formally, we can define the social burden as:

Definition 3.2. Let f be a classifier and g_f a recourse algorithm. Given a sensitive group $s \in \mathcal{S}$, let $\delta((x, S = s), g_f(x))$ be the cost of applying the recourse provided by g_f . Then, the expected social burden $B_{f,g}^s$ can be defined as:

$$\underbrace{\mathbb{E}[\delta((X, s), g_f(X))]}_{\text{Exp. cost for inst. with } y=1, f(x)=0} \underbrace{(1 - P(f(X) = 1 | S = s, Y = 1))}_{\text{True Positive Rate (TPR)}} \quad (3)$$

with the expectation over $P(X | S = s, Y = 1, f(X) = 0)$, and $P(f(X) = 1 | S = s, Y = 1)$ indicates the *true positive rate* (TPR) for the given sensitive group.

Similarly to Eqn. 2, the instances obtaining a positive prediction do not count towards the overall burden. Eqn. 3 underscores the importance of considering the predictive performance of the base classifier when evaluating algorithmic recourse, a factor often overlooked in the recourse literature but central to algorithmic fairness; especially, since numerous studies have shown that machine learning classifiers tend to exhibit unequal error rates (*e.g.*, TPR) across sensitive groups (Pessach and Shmueli 2023). When error rates differ across groups, the burden of recourse is also likely to be unevenly distributed, potentially amplifying existing disparities. However, it is important to note that satisfying *equality of opportunity* (*i.e.*, parity in TPR) does not guarantee a small gap in social burden.

3.3 Gap-Based Metrics can Mask Unfairness

Unfairness in recourse can be quantified by measuring *gaps* in cost or social burden across sensitive groups: $\Delta C_{f,g} = \max_{s,s' \in \mathcal{S}} |C_{f,g}^s - C_{f,g}^{s'}|$ or $\Delta B_{f,g} = \max_{s,s' \in \mathcal{S}} |B_{f,g}^s - B_{f,g}^{s'}|$.

²Some works on algorithmic recourse (*e.g.*, Kavouras et al. (2023)) employ the term *burden* to refer to the average cost itself. However, in an earlier work by Milli et al. (2019), *burden* is defined as the average cost incurred by instances that have a true positive class label. In this paper, we adopt the latter definition.

$B_{f,g}^{s'}$. According to Equations 2 and 3, these disparities stem from three main sources. For cost gaps, the causes include: (a) disparities in AR, as captured by statistical parity (Dwork et al. 2012); (b) differences in the cost function, which may arise from unequal opportunities or socioeconomic conditions; and (c) variations in the distribution $P(X | S, f(X) = 0)$, reflecting potential population skews. For social burden gaps, they arise from: (a) disparities in true positive rates, captured by equality of opportunity (Hardt, Price, and Srebro 2016), (b) unequal cost functions, and (c) skews in the distribution of qualified individuals denied recourse $P(X | S, Y = 1, f(X) = 0)$. Importantly, this illustrates that satisfying equality of opportunity (statistical parity) does not guarantee an equal distribution of the social burden (recourse costs) across groups.

Although comparing the *gap of statistical measures* across groups is a common approach in the fairness literature, it can obscure meaningful underlying unfairness. In fairness for prediction, a well-known limitation of parity-based definitions is that parity can be achieved by *degrading* performance for the privileged group without improving outcomes for the unprivileged one (Martinez, Bertran, and Sapiro 2020; Diana et al. 2021). As a result, groups may exhibit similar expected values for a given metric, even when those values are far from optimal. In recourse, this issue is especially critical when evaluating social burden, which ideally should be as low as possible (preferably zero). A small gap in burden does not necessarily indicate fairness, as parity might be reached by *increasing the burden for the privileged group* rather than reducing it for the unprivileged one. However, in many real-world contexts, such as access to essential goods and services, sacrificing overall well-being is not acceptable. Therefore, fairness efforts should aim to reduce the burden for all groups, rather than simply equalizing it.

We propose drawing inspiration from approaches in the fairness-in-prediction literature that address this limitation and complement difference-based disparity measures with a Rawlsian (*i.e.*, minimax) perspective (Binns 2018). For instance, in addition to evaluating the gap in social burden, we can examine the *worst-group social burden* defined as $\max_{s \in \mathcal{S}} B_{f,g}^s$. This perspective’s benefits are twofold: (a) it helps identify the most vulnerable group, and (b) it serves as a base for the development of fairness-enhancing recourse frameworks that go beyond minimizing burden disparities, aiming instead to reduce the social burden for all groups.

4 MISOB: MInimax SOcial Burden

In this section, we introduce a simple yet effective iterative training procedure, MISOB (MInimax SOcial Burden), that empirically reduces social burden (Eqn. 3) without degrading predictive performance and *without requiring access to sensitive group attributes during training or inference*. The method is outlined in Algorithm 1. Given a pre-trained classifier (line 1)³, the key idea is to incorporate social burden estimates directly into the training loop. At each iteration,

³Empirically, we observed that pretraining a base classifier to achieve reasonable accuracy before applying MISOB contributes to a more stable optimization process.

Algorithm 1: MISOB: MInimax SOcial Burden

Require: $\mathcal{D} = \{(x^i, y^i)\}_{i=1}^N$, $T \in \mathbb{N}$, and $\alpha \in \mathbb{R}^+$

- 1: Pre-train $f^{(0)}$ using \mathcal{D} //Warm-up Phase
- 2: **for** $t = 1$ to T **do**
- 3: $\mathcal{Q} = \{\}$
- 4: **for** $i = 0$ to N **do**
- 5: $b_{f,g}^i \leftarrow \delta(x^i, g_f(x^i)) \mathbb{1}\{y^i = 1\}$ //Compute burden
- 6: $\mathcal{Q} \leftarrow \mathcal{Q} \cup b_{f,g}^i$
- 7: **end for**
- 8: $f^{(t)} \leftarrow \arg \min_{f \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^N \phi(i, \mathcal{Q}, \alpha) \cdot \ell(f(x^i), y^i)$
- 9: **end for**
- 10: **return** $f^{(T)}$

we compute the social burden for each instance (line 5), and we save the result within \mathcal{Q} (line 6). The model is then updated by minimizing a (*weighted*) *loss function* $\ell(f(x^i), y^i)$ (line 8), where weights are determined by the estimated burden $\phi(i, \mathcal{Q})$. Specifically, we propose assigning weights to instances x^i based on the following scheme:

$$\phi(i, \mathcal{Q}, \alpha) = 1 + \alpha N \frac{b_{f,g}^i}{\sum_{j \in |\mathcal{Q}|} b_{f,g}^j} \mathbb{1}\{\beta > 0\} \quad (4)$$

where N is the dataset (or batch) size, $\beta = \sum_{j \in |\mathcal{Q}|} b_{f,g}^j$ is the total burden with $b_{f,g}^i = \delta(x^i, g_f(x^i)) \mathbb{1}\{y^i = 1\}$ being the burden of instance i , and α is a hyperparameter controlling the influence of the burden term. Under this weighting scheme, each instance x^i is assigned a weight proportional to its contribution to the total burden across the dataset (or batch). In the degenerate case, where the total burden is zero, all instances are assigned equal weights. This reweighting scheme amplifies the influence of *high-burden instances* during training, encouraging the classifier to prioritize decisions that reduce unnecessary or costly adjustments imposed on individuals under the current decision rule.

The proposed framework offers several advantages that enhance its practicality and broad applicability.⁴ First, the method is *agnostic* to both the base classifier and the chosen recourse strategy, allowing seamless integration with a wide range of existing approaches. It also operates *without access to sensitive attributes during training*, avoiding legal, ethical, and practical challenges associated with collecting or using such data. Further, MISOB not only improves deployability but also naturally addresses intersectional fairness. Because group definitions are not required a-priori, intersectional subgroups can be introduced post-hoc at *evaluation time*, enabling fairness assessments across multiple dimensions *without retraining*. This flexibility makes the frame-

⁴The runtime of Algorithm 1 is determined by three factors: (i) the number of retraining steps T , (ii) the dataset size N , and (iii) the run-time complexity of the recourse method $O(K)$. In practice, datasets are typically much larger than the number of retraining iterations, and if a simple recourse method is used, we generally have $K \ll T \ll N$. As a result, the overall computational complexity is $O(N^3)$. Efficiency can be further improved through batching, parallelization, or approximating recourse computations.

SENSITIVE ATTRIBUTE(S)	RECOURSE METHOD	FAIRNESS STRATEGY	ACC OVERALL \uparrow	BURDEN (EQ. 3)		TPR		COST (EQ. 2)		AR	
				WORST \downarrow	$\Delta \downarrow \downarrow$	WORST \uparrow	$\Delta \downarrow \downarrow$	WORST \downarrow	$\Delta \downarrow \downarrow$	WORST \uparrow	$\Delta \downarrow \downarrow$
RACE	GS	-	0.81 \pm 0.02	4.56 \pm 0.01	0.03 \pm 0.02	0.27 \pm 0.03	0.08 \pm 0.08	115.69 \pm 1.19	28.37 \pm 1.95	0.06 \pm 0.01	0.06 \pm 0.02
	GS	POSTPRO	0.80 \pm 0.01	4.96 \pm 3.49	0.61 \pm 2.78	0.37 \pm 0.51	0.00 \pm 0.52	98.40 \pm 56.73	17.19 \pm 42.75	0.37 \pm 0.53	0.01 \pm 0.54
	GS	MISOB	0.82 \pm 0.01	3.01 \pm 0.32	0.85 \pm 0.46	0.52 \pm 0.02	0.11 \pm 0.04	93.06 \pm 0.98	27.38 \pm 1.63	0.15 \pm 0.01	0.12 \pm 0.02
	WT	-	0.81 \pm 0.02	1.28 \pm 0.14	0.01 \pm 0.18	0.27 \pm 0.03	0.08 \pm 0.08	38.27 \pm 0.57	12.32 \pm 0.72	0.06 \pm 0.01	0.06 \pm 0.02
	WT	POSTPRO	0.80 \pm 0.01	1.55 \pm 0.27	0.01 \pm 0.47	0.37 \pm 0.51	0.00 \pm 0.52	39.71 \pm 5.03	13.22 \pm 2.10	0.37 \pm 0.53	0.01 \pm 0.54
	WT	MISOB	0.82 \pm 0.01	0.79 \pm 0.00	0.16 \pm 0.00	0.59 \pm 0.04	0.02 \pm 0.06	30.77 \pm 0.23	11.66 \pm 0.45	0.16 \pm 0.02	0.08 \pm 0.04
	CCHVAE	-	0.81 \pm 0.02	6.25 \pm 0.46	0.16 \pm 0.89	0.27 \pm 0.03	0.08 \pm 0.08	119.99 \pm 3.78	27.01 \pm 7.52	0.06 \pm 0.01	0.06 \pm 0.02
	CCHVAE	POSTPRO	0.80 \pm 0.01	11.20 \pm 1.89	3.06 \pm 2.96	0.37 \pm 0.51	0.00 \pm 0.52	120.01 \pm 8.94	21.22 \pm 9.38	0.37 \pm 0.53	0.01 \pm 0.54
	CCHVAE	MISOB	0.81 \pm 0.01	4.03 \pm 0.44	0.42 \pm 0.54	0.48 \pm 0.03	0.19 \pm 0.05	105.10 \pm 1.14	27.63 \pm 1.98	0.12 \pm 0.01	0.16 \pm 0.02
GENDER	GS	-	0.81 \pm 0.02	4.73 \pm 0.39	0.54 \pm 0.58	0.32 \pm 0.05	0.08 \pm 0.04	109.75 \pm 2.96	28.76 \pm 3.72	0.06 \pm 0.03	0.07 \pm 0.04
	GS	POSTPRO	0.80 \pm 0.01	24.46 \pm 25.08	13.98 \pm 8.03	0.00 \pm 0.00	0.00 \pm 0.00	145.47 \pm 47.36	28.17 \pm 63.66	0.00 \pm 0.00	0.00 \pm 0.00
	GS	MISOB	0.82 \pm 0.01	2.59 \pm 0.00	0.33 \pm 0.00	0.55 \pm 0.01	0.04 \pm 0.02	94.26 \pm 0.71	32.49 \pm 1.13	0.10 \pm 0.01	0.16 \pm 0.01
	WT	-	0.81 \pm 0.02	1.29 \pm 0.03	0.04 \pm 0.05	0.32 \pm 0.05	0.08 \pm 0.08	43.06 \pm 0.37	22.88 \pm 0.55	0.06 \pm 0.03	0.07 \pm 0.04
	WT	POSTPRO	0.80 \pm 0.01	1.65 \pm 0.39	0.33 \pm 0.22	0.00 \pm 0.00	0.00 \pm 0.00	42.40 \pm 3.59	20.99 \pm 2.12	0.00 \pm 0.00	0.00 \pm 0.00
	WT	MISOB	0.82 \pm 0.01	0.79 \pm 0.06	0.21 \pm 0.08	0.55 \pm 0.03	0.04 \pm 0.04	34.36 \pm 0.31	20.29 \pm 0.50	0.10 \pm 0.03	0.17 \pm 0.04
	CCHVAE	-	0.81 \pm 0.02	6.43 \pm 0.64	0.57 \pm 1.02	0.32 \pm 0.05	0.08 \pm 0.08	116.69 \pm 4.83	31.12 \pm 6.87	0.06 \pm 0.03	0.07 \pm 0.04
	CCHVAE	POSTPRO	0.80 \pm 0.01	12.97 \pm 4.20	6.61 \pm 4.78	0.00 \pm 0.00	0.00 \pm 0.00	119.88 \pm 7.24	27.32 \pm 12.12	0.00 \pm 0.00	0.00 \pm 0.00
	CCHVAE	MISOB	0.81 \pm 0.01	4.18 \pm 0.27	0.77 \pm 0.43	0.55 \pm 0.03	0.12 \pm 0.05	103.98 \pm 1.52	35.58 \pm 2.69	0.11 \pm 0.02	0.22 \pm 0.03
RACE & GENDER	GS	-	0.81 \pm 0.02	5.13 \pm 0.66	1.79 \pm 1.22	0.20 \pm 0.04	0.22 \pm 0.07	124.86 \pm 2.33	48.63 \pm 3.52	0.02 \pm 0.01	0.11 \pm 0.02
	GS	POSTPRO	0.80 \pm 0.01	5.12 \pm 2.26	3.13 \pm 2.81	0.00 \pm 0.00	0.00 \pm 0.00	128.05 \pm 64.21	91.55 \pm 42.15	0.00 \pm 0.00	0.00 \pm 0.00
	GS	MISOB	0.82 \pm 0.01	3.37 \pm 0.70	1.32 \pm 1.00	0.46 \pm 0.05	0.19 \pm 0.07	109.95 \pm 2.35	51.28 \pm 3.33	0.06 \pm 0.01	0.23 \pm 0.02
	WT	-	0.81 \pm 0.02	1.40 \pm 0.36	0.27 \pm 0.54	0.20 \pm 0.04	0.22 \pm 0.07	49.71 \pm 1.38	30.81 \pm 2.40	0.02 \pm 0.01	0.11 \pm 0.02
	WT	POSTPRO	0.80 \pm 0.01	1.94 \pm 0.75	0.90 \pm 0.48	0.00 \pm 0.00	0.00 \pm 0.00	51.48 \pm 8.83	30.81 \pm 4.33	0.00 \pm 0.00	0.00 \pm 0.00
	WT	MISOB	0.82 \pm 0.01	0.98 \pm 0.08	0.45 \pm 0.12	0.34 \pm 0.05	0.25 \pm 0.08	41.43 \pm 0.66	28.59 \pm 0.93	0.04 \pm 0.03	0.25 \pm 0.04
	CCHVAE	-	0.81 \pm 0.02	7.37 \pm 0.51	3.53 \pm 1.01	0.20 \pm 0.04	0.22 \pm 0.07	133.32 \pm 1.58	53.24 \pm 3.16	0.02 \pm 0.01	0.11 \pm 0.02
	CCHVAE	POSTPRO	0.80 \pm 0.01	18.25 \pm 1.47	13.07 \pm 3.64	0.00 \pm 0.00	0.00 \pm 0.00	149.20 \pm 10.24	66.95 \pm 10.89	0.00 \pm 0.00	0.00 \pm 0.00
	CCHVAE	MISOB	0.81 \pm 0.01	5.36 \pm 0.57	2.11 \pm 1.09	0.35 \pm 0.05	0.33 \pm 0.08	116.93 \pm 0.92	52.78 \pm 1.67	0.04 \pm 0.02	0.31 \pm 0.03

Table 2: Empirical average test results on the ADULT dataset across different sensitive group characterizations. The best results for each setting are highlighted in bold, and we report whether the target metric should increase (\uparrow) or decrease (\downarrow). For MISOB, we set $\alpha = 0.3$ in all experiments (*cf.* Eqn. 4). Each cell shows the average and standard deviation over 10 runs.

work well-suited for real-world applications where fairness concerns span diverse and evolving group identities.

5 Empirical Evaluation

In this section, we evaluate the effectiveness of MISOB in reducing the social burden across sensitive groups. Refer to Appendix B for further details on the experimental setting.

Experimental Setting. We consider the real-world dataset ADULT (Becker and Kohavi 1996) with sensitive groups defined by different characterizations of the sensitive attributes (*race* and *gender*), considering also intersectional groups. For the model, we consider a simple feed-forward neural network. We have chosen this base classifier because it supports multiple different recourse methods proposed in the literature. Further experiments with a linear regression model and other datasets can be found in Appendix C. In our evaluation, we consider MISOB and POSTPRO (Hardt, Price, and Srebro 2016), which enhance prediction fairness in terms of *equality of opportunity*. We employ POSTPRO because it is compatible with any base classifier, enabling fair comparisons with MISOB under identical predictive models. We consider three widely adopted recourse methods, selected for their benchmarking popularity, available implementations, and ease of integration in an agnostic setting: Growing Spheres (GS) (Laugel et al. 2017), the method by Wachter, Mittelstadt, and Russell (2017) (WT), and the Counterfactual Conditional Heterogeneous Variational Autoencoder (CCHVAE) by Pawelczyk, Broelemann, and Kas-

neci (2020). We use the implementations and hyperparameters from the benchmark by Pawelczyk et al. (2021), and we repeat the experiments over 10 random train/test splits. We consider the worst-group values for various metrics:

$$\min_{s \in S} P(f(X = x) = 1 \mid S = s, Y = 1) \text{ (Worst TPR)} \quad (5)$$

$$\min_{s \in S} P(f(X = x) = 1 \mid S = s) \text{ (Worst AR)} \quad (6)$$

$$\min_{s \in S} C_{f,g}^s \text{ (Worst Cost)} \quad \min_{s \in S} B_{f,g}^s \text{ (Worst Burden)} \quad (7)$$

where the recourse cost δ is quantified using the ℓ_2 distance, and their maximum disparity (denoted by Δ) across groups:

$$\max_{s,s' \in S} |C_{f,g}^s - C_{f,g}^{s'}| \quad \max_{s,s' \in S} |B_{f,g}^s - B_{f,g}^{s'}| \quad (8)$$

$$\max_{s,s' \in S} |\text{TPR}^s - \text{TPR}^{s'}| \quad \max_{s,s' \in S} |\text{AR}^s - \text{AR}^{s'}| \quad (9)$$

Importantly, our method operates *entirely without access to sensitive attributes*. Note that computing fairness metrics—whether for prediction or recourse—requires access to sensitive attributes, which is an inherent limitation of these metrics themselves (Buyl and De Bie 2024).

Results. Table 2 shows that training the base classifier under the MISOB framework improves both predictive and recourse-related performance metrics across all groups, without compromising overall accuracy. Notably, MISOB enhances *worst-case group outcomes* and, in many cases, reduces disparities across groups, achieving fairness improvements without degrading the performance for privi-

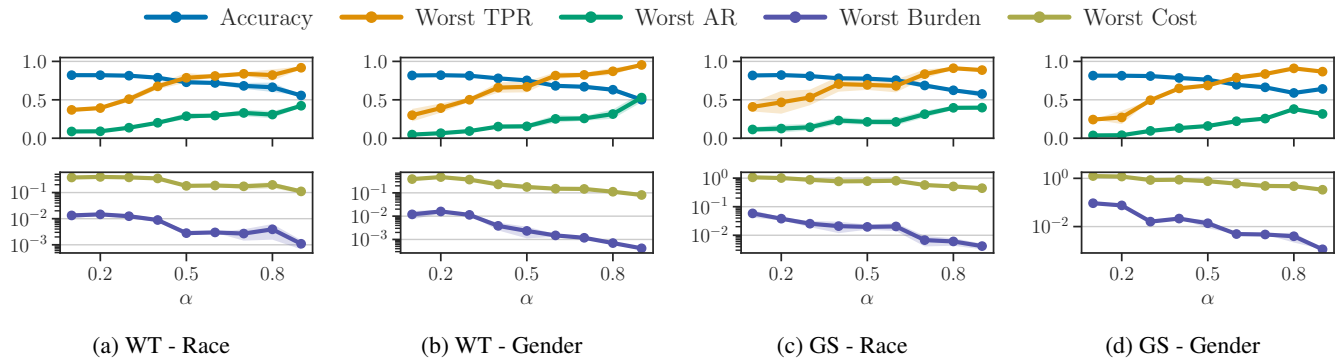


Figure 2: **Empirical evaluation of MISOB for different $\alpha \in \{0.1, \dots, 1.0\}$.** We report the average results and standard deviation (shaded area) over 10 runs, with WT (Wachter, Mittelstadt, and Russell 2017) and GS (Laugel et al. 2017), on the ADULT dataset with *race* and *gender* as the sensitive attributes. The worst burden and cost (Eqn. 7) are in log scale. In brief, increasing α , which favors high-burden instances, improves fairness guarantees, but large α values may hurt overall accuracy.

leged groups. These results also highlight an important limitation of gap-based fairness metrics: small gaps can mask universally poor outcomes across groups, giving a misleading sense of fairness. In contrast, MISOB can produce configurations with slightly larger gaps but better metric values for every group, thereby promoting greater overall well-being. Moreover, the results show that improving fairness at the prediction level does not guarantee fairer recourse outcomes and can even exacerbate them. In particular, enforcing parity in TPR through POSTPRO significantly increases the burden and cost for decision subjects. This occurs when prediction-level fairness improvements are misaligned with fairness improvements at recourse. On the contrary, the holistic perspective of MISOB enables effective improvements in group TPR, while reducing the cost and burden of individuals, underscoring the importance of considering the entire pipeline. Besides, note that POSTPRO requires training a *separate model* for each sensitive group characterization, whereas MISOB, being blind to sensitive information, needs to be trained only once and can be evaluated across multiple group characterizations.

The performance and fairness of MISOB are governed by the emphasis placed on high-burden individuals, controlled by the weighting parameter α from Eqn. 4. Fig. 2 shows that increasing α (*i.e.*, the importance of high-burden instances) strengthens minimax fairness across both prediction and recourse by improving worst-group metrics. However, larger values of α ($\alpha > 0.5$) can reduce overall accuracy due to the overemphasis on high-burden cases. In practice, α governs the trade-off between fairness and accuracy, allowing users to select the value that best aligns with their objectives.

6 Discussion and Limitations

Our work introduces a novel framework to compute fairness metrics in a recourse-aware setting and empirically validates the effectiveness of MISOB. Analyzing the convergence properties of the training procedure (*e.g.*, identifying stationary points) would be a valuable direction for further research. One approach might be to frame recourse

fairness as a two-player Stackelberg game, following the *strategic classification* perspective (Chen, Wang, and Liu 2023). Additionally, applying calibration techniques (Guo et al. 2017) could improve prediction reliability (*cf.* Eqn. 3), which in turn may further enhance fairness when combined with MISOB. Finally, further fairness concerns may emerge due to model updates, dataset shifts (Castelnovo et al. 2021), prediction performativity (Liu et al. 2018; Perdomo et al. 2020), or recourse recommendations that are not robust over time (De Toni et al. 2025). Extending our approach to measure and optimize the social burden of recourse in dynamic settings is, therefore, a natural path for future research. Lastly, we aim to explore more realistic cost models that better reflect the real world’s complexities and assess how these influence fairness guarantees (Tominaga, Yamashita, and Kurashima 2024; Esfahani et al. 2024).

7 Conclusions

In this work, we study the unfairness that can arise in automated decision-making pipelines that provide recourse. We provide a theoretical characterization of the sources of unfairness in algorithmic recourse, *formally linking fairness guarantees in recourse and prediction*, and further highlight the limitations of the standard equal cost paradigm. To address these challenges, we introduce MISOB, a novel framework grounded in the notion of *social burden*, which enables a holistic treatment of fairness and goes beyond a gap-based characterization of fairness guarantees. MISOB is broadly applicable: it is compatible with many classifier-recourse pairs and does not require access to sensitive attributes during training or inference. Empirical results show that improving fairness solely at the prediction level is often insufficient and can even considerably increase the cost and burden of decision subjects. In contrast, our approach leads to consistent joint fairness improvements in prediction and recourse, without compromising overall performance.

Acknowledgments

We would like to thank all the reviewers for their insightful feedback. We also thank Mikel Malagon for the technical support and for reading preliminary versions of the paper. This research was funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HaDEA). Neither the European Union nor the granting authority can be held responsible for them. This work is supported by the European Research Council under the European Union’s Horizon 2020 research and innovation programme Grant Agreement no. 851538 - BayesianGDPR, Horizon Europe research and innovation programme Grant Agreement no. 101120763 - TANGO, and no. 101120237 - ELIAS. This work is also supported by the Basque Government under grant IT1504-22 and through the BERC 2022-2025 program; by the Spanish Ministry of Science and Innovation under the grant PID2022-137442NB-I00, and through BCAM Severo Ochoa accreditation CEX2021-001142-S / MICIN / AEI / 10.13039/501100011033. Lastly, this work was also supported by Ministero delle Imprese e del Made in Italy (IPCEI Cloud DM 27 giugno 2022 – IPCEI-CL-0000007), funded by the European Commission under the NextGeneration EU programme.

References

- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2022. Machine bias. In *Ethics of data and analytics*, 254–264. Auerbach Publications.
- Becker, B.; and Kohavi, R. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- Bell, A.; Fonseca, J.; Abrate, C.; Bonchi, F.; and Stoyanovich, J. 2024. Fairness in Algorithmic Recourse Through the Lens of Substantive Equality of Opportunity. *arXiv preprint arXiv:2401.16088*.
- Binns, R. 2018. Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency*, 149–159. PMLR.
- Buyl, M.; and De Bie, T. 2024. Inherent limitations of AI fairness. *Communications of the ACM*, 67(2): 48–55.
- Castelnovo, A.; Malandri, L.; Mercorio, F.; Mezzanzanica, M.; and Cosentini, A. 2021. Towards fairness through time. In *Joint European conference on machine learning and knowledge discovery in databases*, 647–663. Springer.
- Chen, Y.; Wang, J.; and Liu, Y. 2023. Learning to incentivize improvements from strategic agents. *Transactions on Machine Learning Research*.
- Chouldechova, A.; Benavides-Prado, D.; Fialko, O.; and Vaithianathan, R. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on fairness, accountability and transparency*, 134–148. PMLR.
- De Toni, G.; Lepri, B.; and Passerini, A. 2023. Synthesizing explainable counterfactual policies for algorithmic recourse with program synthesis. *Mach. Learn.*, 112(4): 1389–1409.
- De Toni, G.; Teso, S.; Lepri, B.; and Passerini, A. 2025. Time Can Invalidate Algorithmic Recourse. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency, FAccT 2025, Athens, Greece, June 23-26, 2025*, 89–107. ACM.
- De Toni, G.; Viappiani, P.; Teso, S.; Lepri, B.; and Passerini, A. 2024. Personalized Algorithmic Recourse with Preference Elicitation. *Trans. Mach. Learn. Res.*, 2024.
- Diana, E.; Gill, W.; Kearns, M.; Kenthapadi, K.; and Roth, A. 2021. Minimax group fairness: Algorithms and experiments. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 66–76.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Ehyaei, A.-R.; Karimi, A.-H.; Schölkopf, B.; and Maghsudi, S. 2023. Robustness implies fairness in causal algorithmic recourse. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 984–1001.
- Esfahani, S.; De Toni, G.; Lepri, B.; Passerini, A.; Tentori, K.; and Zancanaro, M. 2024. Preference elicitation in interactive and user-centered algorithmic recourse: an initial exploration. In *Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*, 249–254.
- Eticas Foundation. 2022. Can AI solve gender violence? Auditing the use of AI to assess risk. The case of VioGén. Technical report, Eticas Foundation, Barcelona, Spain. Adversarial Audit report. Project Lead: Gemma Galdon-Clavell; Researchers: Evren Yalaz, Toni Lorente. Produced in collaboration with Fundació Ana Bella.
- EU Commission. 2021. Proposal for a regulation of the European Parliament and the Council laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. *EUR-Lex-52021PC0206*.
- Gao, R.; and Lakkaraju, H. 2023. On the impact of algorithmic recourse on social segregation. In *International Conference on Machine Learning*, 10727–10743. PMLR.
- Grgic-Hlaca, N.; Redmiles, E. M.; Gummadi, K. P.; and Weller, A. 2018. Human Perceptions of Fairness in Algorithmic Decision Making: A Case Study of Criminal Risk Prediction. In *Proceedings of the 2018 World Wide Web Conference, WWW ’18*, 903–912. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee. ISBN 9781450356398.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330. PMLR.
- Gupta, V.; Nokhiz, P.; Roy, C. D.; and Venkatasubramanian, S. 2019. Equalizing recourse across groups. *arXiv preprint arXiv:1909.03166*.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.

- Heidari, H.; Loi, M.; Gummadi, K. P.; and Krause, A. 2019. A moral framework for understanding fair ml through economic models of equality of opportunity. In *Proceedings of the conference on fairness, accountability, and transparency*, 181–190.
- Heikkilä, M. 2022. Dutch scandal serves as a warning for Europe over risks of using algorithms. <https://www.politico.eu/article/dutch-scandal-serves-as-a-warning-for-europe-over-risks-of-using-algorithms/>. Accessed: 19-12-2024.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9): 389–399.
- Karimi, A.-H.; Barthe, G.; Schölkopf, B.; and Valera, I. 2022. A survey of algorithmic recourse: contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5): 1–29.
- Kavouras, L.; Tsopeles, K.; Giannopoulos, G.; Sacharidis, D.; Psaroudaki, E.; Theologitis, N.; Rontogiannis, D.; Fotakis, D.; and Emiris, I. 2023. Fairness aware counterfactuals for subgroups. *Advances in Neural Information Processing Systems*, 36: 58246–58276.
- Kuratomi, A.; Pitoura, E.; Papapetrou, P.; Lindgren, T.; and Tsaparas, P. 2022. Measuring the burden of (un) fairness using counterfactuals. In *Joint European conference on machine learning and knowledge discovery in databases*, 402–417. Springer.
- Laugel, T.; Lesot, M.-J.; Marsala, C.; Renard, X.; and Detryniecki, M. 2017. Inverse classification for comparison-based interpretability in machine learning. *arXiv preprint arXiv:1712.08443*.
- Liu, L. T.; Dean, S.; Rolf, E.; Simchowitz, M.; and Hardt, M. 2018. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, 3150–3158. PMLR.
- Martinez, N.; Bertran, M.; and Sapiro, G. 2020. Minimax pareto fairness: A multi objective perspective. In *International conference on machine learning*, 6755–6764. PMLR.
- Mehrabani, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.*, 54(6).
- Milli, S.; Miller, J.; Dragan, A. D.; and Hardt, M. 2019. The social cost of strategic classification. In *Proceedings of the conference on fairness, accountability, and transparency*, 230–239.
- Pawelczyk, M.; Bielawski, S.; Van den Heuvel, J.; Richter, T.; and Kasneci, G. 2021. CARLA: A Python Library to Benchmark Algorithmic Recourse and Counterfactual Explanation Algorithms. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Pawelczyk, M.; Broelemann, K.; and Kasneci, G. 2020. Learning model-agnostic counterfactual explanations for tabular data. In *Proceedings of the web conference 2020*, 3126–3132.
- Perdomo, J.; Zrnic, T.; Mendler-Dünner, C.; and Hardt, M. 2020. Performative prediction. In *International Conference on Machine Learning*, 7599–7609. PMLR.
- Perdomo, J. C.; Britton, T.; Hardt, M.; and Abebe, R. 2025. Difficult lessons on social prediction from wisconsin public schools. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 2682–2704.
- Pessach, D.; and Shmueli, E. 2023. Algorithmic fairness. In *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*, 867–886. Springer.
- Purificato, E.; Lorenzo, F.; Fallucchi, F.; and De Luca, E. W. 2023. The use of responsible artificial intelligence techniques in the context of loan approval processes. *International Journal of Human-Computer Interaction*, 39(7): 1543–1562.
- Raimondi, F. E.; Lawrence, A. R.; and Chockler, H. 2022. Equality of Effort via Algorithmic Recourse. *arXiv preprint arXiv:2211.11892*.
- Raji, I. D.; Kumar, I. E.; Horowitz, A.; and Selbst, A. 2022. The fallacy of AI functionality. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 959–972.
- Rodríguez-Hernández, C. F.; Cascallar, E.; and Kyndt, E. 2020. Socio-economic status and academic performance in higher education: A systematic review. *Educational Research Review*, 29: 100305.
- Tominaga, T.; Yamashita, N.; and Kurashima, T. 2024. Re-assessing Evaluation Functions in Algorithmic Recourse: An Empirical Study from a Human-Centered Perspective. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, 7913–7921. ijcai.org.
- Ustun, B.; Spangher, A.; and Liu, Y. 2019. Actionable recourse in linear classification. In *Proceedings of the conference on fairness, accountability, and transparency*, 10–19.
- Von Kügelgen, J.; Karimi, A.-H.; Bhatt, U.; Valera, I.; Weller, A.; and Schölkopf, B. 2022. On the fairness of causal algorithmic recourse. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 9584–9594.
- Wachter, S.; Mittelstadt, B.; and Russell, C. 2017. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech.*, 31: 841.
- Yetukuri, J.; Hardy, I.; Vorobeychik, Y.; Ustun, B.; and Liu, Y. 2024. Providing fair recourse over plausible groups. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 21753–21760.