

Equality of Effort via Algorithmic Recourse

Francesca E. D. Raimondi ¹, Andrew R. Lawrence ¹, and Hana Chockler ^{1,2}

¹causaLens

²Department of Informatics, King’s College London

Abstract

This paper proposes a method for measuring fairness through equality of effort by applying algorithmic recourse through minimal interventions. Equality of effort is a property that can be quantified at both the individual and the group level. It answers the counterfactual question: what is the minimal cost for a protected individual or the average minimal cost for a protected group of individuals to reverse the outcome computed by an automated system? Algorithmic recourse increases the flexibility and applicability of the notion of equal effort: it overcomes its previous limitations by reconciling multiple treatment variables, introducing feasibility and plausibility constraints, and integrating the actual relative costs of interventions. We extend the existing definition of equality of effort and present an algorithm for its assessment via algorithmic recourse. We validate our approach both on synthetic data and on the German credit dataset.

1 Introduction

Machine Learning systems are increasingly used in many socially significant applications, such as loan approval, hiring decisions, legal processes, and healthcare, sometimes encoding existing human and historical biases, as well as generating new biases through their algorithms. Fairness is commonly defined as *the absence of any prejudice or favoritism toward an individual or group based on their inherent or acquired characteristics* [1]. The algorithmic fairness literature is growing quickly, but the corresponding conceptualization and applicability need further study and structure [2]. For instance, the prevalent correlation-based fairness algorithms fail to detect discrimination in the presence of statistical anomalies such as Simpson’s paradox [3].

Group fairness notions, mainly based on statistical measures, are aimed at ensuring that groups who differ in their sensitive attributes (e.g. sex, ethnicity, religion) are treated equally or at least similarly [4]. However group fairness, despite its suitability for policies among demographic sub-populations, does not guarantee that individuals are treated fairly and is not suitable for fine-grained groups [5].

In the disparate treatment liability framework (as opposed to disparate impact), discrimination claims require proof of a causal connection between the challenged decision and the sensitive feature. Therefore, *causality* is becoming a fundamental tool in analyzing fairness of a decision [3]. Causal fairness notions are not based exclusively on data but also consider the causal structure of the world, describing how data is generated and how changes in variables propagate in a system.

Ideally, an automated system should produce the same output for every pair of individuals who differ only in their sensitive attributes. This concept is developed in [6] as a measure of *causal discrimination*, i.e. the fraction of inputs for which software discriminates. More generally, *fairness*

through awareness requires that similar individuals will get similar outcomes [7]. Similarity can be defined with respect to a specific task and is based on a trustworthy distance metric between individuals and a distance metric between probability distributions. This notion shows interesting developments for privacy purposes. Similarly, *individual direct discrimination* considers a target individual discriminated against if the difference observed between the rate of positive decisions for similar individuals in two groups: (1) sensitive (or protected) and (2) privileged (or unprotected), is higher than a predefined threshold [8].

Using similarity measures among individuals, [9] introduces the notion of *equality of effort*, that can help answer counterfactual questions like *by how much should an applicant’s credit score improve so that their loan application is approved?*, and judge discrimination from the perspective of equal effort among groups. Equality of effort establishes whether the effort to reach a certain outcome is the same for the protected and unprotected groups. A treatment variable is studied as to what change is required for an individual or group to achieve a certain outcome [9]. A similar direction is pursued in [10] with a focus on characterizing the long-term impact of algorithmic policies on reshaping the underlying populations.

We propose a novel way of estimating equality of effort by applying algorithmic recourse: *the actions with minimal costs required for reversing unfavorable decisions by algorithms and bureaucracies across a range of counterfactual scenarios* [11]. This has several advantages due to the flexibility of the approach in defining feasibility and plausibility constraints on variables, as well as in expressing the minimal cost required to change the outcome. In particular, algorithmic recourse, by including multiple treatment variables and integrating uniquely defined costs of interventions (e.g. relative or asymmetric), increases the expressiveness of the concept of equality of effort. Fairness of recourse was already shown to be complementary to counterfactual fairness of prediction in [12], where the cost of recourse for an individual is compared to the cost of recourse for its counterfactual twin (i.e. a virtual individual with a counterfactual change of the sensitive variable). This opens opportunities to alternative solutions to unfairness issues by societal interventions, as opposed to the direct modification of the classifier as in [13].

In §2 and §2.2 we introduce the concepts of equality of effort and algorithmic recourse, respectively; in §3 we develop our approach and outline our algorithm; in §4 we present results both on synthetic data and on the German credit dataset; finally in §5 we conclude with suggested extensions to our method.

2 Background

2.1 Equality of effort

Unlike most causal fairness notions, which intervene on the sensitive attribute $S \notin X$, equality of effort intervenes on a treatment variable $T \in X$, where X is the set of covariates. We note that any actionable variable in X can be considered as a possible treatment. The definition in [9] belongs to the potential outcome framework:

Definition 1 (γ -Minimum Effort [9]) *For individual i with value $(s_i, t_i, \mathbf{x}_i, y_i)$, given the intervention: $Y_i(t) = (Y_i | do(t_i = t))$ and a scalar $\gamma \in [0, 1]$, the minimum value of the treatment variable to achieve γ -level outcome is defined as ¹*

$$\Psi_i(\gamma) = \arg \min_{t \in T} \{ \mathbb{E}[Y_i(t)] \geq \gamma \} \tag{1}$$

and the *minimum effort* to achieve γ -level outcome is $\Psi_i(\gamma) - t_i$.

¹ \mathbb{E} denotes the expectation mathematical operator.

The strategy proposed in [9] is based on situation testing² and the definition of a similarity to form a subset of individuals, $I = I^+ \cup I^-$, each of whom shares the same or similar characteristics (\mathbf{x} and t) as individual i [8]. We denote by I^+ and I^- the subgroups of elements in I with the binary sensitive attribute value $S = s_0$ and $S = s_1$, respectively. The derivation of the expected outcome under treatment $T = t$ of the two groups, $\mathbb{E}[Y_{I^+}(t)]$ and $\mathbb{E}[Y_{I^-}(t)]$ leads to a comparison of the γ -minimum efforts, $\Psi_{I^+}(\gamma)$ and $\Psi_{I^-}(\gamma)$, with

$$\Psi_{I^*}(\gamma) = \arg \min_{t \in T} \{\mathbb{E}[Y_{I^*}(t)] \geq \gamma\} \quad (2)$$

Equality of effort is achieved when the effort discrepancy $\Psi_{I^+}(\gamma) - \Psi_{I^-}(\gamma)$ between the two sets of individuals/groups is null.

Definition 2 (Equality of Effort) *For a given level γ , there is an equality of effort for individual i if $\Psi_{I^+}(\gamma) = \Psi_{I^-}(\gamma)$. The difference $\delta(\gamma) = \Psi_{I^+}(\gamma) - \Psi_{I^-}(\gamma)$ measures the effort discrepancy at the individual level.*

This definition can be extended to define effort discrepancies at a subgroup level or at the system level, if, instead of an individual i , we consider specific subgroups or the whole dataset, D^+ and D^- .

Definition 3 (Average Effort Discrepancy (AED)) *If $\gamma \in \Gamma$, where Γ is a discrete value set of the minimum expectation of the outcome variable, the Average Effort Discrepancy is defined as*

$$AED = \frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} \delta(\gamma); \quad \text{if } \gamma \in [\gamma_1, \gamma_2] \text{ is continuous, } \quad AED = \frac{1}{\gamma_2 - \gamma_1} \int_{\gamma_1}^{\gamma_2} \delta(\gamma) d\gamma.$$

The work in [9] is based on counterfactual causal inference and develops an optimization-based framework for removing discriminatory effort unfairness from the data if discrimination is detected. Definition 3 is based on the assumptions of no hidden confounders [14], monotonicity of the expectation of the outcome variable, invertibility of the outcome function, the presence of one binary protected attribute and one binary decision. When the invertibility assumption does not hold, causal inference methods such as outcome regression and propensity score inverse weighting are used to estimate the average treatment effect in the subgroups I^+ and I^- . The main limitation of equality of effort is that, generally, a single treatment variable does not appropriately reflect the unfairness between protected and unprotected groups.

2.2 Algorithmic recourse

We find the definition of recourse through minimal interventions in [11], as opposed to recourse via nearest counterfactual explanations in [15], to be particularly suited to our problem of deriving the individual cost to achieve a certain outcome. Assuming causal sufficiency of the structural causal model \mathcal{M} (i.e., no hidden confounders), additive noise, and full specification of an invertible \mathbb{F} such that $\mathbb{F}(\mathbb{F}^{-1}(x)) = x$, endogenous variable X can be uniquely determined given the values of exogenous variables U (and vice-versa, given the invertibility of \mathbb{F}). The outcome Y is related to the variables X by a function $h(\cdot)$: $Y = h(X)$.

Definition 4 (Algorithmic recourse via minimal interventions (MINT) [11]) *is an optimization problem which minimizes the cost of the set of actions A (in the form of structural*

²Situation testing is a *legally grounded technique for analyzing the discriminatory treatment on individuals* [8].

interventions on the factual instance x^F) that results in a structural counterfactual instance x^{SCF} yielding the favorable output from h :

$$\begin{aligned}
 A^*(x^F) &\in \arg \min_A \text{cost}(A; x^F) \quad \text{s.t.} \\
 h(x^F) &\neq h(x^{SCF}) \\
 x^{SCF} &= \mathbb{F}_A(\mathbb{F}^{-1}(x^F)) \\
 x^{SCF} &\in \mathcal{P}, A \in \mathcal{F}
 \end{aligned} \tag{3}$$

with $\text{cost}(\cdot; x^F) : \mathcal{F} \times \mathcal{X} \rightarrow \mathbb{R}_+$ and $x^{*SCF} = \mathbb{F}_{A^*}(\mathbb{F}^{-1}(x^F))$. We can deterministically compute the counterfactual $x^{SCF} = \mathbb{F}_A(\mathbb{F}^{-1}(x^F))$ by performing the Abduction-Action-Prediction steps proposed by [14]. If J is the set of variables that have been intervened upon,

$$x_j^{SCF} = [j \in J](x_j^F + \delta_j) + [j \notin J](x_j^F + f_j(pa_j^{SCF}) - f_j(pa_j^F)) \tag{4}$$

Equation 4 holds for an additive noise model with mutually independent exogenous variables. The structural counterfactual value of the j -th feature, x_j^{SCF} , takes the value $(x_j^F + \delta_j)$ if this feature is intervened upon. Otherwise, x_j^{SCF} is computed as a function of both the factual and counterfactual values of its parents, denoted respectively by pa_j^F and pa_j^{SCF} . The closed-form expression in Equation 4 can replace the counterfactual constraint in Equation 3, i.e., $x^{SCF} = \mathbb{F}_A(\mathbb{F}^{-1}(x^F))$, after which the optimization problem in Equation 3 may be solved by building on existing frameworks for generating nearest counterfactual explanations [15], including gradient-based, evolutionary-based, heuristics-based, or verification-based approaches.

An example of cost provided by [11] for Equation 3 is the ℓ_1 norm over normalized feature changes to make effort comparable across features, i.e.

$$\text{cost}(\delta) = \sum_{j \in J} |\delta_j|/R_j, \quad \text{where } R_j \text{ is the range of feature } j. \tag{5}$$

3 Equality of effort via algorithmic recourse

The authors in [9] limit their algorithms to changing one variable and do not consider weighted versions of the problem, representing the cost and/or the effort required in order to change the value of a variable.

The recourse-based approach has several advantages. We can consider the treatment variable $T \in X$ in Section 2 as part of the feasibility constraints in algorithmic recourse, where we can declare which features are immutable, mutable but non-actionable, and actionable (we can intervene directly on them), and we can also add bounds to define feasible interventions (e.g. age can only increase). Note that in [9] only one treatment variable is considered, whereas algorithmic recourse allows us to extend the interventions to multiple treatment variables, $T \subseteq X$, included in the feasibility constraint \mathcal{F} . Moreover, the calculation of the counterfactual instances allows us to propagate the effect of the change of the treatment variables throughout the structural causal model (cf. Equation 4), and to introduce plausibility constraints on the affected covariates. Finally, the relative costs of changing different variables can be easily reflected in the cost function as weights.

Definition 5 (Average Minimal Effort) *For an individual i belonging to the sensitive group $S = s_0$, we can define an Average Minimal Effort in its neighborhood subsets $I^+ \subset I$ (where $S = s_0$)*

and $I^- \subset I$ (where $S = s_1$):

$$\Phi_{I^+} = \frac{1}{|I^+|} \sum_{x \in I^+} \text{cost}(A^*(x)) \quad x \in I^+ \quad (6)$$

$$\Phi_{I^-} = \frac{1}{|I^-|} \sum_{x \in I^-} \text{cost}(A^*(x)) \quad x \in I^- \quad (7)$$

where A^* is given by Equation 3 and $|I^*|$ denotes the cardinality of I^* .

Definition 6 (Equality of Effort via Algorithmic Recourse) We say that there is an equality of effort for an individual i if $\Phi_{I^+} = \Phi_{I^-}$. The ratio $\Phi = \Phi_{I^+}/\Phi_{I^-}$ measures the **Average Cost Ratio (ACR)** at the individual level.

We say that recourse is possible for an instance x if there exists $A^*(x)$ that is a solution to the recourse problem in Equation 3.

Definition 7 We define the **Ratios of Possible Recourse** for an individual i as

$$\rho_{I^+} = \sum_{x \in I^+} \frac{\mathbb{I}_{\{\exists A^*(x)\}}}{|I^+|}, \quad \rho_{I^-} = \sum_{x \in I^-} \frac{\mathbb{I}_{\{\exists A^*(x)\}}}{|I^-|} \quad (8)$$

where $\mathbb{I}_{\{x\}}$ denotes the characteristic function of x .

Definition 8 The difference $\rho = \rho_{I^-} - \rho_{I^+}$ measures the **Recourse Discrepancy (RD)** at the individual level.

The RD can identify unfairness whenever a significant ratio of the elements of at least one of the two groups are so far from the decision boundaries that recourse is not a possibility for them. In fact, in this case, only considering the ACR for the instances that admit recourse could be misleading.

We can easily extend this to the subgroup level or to the system level as in [9], if instead of individual i we consider specific subgroups or the whole dataset D^+ and D^- , respectively. For the estimation of the ACR, we propose Algorithm 1, which has the benefit of giving an individual definition of fairness for individual i , if applied at an individual level.

For the sake of simplicity, we define Algorithm 1 for binary sensitive variables and for binary outcomes; however, it can be generalized to both categorical sensitive variables and outcomes. In the binary case, we only perform recourse on the instances with unfavorable outcomes.

Note that in [9], Equation 2 relies on the estimation of the average treatment effect on outcome Y (Y will be favorable at least γ fraction of the time) in the individual neighborhoods or group/dataset, and then this quantity is averaged or integrated again on all possible values of γ in order to derive the AED in Definition 3. On the other hand, the proposed Algorithm 1 performs algorithmic recourse for each point of the cluster that has a negative outcome: our Y is always flipped by the recourse for every structural counterfactual instance x^{*SCF} or we find there is no feasible and plausible structural counterfactual to flip the outcome from undesired to desired as there may not always exist an action to surpass the decision boundary given the constraints. Our definition removes the need for further averaging or integration.

Inspired by counterfactual fairness in [16], [12] define the fairness of algorithmic recourse: recourse may be considered fair at the level of the individual if the cost of recourse would have been the same, had the individual belonged to a different protected group, i.e., under a counterfactual change

Algorithm 1 Equality of effort via algorithmic recourse

Input

Dataset $D = \{D^+, D^-\}$, thresholds τ, ϵ , sensitive variable $S \in \{s_0, s_1\}$, unfavorable outcome y , feasibility constraints \mathcal{F} , plausibility constraints \mathcal{P} , individual i (optional), distance d (optional).

Output

Decision: *Is there equality of effort?*

1. **if** individual i and distance d are provided as inputs **then**:
 2. Compute the similarity subsets $I = \{I^+, I^-\}$ using distance d .
 3. **else**:
 4. Set $I = D, I^+ = D^+$ and $I^- = D^-$.
 5. **for** each subset $I^* \in \{I^+, I^-\}$ **do**:
 6. Identify the instances x , s.t. $h(x) = y$, where $y \in \{0, 1\}$ is the unfavorable outcome (i.e. the outcome we wish to change).
 7. Include and bound treatment variables $T \subseteq X$ as actionable variables in the feasibility constraint \mathcal{F} and bound the covariates X in the plausibility constraints \mathcal{P} in Equation 3.
 8. Initialize the count for the Ratio of Possible Recourse in group I^* : $\rho_{I^*} = 0$
 9. **for** $\forall x \in I^*$ **do**:
 10. Perform algorithmic recourse on x following Equations 3 and 4.
 11. **if** there exists a solution to algorithmic recourse for instance x **then**:
 12. Calculate the cost of recourse for instance x , $\text{cost}(A^*(x))$.
 13. **else**:
 14. $\rho_{I^*} = \rho_{I^*} + \frac{1}{|I^*|}$
 15. Calculate the Average Minimal Effort as $\Phi_{I^*} = \mathbb{E}[\text{cost}(A^*(x))]$ for $x \in I^*$ (Definition 5).
 16. Compute ACR as $\Phi = \Phi_{I^+}/\Phi_{I^-}$ (Definition 6).
 17. Compute RD as $\rho = \rho_{I^-} - \rho_{I^+}$ (Definition 8).
 18. **if** $|\rho| \geq \epsilon$ **then: return** False
 19. **else if** $|\Phi| \leq \tau$ **then: return** True
 20. **else: return** False
-

to the sensitive variable S . [12] show that the concepts of counterfactual fairness of prediction and counterfactual fairness of recourse are complementary. We note that counterfactually fair algorithms do not necessarily imply fairness of recourse, and, conversely, fairness of recourse does not guarantee counterfactual fairness, as detailed in [12]. Moreover, when an algorithm is counterfactually unfair (i.e. a flip of the sensitive attribute for an individual reverses an unfavorable outcome, thus removing the need for recourse), counterfactual fairness of recourse, as defined by [12], concludes maximal unfairness at the individual level. Our approach based on situation testing offers a novel way of measuring the fairness of algorithmic recourse, even when the counterfactual twin has already seen its outcome reversed.

Note that Algorithm 1 has to limit the feasibility constraints in Equation 3: the actionable variables or treatment variables $T \subseteq X$ cannot include the sensitive variable S . In the trivial and degenerate case when the treatment variable and the sensitive variable were allowed to coincide, i.e. $T = S$, the mere existence of recourse on the sensitive variable would imply counterfactual unfairness for the individual and would not carry any information about equality of effort. In fact, the notion of counterfactual fairness in [16], [17], albeit an individual notion, offers a different perspective in that the counterfactual change is related to a simple flip of the sensitive variable. In our case, we want to establish the cost of changing other treatment variables that is required to change the outcome.

The choice of the distance [9] extend the situation testing developed in [8] to define the two subsets I^+ and I^- for each assessment. However, the distance function in [8] is unnecessarily complex [3]. We take the ℓ_1 or ℓ_2 norms if all the K variables are continuous, otherwise we follow an idea detailed in [18]. The distance between two domain values x_k and x'_k of variable $X_k, k \in \{1, \dots, K\}$ is defined as the value difference, where the normalized ℓ_1 distance or Manhattan distance is employed for continuous/ordinal attributes, and the overlap measurement is employed for categorical attributes.

Definition 9 *The normalized Manhattan distance is defined as*

$$md(x_k, x'_k) = \frac{|x_k - x'_k|}{range_k} = \frac{|x_k - x'_k|}{\max_\ell(x_k^\ell) - \min_\ell(x_k^\ell)} \quad (9)$$

Definition 10 *The overlap measurement is defined as*

$$overlap(x_k, x'_k) = \begin{cases} 0 & \text{if } x_k = x'_k \\ 1 & \text{otherwise} \end{cases} \quad (10)$$

Definition 11 *The difference between two points x, x' is given by $d(x, x') = \sum_{k=1}^K dist(x_k, x'_k)$*

$$\text{where } dist(x_k, x'_k) = \begin{cases} md(x_k, x'_k) & \text{if } X_k \text{ is continuous/ordinal} \\ overlap(x_k, x'_k) & \text{if } X_k \text{ is categorical} \end{cases} \quad (11)$$

4 Results

We present results of the proposed method on synthetic data and real data from the German credit dataset [19] (cf. Figure 1). In both cases, we learn the structural equations by fitting a linear regression model to the child-parent tuples, and a logistic regression model for the binary outcome. Distances are based on Definition 11 (the choice of the ℓ_1 norm over ℓ_2 norm for distances among neighbors does not seem to have an impact on results). For the sake of simplicity, we do not impose strict bounds on the feasibility constraints, and we fit linear relationships between variables in the causal graph. This reflects in the existence of recourse for all the individuals, meaning that we do not need to leverage the Recourse Discrepancy (RD) and we only have to consider the Average Cost Ratio (ACR) in Algorithm 1. Also, we do not impose different relative costs on variables, but consider them equally important, i.e. we use uniform weights so the same change in each feature results in the same cost. Hence, we use Equation 5 to compute costs.

4.1 Synthetic data

Synthetic data is composed of $n = 1000$ samples following Figure 1a.

- The sensitive variable X_1 is given by the exogenous variable $U_1 \sim Bernoulli(0.5)$.
- The proxy and actionable variable X_2 is given by $X_2 = \alpha \cdot X_1 + U_2$. Exogenous random variable is $U_2 \sim \mathcal{N}(3, 1)$ and parameter α controls the strength of the relationship with X_1 .
- The actionable variable X_3 is given by exogenous random variable $U_3 \sim \mathcal{N}(0, 1)$.
- The outcome is $Y = f(std(X_2 + X_3))$ where $f(\cdot)$ is the logistic sigmoid function and $std(\cdot)$ indicates standardization.
- \hat{Y} is the predicted outcome, according to logistic regression.

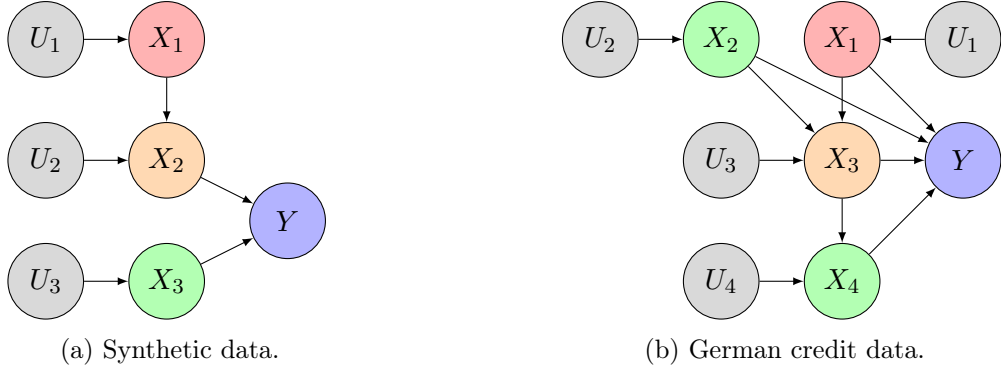


Figure 1: Directed Acyclic Graphs (DAGs).

Figure 2a refers to the ACR in Definition 6, whereas Figure 2b refers to the ratio of protected individuals in the neighborhood $I = I^+ \cup I^-$, both as functions of the quantile threshold on distance for the subset I centered on each individual. The center line refers to the average value, enclosed by the 95% confidence interval. The variance tends to be higher for smaller quantiles (as there are fewer elements in each subset), whereas for unitary quantile the variance is zero as we always consider the whole dataset for comparisons and can therefore derive the effort discrepancy at system level from that point. The ACR steadily remains over 1.2 for all radii, thus confirming the unfairness (inequality of effort) at both the individual level ($ACR(0.2) \approx 1.5$) and the system level ($ACR(1) \approx 2$). On the contrary, for the unprotected group $ACR(q) \leq 0.8 \forall q \in Q$ and $ACR(1) \approx 0.5$. Note that the two curves of Figure 5a diverge at the system level (i.e. for high quantiles).

Figure 3a describes the ACR averaged over all the protected individuals ($S = s_0$), whereas Figure 3b presents the corresponding average ratio of elements with the sensitive attribute in the subset $I = I^+ \cup I^-$, both for different values of α . We notice that as α increases, the unfairness also grows in terms of the inequality of effort. When α is zero, we see that ACR stays at one and we can conclude the system has an equality of effort.

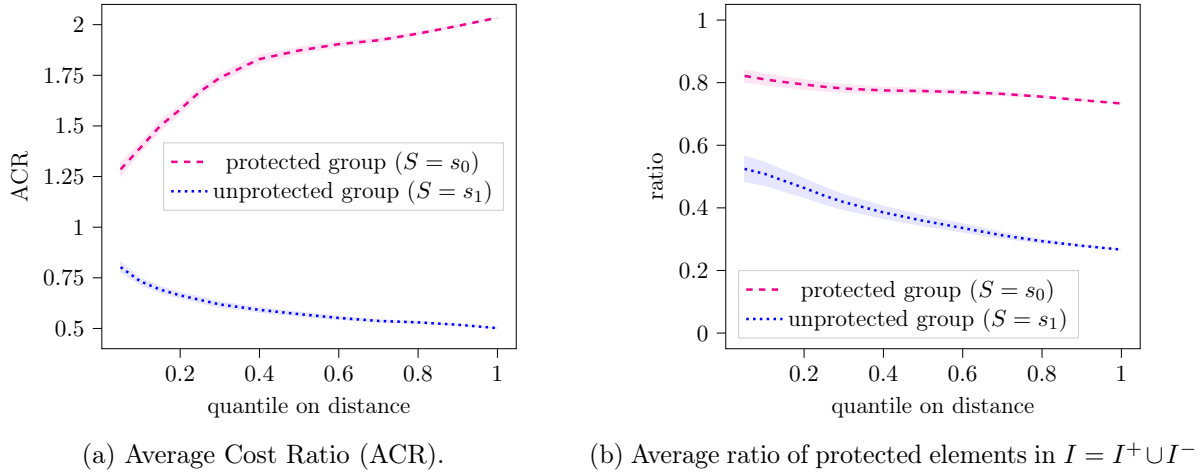
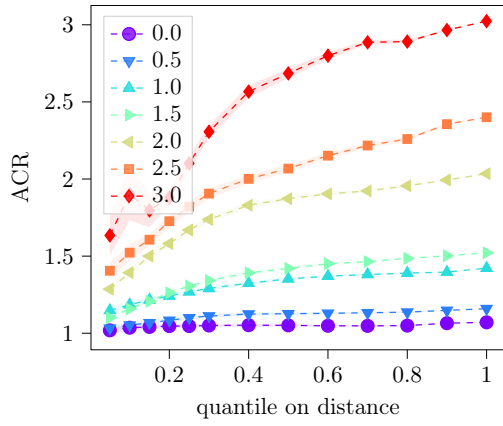
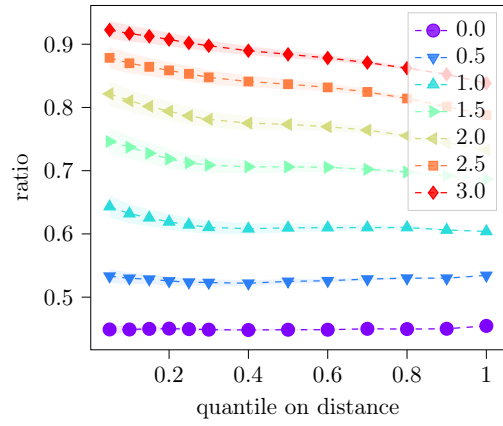


Figure 2: Equality of Effort for $\alpha = 2$ - synthetic data.

Figure 4 describes in more detail the average costs for the subsets I^+ and I^- , respectively averaged over all the protected individuals ($S = s_0$) in Figure 4a) and all the unprotected individuals

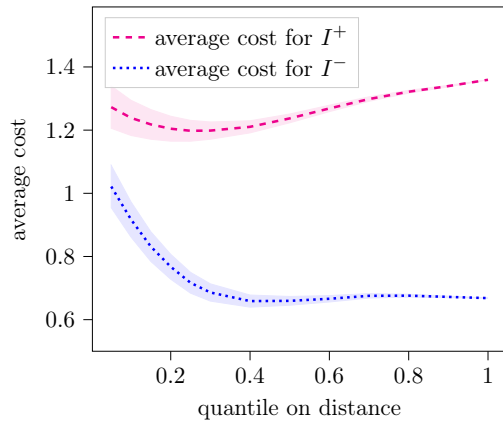


(a) Average Cost Ratio (ACR).

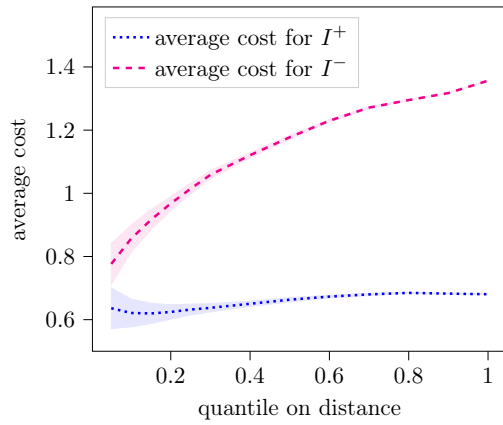


(b) Average ratio of protected elements in $I = I^+ \cup I^-$.

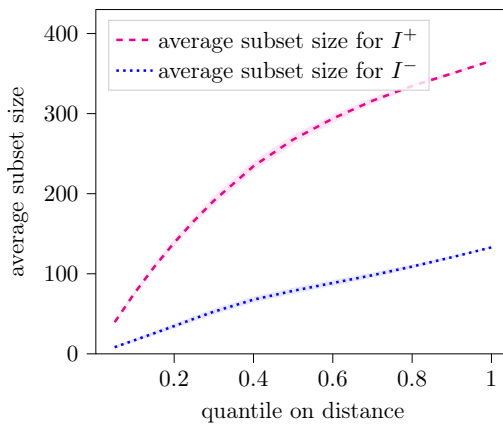
Figure 3: Plots for different values of α - synthetic data.



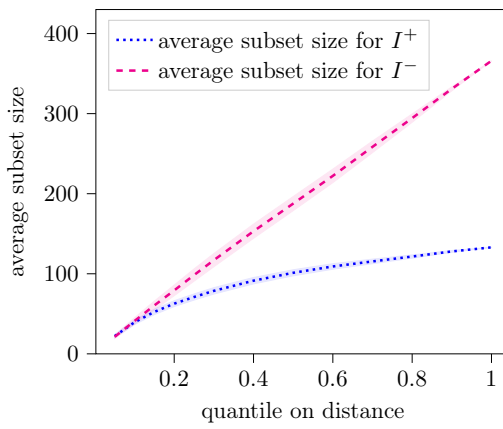
(a) Cost of recourse, protected group ($S = s_0$).



(b) Cost of recourse, unprotected group ($S = s_1$).



(c) Size of subsets, protected group ($S = s_0$).



(d) Size of subsets, unprotected group ($S = s_1$).

Figure 4: Details of recourse costs and subset size - synthetic data for $\alpha = 2$.

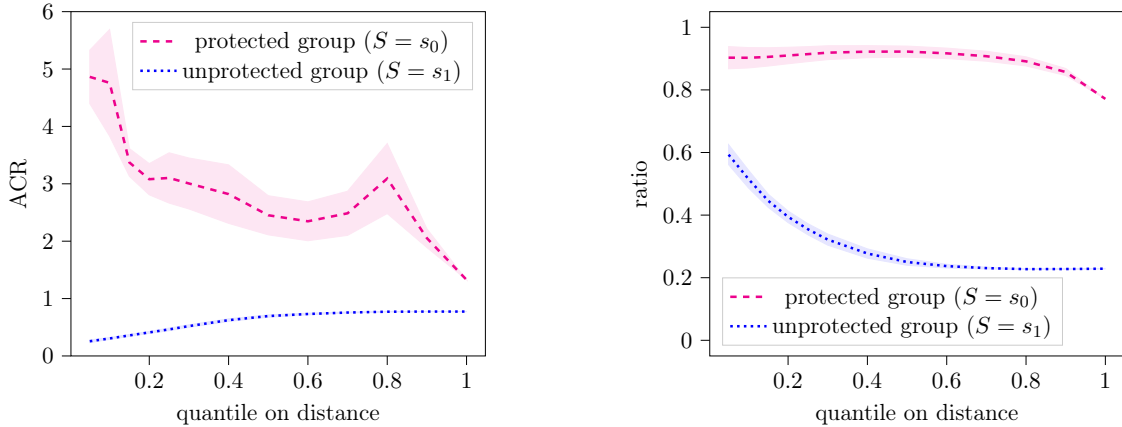
($S = s_1$) in Figure 4b), both with unfavorable outcome ($Y = 0$); whereas Figure 4c and 4d represent the corresponding number of elements in I^+ and I^- .

4.2 Real data

The German Credit dataset from the UCI repository contains 20 attributes of 1000 individuals applying for loans [19]. The outcome is binary: a low or high credit risk, representing the likelihood of repaying the loan. We consider a subset of the features, similarly to [11]. The setup is depicted in Figure 1b:

- X_1 is the individual’s sex (the sensitive attribute, treated as immutable).
- X_2 is the individual’s age (actionable but can only increase).
- X_3 is the credit amount given by the bank (actionable).
- X_4 is the repayment duration of the credit (non-actionable but mutable).
- The outcome Y is the ground truth for customer risk.
- \hat{Y} is the predicted customer risk, according to logistic regression.

Similarly, Figure 5a refers to the ACR in Definition 6, i.e. the ratio of the costs of the two subsets, whereas Figure 5b refers to the ratio of protected individuals in the neighborhood $I = I^+ \cup I^-$. The ACR steadily remains well over 1.2 for all radii in the protected group, thus confirming the unfairness (inequality of effort) at both the individual level with $ACR(0.2) \approx 3$ and at the system level with $ACR(1) \approx 1.3$. On the contrary, for the unprotected group $ACR(q) < 0.8 \forall q \in Q$ and $ACR(1) \approx 0.77$. The two curves of Figure 5a diverge at the individual level while converging at the system level, showing that group fairness notions might struggle to capture individual unfairness.

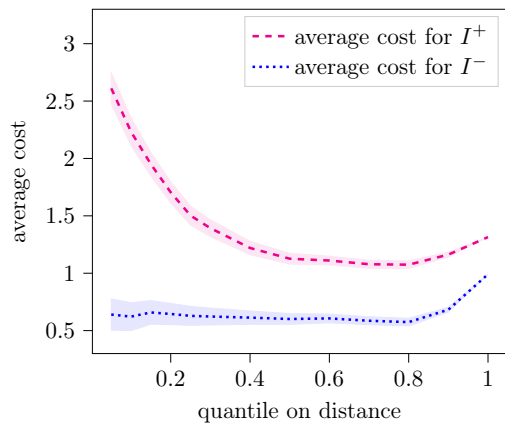


(a) Average Cost Ratio (ACR).

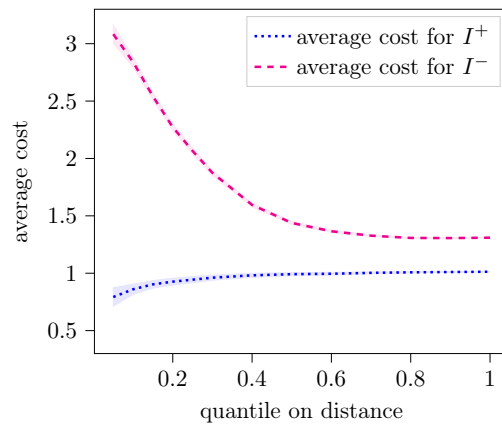
(b) Average ratio of protected elements in $I = I^+ \cup I^-$.

Figure 5: Equality of Effort - German credit dataset.

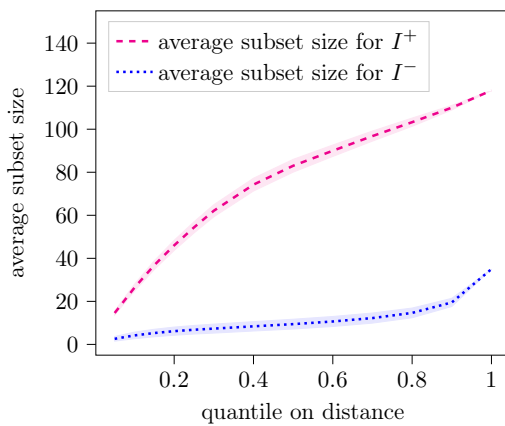
Figure 6 describes in more detail the average costs for the subsets I^+ and I^- , respectively averaged over all the protected individuals ($S = s_0$) in Figure 6a) and all the unprotected individuals ($S = s_1$) in Figure 6b), both with unfavorable outcome ($Y = 0$); whereas Figure 6c and 6d represents the corresponding number of elements in I^+ and I^- .



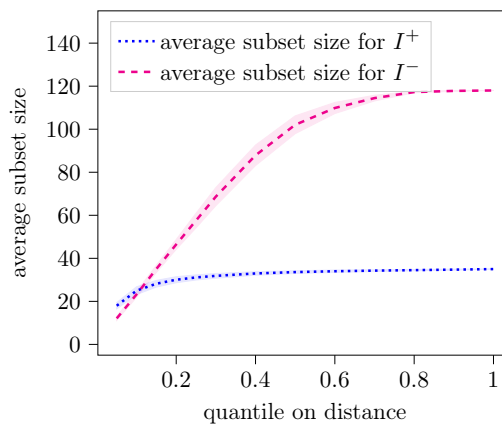
(a) Costs of recourse, protected group ($S = s_0$).



(b) Costs of recourse, unprotected group ($S = s_1$).



(c) Size of subsets, protected group ($S = s_0$).

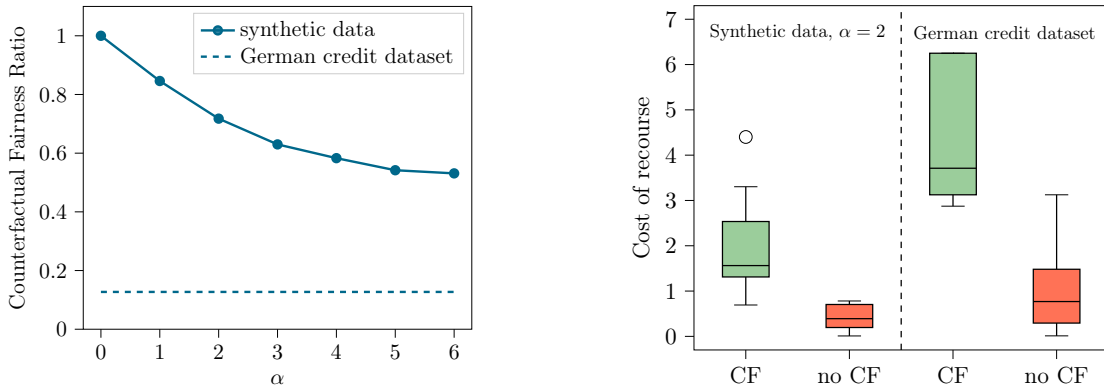


(d) Size of subsets, unprotected group ($S = s_1$).

Figure 6: Details of recourse costs and subset size - German credit dataset.

4.3 Comparison between counterfactual fairness and equality of effort

We observe that the notions of counterfactual fairness and of equality of effort are complementary as each offers different insights into potential unfairness. Figure 7a shows results of counterfactual fairness ratio (CFR), i.e. the ratio of individuals towards whom our classification algorithm is counterfactually fair as in [17], for some configurations of synthetic data and the German credit dataset. Note that, as an inverse to ACR in Figure 3a, the CFR is a monotonically decreasing function of α , because it is a direct measure of fairness (as opposed to unfairness). Figure 7b shows the distribution of recourse costs, if we consider the subsets of counterfactually fairly treated individuals (CF) and counterfactually unfairly treated individuals (no CF) for both synthetic data with $\alpha = 2$ and for the German credit dataset. The costs are considerably lower under counterfactual unfairness, because counterfactually unfairly treated individuals tend to be closer to the decision boundary as their counterfactual twin would cross the boundary.



(a) CFR = 1 for complete fairness.

(b) Recourse cost by counterfactual fairness outcome.

Figure 7: Comparison with counterfactual fairness.

5 Conclusion

We presented a novel approach to assess equality of effort, based on the cost of algorithmic recourse via minimal interventions. The proposed algorithm has several advantages: the flexibility of the recourse actions (multiple variables can be actionable at the same time) and of the feasibility and plausibility constraints that can be imposed on them; the cost of recourse can be derived directly from the optimization problem; finally relative costs of variables could be directly included in the model as weights in the cost expression, thus capturing a real cost or effort for each intervention. We discussed how this notion relates to other causal notions of fairness (equality of effort in [9] and counterfactual fairness of prediction and recourse in [16], [17]) and [12], respectively. We applied this notion to synthetic data as well as real data from the German credit dataset, averaging results for different values of hyperparameters and showing equality and inequality of effort at both the individual and the system level. Our idea can be developed further by including plausibility constraints on covariates and relative costs on the actionable variables in the recourse problem; by enforcing more complex feasibility constraints and taking into account the Recourse Discrepancy in Definition 8; and by studying the complexity of the algorithm and introducing matching techniques to deal with large datasets.

References

- [1] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, “A survey on bias and fairness in machine learning,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1–35, 2021.
- [2] A. Chouldechova and A. Roth, “A snapshot of the frontiers of fairness in machine learning,” *Communications of the ACM*, vol. 63, no. 5, pp. 82–89, 2020.
- [3] K. Makhoulouf, S. Zhioua, and C. Palamidessi, “Survey on causal-based machine learning fairness notions,” *arXiv preprint arXiv:2010.09553*, 2020.
- [4] D. Pessach and E. Shmueli, “Algorithmic fairness,” *arXiv preprint arXiv:2001.09784*, 2020.
- [5] K. Makhoulouf, S. Zhioua, and C. Palamidessi, “Machine learning fairness notions: Bridging the gap with real-world applications,” *Information Processing & Management*, vol. 58, no. 5, p. 102642, 2021.
- [6] S. Galhotra, Y. Brun, and A. Meliou, “Fairness testing: Testing software for discrimination,” in *Proceedings of the 2017 11th Joint meeting on foundations of software engineering*, 2017, pp. 498–510.
- [7] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, “Fairness through awareness,” in *Proceedings of the 3rd innovations in theoretical computer science conference*, 2012, pp. 214–226.
- [8] L. Zhang, Y. Wu, and X. Wu, “Situation testing-based discrimination discovery: A causal inference approach,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, vol. 16, 2016, pp. 2718–2724.
- [9] W. Huan, Y. Wu, L. Zhang, and X. Wu, “Fairness through equality of effort,” in *Companion Proceedings of the Web Conference 2020*, 2020, pp. 743–751.
- [10] H. Heidari, V. Nanda, and K. Gummadi, “On the long-term impact of algorithmic decision policies: Effort unfairness and feature segregation through social learning,” in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 97, PMLR, 2019, pp. 2692–2701.
- [11] A.-H. Karimi, B. Schölkopf, and I. Valera, “Algorithmic recourse: From counterfactual explanations to interventions,” in *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 2021, pp. 353–362.
- [12] J. Von Kügelgen, A.-H. Karimi, U. Bhatt, I. Valera, A. Weller, and B. Schölkopf, “On the fairness of causal algorithmic recourse,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, 2022, pp. 9584–9594.
- [13] V. Gupta, P. Nokhiz, C. D. Roy, and S. Venkatasubramanian, “Equalizing recourse across groups,” *arXiv preprint arXiv:1909.03166*, 2019.
- [14] J. Pearl, *Causality*. Cambridge university press, 2009.
- [15] A.-H. Karimi, G. Barthe, B. Balle, and I. Valera, “Model-agnostic counterfactual explanations for consequential decisions,” in *International Conference on Artificial Intelligence and Statistics*, PMLR, 2020, pp. 895–905.
- [16] M. J. Kusner, J. Loftus, C. Russell, and R. Silva, “Counterfactual fairness,” *Advances in neural information processing systems*, vol. 30, 2017.

- [17] H. Chockler and J. Y. Halpern, “On testing for discrimination using causal models,” in *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*, 2022.
- [18] L. Zhang, Y. Wu, and X. Wu, “A causal framework for discovering and removing direct and indirect discrimination,” *arXiv preprint arXiv:1611.07509*, 2016.
- [19] D. Dua and C. Graff, *UCI machine learning repository*, 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>.