

Decisions, Counterfactual Explanations and Strategic Behavior

Stratis Tsirtsis and Manuel Gomez-Rodriguez

Max Planck Institute for Software Systems
{stsirtsis, manuelgr}@mpi-sws.org

Abstract

Data-driven predictive models are increasingly used to inform decisions that hold important consequences for individuals and society. As a result, decision makers are often obliged, even legally required, to provide explanations about their decisions. In this context, it has been increasingly argued that these explanations should help individuals understand what would have to change for these decisions to be beneficial ones. However, there has been little discussion on the possibility that individuals may use the above *counterfactual explanations* to invest effort strategically in order to maximize their chances of receiving a beneficial decision. In this paper, our goal is to find policies and counterfactual explanations that are optimal in terms of utility in such a strategic setting. To this end, we first show that, given a pre-defined policy, the problem of finding the optimal set of counterfactual explanations is NP-hard. However, we further show that the corresponding objective is nondecreasing and satisfies submodularity. Therefore, a standard greedy algorithm offers an approximation factor of $(1 - 1/e)$ at solving the problem. Additionally, we also show that the problem of jointly finding both the optimal policy and set of counterfactual explanations reduces to maximizing a non-monotone submodular function. As a result, we can use a recent randomized algorithm to solve the problem, which offers an approximation factor of $1/e$. Finally, we illustrate our theoretical findings by performing experiments on synthetic and real lending data.

1 Introduction

Whenever a bank decides to offer a loan to a customer, a judge decides to grant bail to a person, or a company decides to hire a new employee, the decision is increasingly informed by a data-driven predictive model. In all these high-stakes applications, the goal of the decision maker is to take decisions that maximize a given utility function while the goal of the predictive model is to provide accurate predictions of the outcomes from a set of observable features. For example, a bank may decide whether or not to offer a loan to a customer using the model’s estimate of the probability that the customer would repay the loan.

In this context, there has been a tremendous excitement on the potential of data-driven predictive models to enhance decision making in high-stakes applications. However, there has also been a heated debate about their lack of transparency and explainability (Doshi-Velez and Kim, 2017; Weller, 2017; Lipton, 2018; Gunning and Aha, 2019; Rudin, 2019). As a result, there already exists a legal requirement to grant individuals who are subject to (semi)-automated decision making the *right-to-explanation* in the European Union (Voigt and Von dem Bussche, 2017; Wachter et al., 2017a). With this motivation, there has been a flurry of work on interpretable machine learning (Ribeiro et al., 2016; Koh and Liang, 2017; Lundberg and Lee, 2017; Chakraborty et al., 2017; Murdoch et al., 2019; Wachter et al., 2017b; Karimi et al., 2019; Mothilal et al., 2020), which has predominantly focused on developing methods to find explanations for the predictions made by a predictive model. Within this line of work, the work most closely related to ours (Wachter et al., 2017b; Karimi et al., 2019; Mothilal et al., 2020) aims to find counterfactual explanations that help individuals understand what would have to change for a predictive model to make a positive prediction about them. In our work, rather than explaining predictions, we pursue the development of methods to find counterfactual explanations for the decisions taken by a decision maker¹, which are ultimately what individuals who are subject to (semi)-automated decision

¹These counterfactual explanations help individuals understand what would have to change in order to receive a beneficial decision, rather than a positive prediction.

making typically care about.

Once we focus on explaining decisions, we cannot overlook the possibility that individuals may use these explanations to invest effort strategically in order to maximize their chances of receiving a beneficial decision. However, this is also an opportunity for us to find counterfactual explanations that help individuals to self-improve and eventually increase the utility of a decision policy, as noted by several studies in economics (Coate and Loury, 1993; Fryer and Loury, 2013; Hu and Chen, 2018) and, more recently, in the computer science literature (Kleinberg and Raghavan, 2019; Tabibian et al., 2020). For example, if a bank explains to a customer that, if she reduces her credit card debt by 20%, she will receive the loan she is applying for, she may feel compelled to reduce her overall credit card debt by the proposed percentage to pay less interest, improving her financial situation, and this will eventually increase the profit the bank makes when she is able to successfully return the loan. This is in contrast with previous work on interpretable machine learning, which have ignored the influence that (counterfactual) explanations (of predictions by a predictive model) may have on the accuracy of predictive models and the utility of the decision policies.

Our contributions. We cast the above problem as a Stackelberg game in which the decision maker moves first and shares her counterfactual explanations before individuals best-respond to these explanations and invest effort to receive a beneficial decision. Under this problem formulation, we first show that, given a pre-defined policy, the problem of finding the optimal set of counterfactual explanations is NP-hard by using a novel reduction of the Set Cover problem (Karp, 1972). However, we further show that the corresponding objective function is monotone and submodular and, as a direct consequence, it readily follows that a standard greedy algorithm offers a $(1 - 1/e)$ approximation ratio at solving the problem. In addition, we show that, given a pre-defined set of counterfactual explanations, the optimal policy is deterministic and can be computed in polynomial time. Building on this result, we can reduce the problem of jointly finding both the optimal policy and set of counterfactual explanations to maximizing a non-monotone submodular function. As a consequence, we can use a recent randomized algorithm to solve the problem, which offers an approximation factor of $1/e$. Finally, we perform a series of experiments with synthetic and real lending data to illustrate our theoretical findings and show that the counterfactual explanations and decision policies found by the above algorithms achieve higher utility than several competitive baselines.

Further related work. Our work builds upon previous work on interpretable machine learning, strategic machine learning, and machine-assisted decision making.

There is not yet an agreement on what constitutes a *good* post-hoc explanation in the literature on interpretable machine learning. However, most previous work focuses on one of the two following types of explanations: (i) feature-based explanations (Ribeiro et al., 2016; Koh and Liang, 2017; Lundberg and Lee, 2017) or (ii) counterfactual explanations (Wachter et al., 2017b; Karimi et al., 2019; Mothilal et al., 2020). Feature-based explanations help individuals understand the importance each feature has on a particular prediction, typically through local approximation, while counterfactual explanations help them understand what features would have to change for a predictive model to make a positive prediction about them. While our work also focuses on counterfactual explanations, it focuses on explaining decisions, rather than predictions. By doing so, it sheds light on the possibility of using explanations to increase the utility of a decision policy, uncovering a previously unexplored connection between interpretable machine learning and the nascent field of strategic machine learning.

Similarly as in our work, previous work on strategic machine learning also assumes that individuals may use knowledge, gained by transparency, to invest effort strategically in order to receive either a positive prediction (Dalvi et al., 2004; Brückner and Scheffer, 2011; Hardt et al., 2016; Dong et al., 2018; Hu et al., 2019; Milli et al., 2019; Miller et al., 2019) or a beneficial decision (Kleinberg and Raghavan, 2019; Tabibian et al., 2020). However, none of this previous work focuses on finding (counterfactual) explanations and they assume full transparency—individuals who are subject to (semi)-automated decision making can observe the entire predictive model or the decision policy. As a result, their formulation is fundamentally different and their technical contributions are orthogonal to ours.

In the machine-assisted decision making literature, the distinction between decisions and predictions has not been made explicit until very recently (Corbett-Davies et al., 2017; Kleinberg et al., 2018; Mitchell et al., 2018; Valera et al., 2018; Kilbertus et al., 2019; Tabibian et al., 2020). However, previous work has focused on the design of optimal decision policies rather than (counterfactual) explanations.

2 Preliminaries

Given an individual with a feature vector $\mathbf{x} \in \{1, \dots, n\}^d$ and a (*ground-truth*) label $y \in \{0, 1\}$, we assume a decision $d(\mathbf{x}) \in \{0, 1\}$ controls whether the corresponding label is *realized*². This setting fits a variety of real-world scenarios. For example, in university admissions, the decision specifies whether a student is admitted ($d(\mathbf{x}) = 1$) or rejected ($d(\mathbf{x}) = 0$); the label indicates whether the student completes the program ($y = 1$) or drops out ($y = 0$) upon acceptance; and the feature vector (\mathbf{x}) may include her GRE scores, undergraduate GPA, or research experience. Throughout the paper, we denote the set of feature values as $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$, where $m = n^d$ denotes the number of feature values.

Each decision is sampled from a decision policy $d(\mathbf{x}) \sim \pi(d|\mathbf{x})$, where, for brevity, we will write $\pi(\mathbf{x}) = \pi(d = 1|\mathbf{x})$. For each individual, the label y is sampled from a conditional probability distribution $y \sim P(y|\mathbf{x})$ and, without loss of generality, we index the feature values in decreasing order with respect to their corresponding outcome, *i.e.*, $i < j \Rightarrow P(y = 1|\mathbf{x}_i) \geq P(y = 1|\mathbf{x}_j)$. Moreover, we adopt a Stackelberg game-theoretic formulation in which each individual with initial feature value \mathbf{x}_i receives a (counterfactual) explanation from the decision maker by means of a feature value $\mathcal{E}(\mathbf{x}_i) \in \mathcal{A} \subseteq \mathcal{P}_\pi := \{\mathbf{x} \in \mathcal{X} : \pi(\mathbf{x}) = 1\}$ before she (best-)responds³. Here, note that each individual is guaranteed to receive a positive decision if she changes from her initial feature value \mathbf{x}_i to the feature value $\mathcal{E}(\mathbf{x}_i)$. Inspired by Tabibian et al. (2020), we will assume that each individual’s best response is to change from her initial feature value \mathbf{x}_i to $\mathcal{E}(\mathbf{x}_i)$ if and only if

$$\mathcal{E}(\mathbf{x}_i) \in \mathcal{R}(\mathbf{x}_i) := \{\mathbf{x}_j \in \mathcal{X} : b(\pi, \mathbf{x}_j) - c(\mathbf{x}_i, \mathbf{x}_j) \geq b(\pi, \mathbf{x}_i)\} \quad (1)$$

and it is to keep her initial feature value \mathbf{x}_i otherwise. In the above equation, $c(\mathbf{x}, \mathcal{E}(\mathbf{x}_i))$ is the cost⁴ an individual pays for changing from \mathbf{x}_i to $\mathcal{E}(\mathbf{x}_i)$ and $b(\pi, \mathbf{x}) = \mathbb{E}_{d \sim \pi(d|\mathbf{x})}[f(d(\mathbf{x}))]$ is the (immediate) benefit she obtains from a policy π , where the function $f(\cdot)$ is problem dependent. Here, for simplicity, we assume that $f(d(\mathbf{x})) = \pi(\mathbf{x})$ and thus the benefit is proportional to the probability that an individual receives a positive decision and, for each feature value \mathbf{x} , we refer to $\mathcal{E}(\mathbf{x})$ as its counterfactual explanation and $\mathcal{R}(\mathbf{x})$ as its *region of adaptation*.

At a population level, the above best response results into a transportation of mass between the original feature distribution $P(\mathbf{x})$ and a new feature distribution $P(\mathbf{x}|\pi, \mathcal{A})$ induced by the policy π and the counterfactual explanations \mathcal{A} . More specifically, we can readily derive an analytical expression for the induced feature distribution in terms of the original feature distribution, *i.e.*,

$$P(\mathbf{x}_j|\pi, \mathcal{A}) = P(\mathbf{x}_j)\mathbb{I}(\mathcal{R}(\mathbf{x}_j) \cap \mathcal{A} = \emptyset) + \sum_{i \in [m]} P(\mathbf{x}_i)\mathbb{I}(\mathcal{E}(\mathbf{x}_i) = \mathbf{x}_j \wedge \mathbf{x}_j \in \mathcal{R}(\mathbf{x}_i)) \text{ for all } \mathbf{x}_j \in \mathcal{X}. \quad (2)$$

Similarly as in previous work (Corbett-Davies et al., 2017; Valera et al., 2018; Kilbertus et al., 2019; Tabibian et al., 2020), we will assume that the decision maker is rational and aims to maximize the (immediate) utility $u(\pi, \gamma)$, which is the expected overall profit she obtains, *i.e.*,

$$u(\pi, \mathcal{A}) = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x}|\pi, \mathcal{A}), y \sim P(y|\mathbf{x}), d \sim \pi(\mathbf{x})} [yd(\mathbf{x}) - \gamma d(\mathbf{x})] = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x}|\pi, \mathcal{A})} [\pi(\mathbf{x})(P(y = 1|\mathbf{x}) - \gamma)], \quad (3)$$

where $\gamma \in (0, 1)$ is a given constant reflecting economic considerations of the decision maker. For example, in university admissions, the term $\pi(\mathbf{x})P(y = 1|\mathbf{x})$ is proportional to the expected number of students who are admitted and complete the program, the term $\pi(\mathbf{x})\gamma$ is proportional to the number of students who are admitted, and γ measures the cost of education in units of graduated students. As a direct consequence, given a feature value \mathbf{x}_i and a set of counterfactual explanations \mathcal{A} , we can conclude that, if $\mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} \neq \emptyset$, the decision maker will decide to provide the counterfactual explanation $\mathcal{E}(\mathbf{x}_i)$ that provides the largest utility gain under the assumption that individuals best respond, *i.e.*,

$$\mathcal{E}(\mathbf{x}_i) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{A} \cap \mathcal{R}(\mathbf{x}_i)} P(y|\mathbf{x}) \text{ for all } \mathbf{x}_i \in \mathcal{X} \setminus \mathcal{P}_\pi \text{ such that } \mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} \neq \emptyset, \quad (4)$$

and, if $\mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} = \emptyset$, we arbitrarily assume that $\mathcal{E}(\mathbf{x}_i) = \operatorname{argmin}_{\mathbf{x} \in \mathcal{A}} c(\mathbf{x}_i, \mathbf{x})$ ⁵.

²Without loss of generality, we assume each feature takes n different values.

³In practice, individuals with initial feature values \mathbf{x}_i such that $\pi(\mathbf{x}) = 1$ may not receive any explanation since they are guaranteed to receive a positive decision.

⁴In practice, the cost for each pair of feature values may be given by a parameterized function.

⁵Note that, if $\mathcal{A} \cap \mathcal{R}(\mathbf{x}_i) = \emptyset$, the individual’s best response is to keep her initial feature value \mathbf{x}_i and thus any choice of counterfactual explanation $\mathcal{E}(\mathbf{x}_i)$ leads to the same utility.

Given the above preliminaries, our goal is to help the decision maker to first find the optimal set of counterfactual explanations \mathcal{A} for a pre-defined policy in Section 3 and then both the optimal policy π and set of counterfactual explanations \mathcal{A} in Section 4.

Remarks. Given an individual with initial feature value \mathbf{x} , one may think that, by providing the counterfactual explanation $\mathcal{E}(\mathbf{x})$ that gives the largest utility gain, the decision maker is not acting in the individual’s best interest but rather selfishly. This is because there may exist another counterfactual explanation $\mathbf{x}' \in \mathcal{A}$ with lower cost for the individual, *i.e.*, $c(\mathbf{x}, \mathbf{x}') \leq c(\mathbf{x}, \mathcal{E}(\mathbf{x}))$. However, in our work, we argue that the provided counterfactual explanations help the individual to achieve a greater self-improvement and this is likely to result in a superior long-term well-being. For example, if a bank explains to a customer that she will receive the loan she is applying for if she reduces her credit card debt by 10%, rather than 20%, even though the corresponding feature values are both within the region of adaptation $\mathcal{R}(\mathbf{x})$ of her original features \mathbf{x} , the customer will be more likely to default and this will very negatively impact her long-term well-being.

As argued very recently (Miller et al., 2019; Tabibian et al., 2020), due to Goodhart’s law, the conditional probability $P(y | \mathbf{x})$ may change after individuals (best)-respond if the features \mathbf{x} are noncausal. Moreover, Miller et al. (2019) have argued that (best)-responses to noncausal and causal features correspond to gaming and improvement, respectively. In this work, for simplicity, we assume that $P(y | \mathbf{x})$ does not change, however, it would be very interesting to lift this assumption in future work.

3 Finding the optimal counterfactual explanations for a policy

In this section, our goal is to find the optimal set of counterfactual explanations \mathcal{A}^* for a pre-defined policy π , *i.e.*,

$$\mathcal{A}^* = \underset{\mathcal{A} \subseteq \mathcal{P}_\pi : |\mathcal{A}| \leq k}{\operatorname{argmax}} u(\pi, \mathcal{A}), \quad (5)$$

where k is the maximum number of counterfactual explanations the decision maker is willing to provide to the population to balance the right to explanation with trade secrets. As it will become clearer in the experimental evaluation in Section 5, our results may persuade decision makers to be transparent about their decision policies, something they are typically reluctant to be despite the increasing legal requirements, since we show that transparency increases the utility of the policies. Moreover, throughout this section, we will assume that the decision maker picking the pre-defined policy is rational⁶ and the policy is outcome monotonic^{7,8} (Tabibian et al., 2020). Outcome monotonicity just implies that, the higher an individual’s outcome $P(y = 1 | \mathbf{x})$, the higher their chances of receiving a positive decision $\pi(\mathbf{x})$.

Unfortunately, using a novel reduction of the Set Cover problem (Karp, 1972), the following theorem reveals that we cannot expect to find the optimal set of counterfactual explanations in polynomial time:

Theorem 1 *The problem of finding the optimal set of counterfactual explanations that maximizes utility under a cardinality constraint is NP-Hard.*

Proof Consider an instance of the Set Cover problem with a set of elements $\mathcal{U} = \{u_1, \dots, u_n\}$ and a collection $\mathcal{S} = \{\mathcal{S}_1, \dots, \mathcal{S}_m\} \subseteq 2^{\mathcal{U}}$ such that $\bigcup_{i \in [m]} \mathcal{S}_i = \mathcal{U}$. In the decision version of the problem, given a constant k , we need to answer the question whether there are at most k sets from the collection \mathcal{S} such that their union is equal to \mathcal{U} or not. With the following procedure, we show that any instance of that problem can be transformed to an instance of the problem of finding the optimal set of counterfactual explanations, defined in Eq. 5, in polynomial time.

Consider $n + m$ feature values corresponding to the n elements of \mathcal{U} and the m sets of \mathcal{S} . Moreover, denote the first n feature values as $\mathbf{x}_{u_1}, \dots, \mathbf{x}_{u_n}$ and the remaining m as $\mathbf{x}_{\mathcal{S}_1}, \dots, \mathbf{x}_{\mathcal{S}_m}$. We set the decision maker’s parameter γ to some positive constant less than 1. Then, we set the outcome probabilities $P(y = 1 | \mathbf{x}_{u_i}) = \gamma \forall i \in [n]$ and $P(y = 1 | \mathbf{x}_{\mathcal{S}_i}) = 1 \forall i \in [m]$ and the policy values $\pi(\mathbf{x}_{u_i}) = 0 \forall i \in [n]$ and $\pi(\mathbf{x}_{\mathcal{S}_i}) = 1 \forall i \in [m]$. This way, the portion of utility the decision-maker obtains from the first n feature values is zero, while the portion of utility she obtains from the remaining m is proportional to $1 - \gamma$.

⁶Note that, if the decision maker is rational and her goal is to maximize the utility, as defined in Eq. 3, then, for all $\mathbf{x} \in \mathcal{X}$ such that $P(y = 1 | \mathbf{x}) < \gamma$, it holds that $\pi(\mathbf{x}) = 0$.

⁷A policy π is called outcome monotonic if $P(y = 1 | \mathbf{x}_i) \geq P(y = 1 | \mathbf{x}_j) \Leftrightarrow \pi(\mathbf{x}_i) \geq \pi(\mathbf{x}_j) \forall \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$.

⁸If the policy π is deterministic, our results also hold for non outcome monotonic policies.




Figure 1: Consider that $\mathcal{U} = \{u_1, u_2\}$ and $\mathcal{S} = \{S_1, S_2\}$ with $S_1 = \{u_1, u_2\}$, $S_2 = \{u_2\}$. The red feature values have initial population $P(\mathbf{x}) = 1/2$, $\pi(\mathbf{x}) = 0$ and $P(y = 1 | \mathbf{x}) = \gamma$ while for the green feature values it is $P(\mathbf{x}) = 0$, $\pi(\mathbf{x}) = 1$ and $P(y = 1 | \mathbf{x}) = 1$. The edges represent the cost between feature values corresponding to sets and their respective elements while all the non-visible pairwise costs are equal to 2.

Regarding the cost function, we set $c(\mathbf{x}_{u_i}, \mathbf{x}_{S_j}) = 0 \forall (\mathbf{x}_{u_i}, \mathbf{x}_{S_j}) : u_i \in S_j$, $c(\mathbf{x}_{u_i}, \mathbf{x}_{u_i}) = 0 \forall i \in [n]$, and all the remaining values of the cost function to 2. Finally, we set the initial feature value distribution to $P(\mathbf{x}_{u_i}) = \frac{1}{n} \forall i \in [n]$ and $P(\mathbf{x}_{S_i}) = 0 \forall i \in [m]$. A toy example of this transformation is presented in Figure 1.

In this setting, it is easy to observe that an individual with initial feature value \mathbf{x}_{u_i} is always rejected at first and has the ability to move to a new feature value \mathbf{x}_{S_j} recommended to her iff $c(\mathbf{x}_{u_i}, \mathbf{x}_{S_j}) \leq 1 \Leftrightarrow u_i \in S_j$. Also, we can easily see that the transformation of instances can be done in $\mathcal{O}((m+n)^2)$ time.

Now, assume there exists an algorithm that optimally solves the problem of finding the optimal set of counterfactual explanations in polynomial time. Given the aforementioned instance and a maximum number of counterfactual explanations k , the utility $u(\pi, \mathcal{A})$ achieved by the set of counterfactual explanations \mathcal{A} the algorithm returns can fall into one of the following two cases:

1. $u(\pi, \mathcal{A}) = 1 - \gamma$. This can happen only if all individuals, according to the induced distribution $P(\mathbf{x} | \pi, \mathcal{A})$, have moved to some of the feature values \mathbf{x}_{S_j} , *i.e.*, for all \mathbf{x}_{u_i} with $i \in [n]$, there exists \mathbf{x}_{S_j} with $j \in [m]$ such that $\mathbf{x}_{S_j} \in \mathcal{A} \wedge c(\mathbf{x}_{u_i}, \mathbf{x}_{S_j}) \leq 1$ with $|\mathcal{A}| \leq k$. As a consequence, if we define $\mathcal{S}' = \{S_j : \mathbf{x}_{S_j} \in \mathcal{A}\}$, it holds that for all u_i with $i \in [n]$, there exists S_j with $j \in [m]$ such that $S_j \in \mathcal{S}' \wedge u_i \in S_j$ and therefore \mathcal{S}' is a set cover with $|\mathcal{S}'| = |\mathcal{A}| \leq k$.
2. $u(\pi, \mathcal{A}) < 1 - \gamma$. This can happen only if every possible set of k counterfactual explanations leaves the individuals of at least one feature value \mathbf{x}_{u_i} with a best-response of not following the counterfactual explanation they were given, *i.e.*, for all $\mathcal{A} \subseteq \mathcal{P}_\pi$ such that $|\mathcal{A}| \leq k$, there exists \mathbf{x}_{u_i} with $i \in [n]$ such that, for all $\mathbf{x}_{S_j} \in \mathcal{A}$, it holds that $c(\mathbf{x}_{u_i}, \mathbf{x}_{S_j}) > 1$. Equivalently, it holds that for all $\mathcal{S}' \subseteq \mathcal{S}$ such that $|\mathcal{S}'| \leq k$, there exists u_i with $i \in [n]$ such that for all $S_j \in \mathcal{S}'$, it holds that $u_i \notin S_j$ and therefore there does not exist a set cover of size less or equal than k .

The above directly implies that we can have a decision about any instance of the Set Cover problem in polynomial time, which is a contradiction unless $P = NP$. This concludes the reduction and proves that the problem of finding the optimal set of counterfactual explanations for a given policy is NP-Hard. ■

Even though Theorem 1 is a negative result, we will now show that the objective function in Eq. 5 satisfies a set of desirable properties, *i.e.*, non-negativity, monotonicity and submodularity⁹, which allow a standard greedy algorithm to enjoy a $(1 - 1/e)$ approximation ratio at solving the problem. To this aim, with a slight abuse of notation, we first express the objective function as a set function $f(\mathcal{A}) = u(\pi, \mathcal{A})$, which takes values over the ground set of counterfactual explanations, \mathcal{P}_π . Then, we have the following proposition:

Proposition 2 *The function f is non-negative, submodular and monotone.*

Proof It readily follows that the function f is non-negative from the fact that, if the decision maker is rational, it holds that $\pi(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{X}$ such that $P(y = 1 | \mathbf{x}) < \gamma$.

⁹A function $f : 2^{\mathcal{X}} \rightarrow \mathbb{R}$ is submodular if for every $\mathcal{A}, \mathcal{B} \subseteq \mathcal{X} : \mathcal{A} \subseteq \mathcal{B}$ and $x \in \mathcal{X} \setminus \mathcal{B}$ it holds that $f(\mathcal{A} \cup \{x\}) - f(\mathcal{A}) \geq f(\mathcal{B} \cup \{x\}) - f(\mathcal{B})$.

Now, consider two sets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{P}_\pi : \mathcal{A} \subseteq \mathcal{B}$ and a feature value $\mathbf{x} \in \mathcal{P}_\pi \setminus \mathcal{B}$. Also, let $\mathcal{E}_\mathcal{S}(\mathbf{x}_i)$ be the counterfactual explanation given to the individuals with initial feature value \mathbf{x}_i under a set of counterfactual explanations \mathcal{S} . It is easy to see that the marginal difference $f(\mathcal{S} \cup \{\mathbf{x}\}) - f(\mathcal{S})$ can only be affected by individuals with initial features \mathbf{x}_i such that $\mathbf{x}_i \notin \mathcal{P}_\pi$, $\mathbf{x} \in \mathcal{R}(\mathbf{x}_i)$ and $\mathbf{x} = \mathcal{E}_{\mathcal{S} \cup \{\mathbf{x}\}}(\mathbf{x}_i)$. Moreover, we can divide all of these individuals into two cases:

1. $\mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} = \emptyset$: in this case, the addition of \mathbf{x} to \mathcal{A} causes a change in their best-response from \mathbf{x}_i to \mathbf{x} contributing to the marginal difference of f by a factor $P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - \gamma - \pi(\mathbf{x}_i)(P(y = 1 | \mathbf{x}_i) - \gamma)]$. However, considering the marginal difference of f under the set of counterfactual explanations \mathcal{B} , three subcases are possible:

- (a) $\mathcal{E}_\mathcal{B}(\mathbf{x}_i) \in \mathcal{R}(\mathbf{x}_i) \wedge P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) > P(y = 1 | \mathbf{x})$: the contribution to the marginal difference of f is zero.
- (b) $\mathcal{E}_\mathcal{B}(\mathbf{x}_i) \in \mathcal{R}(\mathbf{x}_i) \wedge P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) \leq P(y = 1 | \mathbf{x})$: the contribution to the marginal difference of f is $P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i))]$. Since π is outcome monotonic, $\mathcal{E}_\mathcal{B}(\mathbf{x}_i) \in \mathcal{P}_\pi$ and $\mathbf{x}_i \notin \mathcal{P}_\pi$, it holds that

$$\begin{aligned} P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) &\geq P(y = 1 | \mathbf{x}_i) \Rightarrow \\ P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) - \gamma &\geq P(y = 1 | \mathbf{x}_i) - \gamma > \pi(\mathbf{x}_i)[P(y = 1 | \mathbf{x}_i) - \gamma]. \end{aligned}$$

Therefore, it readily follows that

$$P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i))] < P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - \gamma - \pi(\mathbf{x}_i)(P(y = 1 | \mathbf{x}_i) - \gamma)].$$

- (c) $\mathcal{R}(\mathbf{x}_i) \cap \mathcal{B} = \emptyset$: the contribution to the marginal difference of f is $P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - \gamma - \pi(\mathbf{x}_i)(P(y = 1 | \mathbf{x}_i) - \gamma)]$.

2. $\mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} \neq \emptyset \wedge P(y = 1 | \mathbf{x}) > P(y = 1 | \mathcal{E}_\mathcal{A}(\mathbf{x}_i))$: In this case, the addition of \mathbf{x} to \mathcal{A} causes a change in their best-response from $\mathcal{E}_\mathcal{A}(\mathbf{x}_i)$ to \mathbf{x} contributing to the marginal difference of f by a factor $P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_\mathcal{A}(\mathbf{x}_i))]$. Considering the marginal difference of f under the set of counterfactual explanations \mathcal{B} , two subcases are possible:

- (a) $\mathcal{E}_\mathcal{B}(\mathbf{x}_i) \in \mathcal{R}(\mathbf{x}_i) \wedge P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) > P(y = 1 | \mathbf{x})$: the contribution to the marginal difference of f is zero.
- (b) $\mathcal{E}_\mathcal{B}(\mathbf{x}_i) \in \mathcal{R}(\mathbf{x}_i) \wedge P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) \leq P(y = 1 | \mathbf{x})$. Then, the contribution of those individuals to the marginal difference of f is $P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i))]$. Since $\mathcal{A} \subseteq \mathcal{B}$ and $\mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} \neq \emptyset$, it readily follows that

$$\begin{aligned} P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i)) &\geq P(y = 1 | \mathcal{E}_\mathcal{A}(\mathbf{x}_i)) \Rightarrow \\ P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_\mathcal{A}(\mathbf{x}_i))] &\geq P(\mathbf{x}_i)[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_\mathcal{B}(\mathbf{x}_i))]. \end{aligned}$$

Finally, because $\mathcal{A} \subseteq \mathcal{B}$, we can conclude that $f(\mathcal{B} \cup \{\mathbf{x}\}) - f(\mathcal{B}) \neq 0 \Rightarrow f(\mathcal{A} \cup \{\mathbf{x}\}) - f(\mathcal{A}) \neq 0$ and therefore the aforementioned cases are sufficient. Combining all cases, we can see that the contribution of each individual to the marginal difference of f is always greater or equal under the set of counterfactual explanations \mathcal{A} than under the set of counterfactual explanations \mathcal{B} . As a direct consequence, it follows that f is submodular. Additionally, we can easily see that this contribution is always greater or equal than zero, leading to the conclusion that f is also monotone. \blacksquare

The above result directly implies that the standard greedy algorithm by Nemhauser et al. (1978) will find a solution \mathcal{A} to the problem such that $f(\mathcal{A}) \geq (1 - 1/e)f(\mathcal{A}^*)$, where \mathcal{A}^* is the optimal set of counterfactual explanations.

Algorithm 1 summarizes the overall procedure. It starts from a solution set $\mathcal{A} = \emptyset$ and it iteratively adds to \mathcal{A} the counterfactual explanation $\mathbf{x} \in \mathcal{P}_\pi \setminus \mathcal{A}$ that provides the maximum marginal difference $f(\mathcal{A} \cup \{\mathbf{x}\}) - f(\mathcal{A})$. Finally, note that, while the greedy algorithm has a complexity of $\mathcal{O}(km)$ under oracle access to the utility function f , it is easy to show that, given a set of counterfactual explanations, it takes $\mathcal{O}(km)$ to compute the utility $f(\mathcal{S})$. As a consequence, for our problem, the greedy algorithm has a complexity of $\mathcal{O}(k^2m^2)$

ALGORITHM 1: Greedy algorithm

Input: Ground set of counterfactual explanations \mathcal{P}_π , maximum number of counterfactual explanations k and utility function f

Output: Set of counterfactual explanations \mathcal{A}

- 1: $\mathcal{A} \leftarrow \emptyset$
 - 2: **while** $|\mathcal{A}| \leq k$ **do**
 - 3: $\mathbf{x}^* \leftarrow \operatorname{argmax}_{\mathbf{x} \in \mathcal{P}_\pi \setminus \mathcal{A}} f(\mathcal{A} \cup \{\mathbf{x}\}) - f(\mathcal{A})$
 - 4: $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{x}^*\}$
 - 5: **end while**
 - 6: **return** \mathcal{A}
-

Remarks on fair decision making. In many cases, decision makers may like to ensure that individuals across the whole spectrum of the population are incentivized to self-improve. For example, in a loan scenario, the bank may use age group as a feature to estimate the probability that a customer repays the loan, however, it may like to deploy a decision policy that incentivizes individuals across all age groups in order to improve the financial situation of all. To this aim, rather than the cardinality constraint $|\mathcal{A}| < k$, the decision maker could incorporate a matroid constraint into the problem formulation. Formally, consider disjoint sets $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_l$ such that $\bigcup_i \mathcal{X}_i = \mathcal{X}$ and integers d_1, d_2, \dots, d_l such that $k = \sum_i d_i$. Then, a partition matroid is the collection of sets $\{S \subseteq \mathcal{X} : |S \cap \mathcal{X}_i| \leq d_i \forall i \in [l]\}$. In the loan example, the decision maker could search for a set of counterfactual explanations \mathcal{A} within a partition matroid where each one of the \mathcal{X}_i 's corresponds to the feature values covered by each age group and $d_i = k/l \forall i \in [l]$. This way, the set of counterfactual explanations \mathcal{A} would include explanations for every age group. In this case, the decision maker could rely on a variety of polynomial time algorithms with global guarantees for submodular function maximization under matroid constraints, *e.g.*, the algorithm by Calinescu et al. (2011).

4 Finding the optimal policy and counterfactual explanations

In this section, our goal is to jointly find the optimal decision policy and set of counterfactual explanations \mathcal{A}^* , *i.e.*,

$$\pi^*, \mathcal{A}^* = \operatorname{argmax}_{(\pi, \mathcal{A}): \mathcal{A} \subseteq \mathcal{P}_\pi \wedge |\mathcal{A}| \leq k} u(\pi, \mathcal{A}) \quad (6)$$

where, similarly as in the previous section, k is the maximum number of counterfactual explanations the decision maker is willing to provide to the population to balance the right to explanation with trade secrets. By jointly optimizing both the decision policy and the counterfactual explanations, Figure 2 shows that we may obtain an additional gain in terms of utility in comparison with just optimizing for the set of counterfactual explanations given the optimal decision policy in a non-strategic setting. Moreover, as we will show in the experimental evaluation in Section 5, this additional gain will be significant.

Similarly as in Section 3, we cannot expect to find the optimal policy and set of counterfactual explanations in polynomial time. More specifically, we have the following negative result:

Theorem 3 *The problem of jointly finding both the optimal policy and the set of counterfactual explanations that maximize utility under a cardinality constraint is NP-hard.*

Proof The theorem can be easily proven by using Proposition 4 and slightly extending the proof of Theorem 1. ■

However, while the problem of finding both the policy and the set of counterfactual explanations appears significantly more challenging than the problem of finding just the set of counterfactual explanations given a pre-defined policy (refer to Eq. 5), the following proposition shows that the problem is not inherently *harder*. More specifically, for each possible set of counterfactual explanations, it shows that the policy that maximizes the utility can be easily computed:

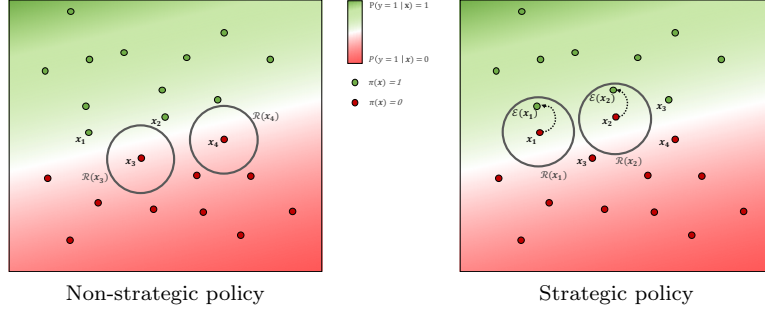


Figure 2: Jointly optimizing the decision policy and the counterfactual explanations can offer additional gains. The left panel shows the optimal (deterministic) decision policy π under non-strategic behavior, as given by Eq. 8. Here, there does not exist a set of counterfactual explanations $\mathcal{A} \in \mathcal{P}_\pi$ that increases the utility of the policy. This happens because the area of adaption of \mathbf{x}_3 and \mathbf{x}_4 does not include any feature value that receives a positive decision. The right panel shows the decision policy and counterfactual explanations that are (jointly) optimal in terms of utility, as given by Eq. 6. Here, the individuals with feature values \mathbf{x}_1 and \mathbf{x}_2 receive $\mathcal{E}(\mathbf{x}_1)$ and $\mathcal{E}(\mathbf{x}_2)$, respectively, as counterfactual explanations. Since these explanations are within their areas of adaptation $\mathcal{R}(\mathbf{x}_1)$ and $\mathcal{R}(\mathbf{x}_2)$, they change their initial feature values in order to receive a positive decision. Both panels share the same $P(\mathbf{x})$, $P(y|\mathbf{x})$ and $c(\mathbf{x}, \mathbf{x}')$.

Proposition 4 Given a set of counterfactual explanations $\mathcal{A} \subseteq \mathcal{Y} := \{\mathbf{x} \in \mathcal{X} : P(y|\mathbf{x}) \geq \gamma\}$ ¹⁰, the policy $\pi_{\mathcal{A}}^* = \operatorname{argmax}_{\pi: \mathcal{A} \subseteq \mathcal{P}_\pi} u(\pi, \mathcal{A})$ that maximizes the utility is deterministic and can be found in polynomial time, i.e.,

$$\pi_{\mathcal{A}}^*(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in \mathcal{A} \vee \{P(y=1|\mathbf{x}) \geq \gamma \wedge c(\mathbf{x}, \mathbf{x}') > 1 \forall \mathbf{x}' \in \mathcal{A}\} \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Proof By definition, since $\mathcal{A} \subseteq \mathcal{P}_{\pi_{\mathcal{A}}^*}$, it readily follows that $\pi_{\mathcal{A}}^*(\mathbf{x}) = 1$ for all $\mathbf{x} \in \mathcal{A}$. To find the remaining values of the decision policy, we first observe that, for each $\mathbf{x} \notin \mathcal{A}$, the value of the decision policy $\pi_{\mathcal{A}}^*(\mathbf{x})$ does not affect the best-responses of the individuals with initial feature values $\mathbf{x}' \neq \mathbf{x}$. As a result, we can just set $\pi_{\mathcal{A}}^*(\mathbf{x})$ for all $\mathbf{x} \notin \mathcal{A}$ independently for each feature value \mathbf{x} such that the best-response of the respective individuals is the one that contributes maximally to the overall utility.

First, it is easy to see that, for all $\mathbf{x} \notin \mathcal{A}$ such that $P(y=1|\mathbf{x}) < \gamma$, we should set $\pi_{\mathcal{A}}^*(\mathbf{x}) = 0$. Next, consider the feature values $\mathbf{x} \notin \mathcal{A}$ such that $P(y=1|\mathbf{x}) \geq \gamma$. Here, we distinguish two cases. If there exists $\mathbf{x}' \in \mathcal{A}$ such that $c(\mathbf{x}, \mathbf{x}') \leq 1 \wedge P(y=1|\mathbf{x}') > P(y=1|\mathbf{x})$, then, if the individuals move to that \mathbf{x}' , the corresponding contribution to the utility will be higher. Moreover, the value of the decision policy that maximizes their region of adaption (and thus increases their chances of moving to \mathbf{x}') is clearly $\pi_{\mathcal{A}}^*(\mathbf{x}) = 0$. If there does not exist $\mathbf{x}' \in \mathcal{A}$ such that $c(\mathbf{x}, \mathbf{x}') \leq 1 \wedge P(y=1|\mathbf{x}') > P(y=1|\mathbf{x})$, then, the contribution of the corresponding individuals to the utility will be higher if they keep their initial feature values. Moreover, the value of the decision policy that will maximize this contribution will be clearly $\pi_{\mathcal{A}}^*(\mathbf{x}) = 1$. Finally, note that, to set all the values of the decision policy, we only need to perform $\mathcal{O}(km)$ comparisons. ■

The above result reveals that, in contrast with the non strategic setting, the optimal policy given a set of counterfactual explanations is not a deterministic threshold rule with a single threshold (Corbett-Davies et al., 2017; Valera et al., 2018), i.e.,

$$\pi(\mathbf{x}) = \begin{cases} 1 & \text{if } P(y=1|\mathbf{x}) \geq \gamma \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

but rather a more conservative deterministic decision policy that does not depend only on the outcome $P(y=1|\mathbf{x})$ and γ but also on the cost individuals pay to change features. Moreover, we can build up on the above result to prove that the problem of finding the optimal decision policy and set of counterfactual

¹⁰Since the decision maker is rational, she will never provide an explanation that contributes negatively to her utility.

explanations can be reduced to maximizing a non-monotone submodular function. To this aim, let $\pi_{\mathcal{A}}^*$ be the optimal policy induced by a given set of counterfactual explanations \mathcal{A} , as in Proposition 4, and define the set function $h(\mathcal{A}) = u(\pi_{\mathcal{A}}^*, \mathcal{A})$ over the ground set \mathcal{Y} . Then, we have the following proposition:

Proposition 5 *The function h is non-negative and submodular.*

Proof It readily follows that the function h is non-negative from the fact that, if the decision maker is rational, $\pi(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{X}$ such that $P(y = 1 | \mathbf{x}) < \gamma$.

Next, consider two sets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{Y}$ such that $\mathcal{A} \subseteq \mathcal{B}$ and a feature value $\mathbf{x} \in \mathcal{Y} \setminus \mathcal{B}$. Also, let $\mathcal{E}_{\mathcal{S}}(\mathbf{x}_i)$ be the counterfactual explanation given to the individuals with initial feature value \mathbf{x}_i under a set of counterfactual explanations \mathcal{S} . Then, it is clear that the marginal difference $h(\mathcal{S} \cup \{\mathbf{x}\}) - h(\mathcal{S})$ only depends on individuals with initial features \mathbf{x}_i such that either $1 - c(\mathbf{x}_i, \mathbf{x}) \geq 0$ and $\mathbf{x} = \mathcal{E}_{\mathcal{S} \cup \{\mathbf{x}\}}(\mathbf{x}_i)$ or $\mathbf{x}_i = \mathbf{x}$. Moreover, if $1 - c(\mathbf{x}_i, \mathbf{x}) \geq 0$ and $\mathbf{x} = \mathcal{E}_{\mathcal{S} \cup \{\mathbf{x}\}}(\mathbf{x}_i)$, the contribution to the marginal difference is positive and, if $\mathbf{x}_i = \mathbf{x}$, the contribution to the marginal difference is negative.

Consider first the individuals with initial features \mathbf{x}_i such that $1 - c(\mathbf{x}_i, \mathbf{x}) \geq 0$ and $\mathbf{x} = \mathcal{E}_{\mathcal{A} \cup \{\mathbf{x}\}}(\mathbf{x}_i)$. We can divide all of these individuals into three cases:

1. $\pi_{\mathcal{B}}(\mathbf{x}_i) = 0$: in this case, $\mathbf{x}_i \notin \mathcal{B}$ and the individuals change their best-response from $\mathcal{E}_{\mathcal{B}}(\mathbf{x}_i)$ to \mathbf{x} . Moreover, under the set of counterfactual explanations \mathcal{A} , their best-response is either \mathbf{x}_i or $\mathcal{E}_{\mathcal{A}}(\mathbf{x}_i)$ and it changes to \mathbf{x} . Then, using a similar argument as in the proof of proposition 2, we can conclude that the contribution of the individuals to the marginal difference is greater or equal under the set of counterfactual explanations \mathcal{A} than under \mathcal{B} .
2. $\pi_{\mathcal{B}}(\mathbf{x}_i) = 1 \wedge \pi_{\mathcal{A}}(\mathbf{x}_i) = 0$: in this case, $\mathbf{x}_i \notin \mathcal{A}$ and $\mathbf{x}_i \in \mathcal{B}$. Therefore, under the set of counterfactual explanations \mathcal{A} , the individuals' best-response changes from $\mathcal{E}_{\mathcal{A}}(\mathbf{x}_i)$ to \mathbf{x} and there is a positive contribution to the marginal difference while, under \mathcal{B} , the individuals' best response does not change and the contribution to the marginal difference is zero.
3. $\pi_{\mathcal{B}}(\mathbf{x}_i) = 1 \wedge \pi_{\mathcal{A}}(\mathbf{x}_i) = 1$: in this case, $\mathbf{x}_i \notin \mathcal{B}$. Therefore, the best-response changes from \mathbf{x}_i to \mathbf{x} under both sets of counterfactual explanations and there is an equal positive contribution to the marginal difference.

Now, consider the individuals with initial features \mathbf{x}_i such that $\mathbf{x}_i = \mathbf{x}$. We can divide all of these individuals also into three cases:

1. $\pi_{\mathcal{A}}(\mathbf{x}) = \pi_{\mathcal{B}}(\mathbf{x}) = 0$: in this case, under both sets of counterfactual explanations, the counterfactual explanation \mathbf{x} changes the value of the decision policy to $\pi_{\mathcal{A} \cup \{\mathbf{x}\}}(\mathbf{x}) = \pi_{\mathcal{B} \cup \{\mathbf{x}\}}(\mathbf{x}) = 1$. Moreover, the contribution to the marginal difference is less negative under the set of counterfactual explanations \mathcal{A} than under \mathcal{B} since $P(y = 1 | \mathcal{E}_{\mathcal{A}}(\mathbf{x})) \leq P(y = 1 | \mathcal{E}_{\mathcal{B}}(\mathbf{x}))$ and thus $P(\mathbf{x})[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_{\mathcal{A}}(\mathbf{x}))] \geq P(\mathbf{x})[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_{\mathcal{B}}(\mathbf{x}))]$.
2. $\pi_{\mathcal{A}}(\mathbf{x}) = 1 \wedge \pi_{\mathcal{B}}(\mathbf{x}) = 0$: in this case, under the set of counterfactual explanations \mathcal{A} , the individuals' best response does not change and thus the contribution to the marginal difference is zero and, under the set of counterfactual explanations \mathcal{B} , their best-response changes from $\mathcal{E}_{\mathcal{B}}(\mathbf{x})$ to \mathbf{x} and thus there is a negative contribution to the marginal difference *i.e.*, $P(\mathbf{x})[P(y = 1 | \mathbf{x}) - P(y = 1 | \mathcal{E}_{\mathcal{B}}(\mathbf{x}))] < 0$.
3. $\pi_{\mathcal{A}}(\mathbf{x}) = \pi_{\mathcal{B}}(\mathbf{x}) = 1$: in this case, under both sets of counterfactual explanations, the individuals' best response does not change and thus the contribution to the marginal difference is zero.

As a direct consequence of the above observations, it readily follows that $h(\mathcal{A} \cup \{\mathbf{x}\}) - h(\mathcal{A}) \geq h(\mathcal{B} \cup \{\mathbf{x}\}) - h(\mathcal{B})$ and therefore the function h is submodular. ■

However, in contrast with Section 3, the function h is non-monotone since it can happen that the negative marginal contribution exceeds the positive one. For example, consider the following instance of the problem, where $\mathbf{x} \in \{1, 2, 3\}$ with $\gamma = 0.1$:

$$\begin{aligned} P(\mathbf{x}) &= 0.1 \mathbb{I}(\mathbf{x} = 1) + 0.8 \mathbb{I}(\mathbf{x} = 2) + 0.1 \mathbb{I}(\mathbf{x} = 3), \\ P(y = 1 | \mathbf{x}) &= 1.0 \mathbb{I}(\mathbf{x} = 1) + 0.5 \mathbb{I}(\mathbf{x} = 2) + 0.4 \mathbb{I}(\mathbf{x} = 3), \end{aligned}$$

ALGORITHM 2: Randomized algorithm

Input: Ground set of counterfactual explanations \mathcal{Y} , maximum number of counterfactual explanations k and utility function f

Output: Set of counterfactual explanations \mathcal{A}

```
1:  $\mathcal{A} \leftarrow \emptyset$ 
2: while  $|\mathcal{A}| \leq k$  do
3:    $\mathcal{B} \leftarrow \text{GetTopK}(\mathcal{Y}, \mathcal{A}, f)$ 
4:    $\mathbf{x}^* \sim \mathcal{B}$ 
5:    $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{x}^*\}$ 
6: end while
7: return  $\mathcal{A}$ 
```

and

$$c(\mathbf{x}_i, \mathbf{x}_j) = \begin{bmatrix} 0.0 & 0.2 & 0.3 \\ 0.3 & 0.0 & 0.7 \\ 0.4 & 0.5 & 0.0 \end{bmatrix}.$$

Assume there is a set of counterfactual explanations $\mathcal{A} = \{1\}$. Then, the optimal policy is given by $\pi_{\mathcal{A}}^*(1) = 1, \pi_{\mathcal{A}}^*(2) = 0, \pi_{\mathcal{A}}^*(3) = 0$ inducing a movement from feature values 2, 3 to feature value 1, giving a utility equal to 0.9. Now, add $\mathbf{x} = 2$ to the set of counterfactual explanations *i.e.*, $\mathcal{A} = \{1, 2\}$. Then, the optimal policy is given by $\pi_{\mathcal{A}}^*(1) = 1, \pi_{\mathcal{A}}^*(2) = 1, \pi_{\mathcal{A}}^*(3) = 0$ inducing a movement from feature value 3 to feature value 1, giving a lower utility, equal to 0.5. Therefore, the function h is non-monotone.

Fortunately, there exist efficient algorithms with global approximation guarantees for maximizing a non-monotone submodular function under cardinality constraints. For example, Buchbinder et al. (2014) have proposed a randomized polynomial time algorithm that can find a solution \mathcal{A} such that $h(\mathcal{A}) \geq (1/e)h(\mathcal{A}^*)$, where \mathcal{A}^* and $\pi_{\mathcal{A}^*}^*$ are the optimal set of counterfactual explanations and decision policy, respectively. Algorithm 2 summarizes the whole procedure, which is just a randomized variation of the standard greedy algorithm (Algorithm 1). It starts from a solution set $\mathcal{A} = \emptyset$ and it iteratively adds one counterfactual explanation $\mathbf{x} \in \mathcal{Y} \setminus \mathcal{A}$. However, instead of greedily choosing the element \mathbf{x} that provides the maximum marginal difference $h(\mathcal{A} \cup \{\mathbf{x}\}) - h(\mathcal{A})$, it sorts all the candidate elements with respect to their marginal difference (line 3) and picks one at random among the top k (line 4). Finally, note that, since the above randomized algorithm has a complexity of $\mathcal{O}(km)$, similarly as the standard greedy algorithm, we can use Proposition 4 to conclude that, for our problem, the algorithm has a complexity of $\mathcal{O}(k^2m^2)$.

Remarks. To enjoy a $1/e$ approximation guarantee, Algorithm 2 requires that there are $2k < m$ candidate feature values whose marginal contribution to any set is zero. In our problem, this can be trivially satisfied by adding $2k$ feature values \mathbf{x} to \mathcal{X} such that $P(y = 1 | \mathbf{x}) = \gamma$, $P(\mathbf{x}) = 0$ and $c(\mathbf{x}, \mathbf{x}_j) = c(\mathbf{x}_j, \mathbf{x}) = 2 \forall \mathbf{x}_j \in \mathcal{X}$. If the algorithm adds some of those counterfactual explanations to the set \mathcal{A} , it is easy to see that we can ignore them without causing any difference in utility or the individuals' best-responses.

5 Experiments

In this section, we perform a series of experiments on both synthetic and real data to evaluate Algorithms 1 and 2. To this aim, we compare the utility of the following decision policies and counterfactual explanations:

— *Black box*: decisions are taken by the optimal decision policy in the non-strategic setting, given by Eq. 8, and individuals do not receive any counterfactual explanations.

— *Minimum cost*: decisions are taken by the optimal decision policy in the non-strategic setting, given by Eq. 8, and individuals receive counterfactual explanations of minimum cost with respect to their initial feature values, similarly as in previous work (Tolomei et al., 2017; Karimi et al., 2019; Ustun et al., 2019). More specifically, we cast the problem of finding the set of counterfactual explanations as the minimization of the weighted average cost individuals pay to change their feature values to the

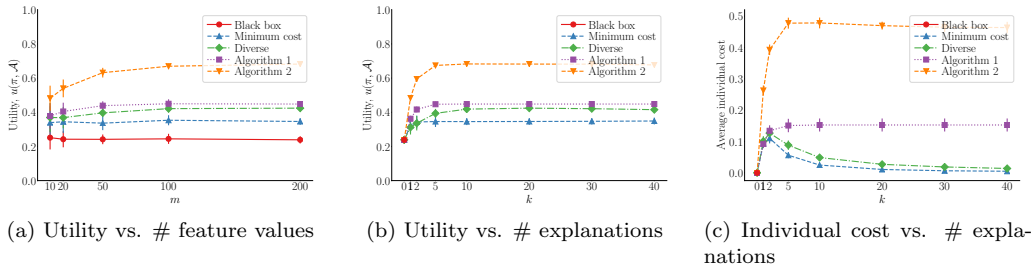


Figure 3: Results on synthetic data. Panels (a) and (b) show the utility achieved by six types of decision policies and counterfactual explanations against the total number of feature values m and the number of counterfactual explanations k , respectively. Panel (c) shows the average cost individuals had to pay to change from their initial features to the feature value of the counterfactual explanation they receive under the same five types of decision policies and counterfactual explanations. In Panel (a), we set $k = 0.1m$ and, in Panels (b) and (c), we set $m = 200$. In all panels, we repeat each experiment 20 times.

closest counterfactual explanation, *i.e.*,

$$\mathcal{A}_{mc} = \underset{\mathcal{A} \subseteq \mathcal{P}_\pi : |\mathcal{A}| \leq k}{\operatorname{argmin}} \sum_{\mathbf{x}_i \in \mathcal{X} \setminus \mathcal{P}_\pi} P(\mathbf{x}_i) \min_{\mathbf{x}_j \in \mathcal{A}} c(\mathbf{x}_i, \mathbf{x}_j),$$

and realize that this problem is a version of the k-median problem, which we can solve using a greedy heuristic (Solis-Oba, 2006).

— *Diverse*: decisions are taken by the optimal decision policy in the non-strategic setting, given by Eq. 8, and individuals receive a set of diverse counterfactual explanations of minimum cost with respect to their initial feature values, similarly as in previous work (Russell, 2019; Mothilal et al., 2020), *i.e.*,

$$\mathcal{A}_d = \underset{\mathcal{A} \subseteq \mathcal{P}_\pi : |\mathcal{A}| \leq k}{\operatorname{argmax}} \sum_{\mathbf{x}_i \in \mathcal{X} \setminus \mathcal{P}_\pi} P(\mathbf{x}_i) \mathbb{I}(\mathcal{R}(\mathbf{x}_i) \cap \mathcal{A} \neq \emptyset),$$

To solve the above problem, we realize it can be reduced to the weighted version of the maximum coverage problem, which can be solved using a well-known greedy approximation algorithm (Hochbaum and Pathria, 1998).

— *Algorithm 1*: decisions are taken by the optimal decision policy in the non-strategic setting, given by Eq. 8, and individuals receive counterfactual explanations given by Eq. 4, where \mathcal{A} is found using Algorithm 1.

— *Algorithm 2*: decisions are taken by the decision policy given by Eq. 7 and individuals receive counterfactual explanations given by Eq. 4, where \mathcal{A} is found using Algorithm 2.

5.1 Experiments on Synthetic Data

Experimental setup. For simplicity, we consider feature values $\mathbf{x} \in \{0, \dots, m-1\}$ and $P(\mathbf{x} = i) = p_i / \sum_j p_j$ where p_i is sampled from a Gaussian distribution $N(\mu = 0.5, \sigma = 0.1)$ truncated from below at zero. We also sample $P(y = 1 | \mathbf{x}) \sim U[0, 1]$, $c(\mathbf{x}_i, \mathbf{x}_j) \sim U[0, 1]$ for 50% of all pairs and $c(\mathbf{x}_i, \mathbf{x}_j) = 2$ for the rest. Finally, we set $\gamma = 0.3$.

Results. Figures 3(a,b) show the utility achieved by each of the decision policies and counterfactual explanations for several numbers of feature values m and counterfactual explanations k . We find several interesting insights: (i) the decision policies given by Eq. 7 and the counterfactual explanations found by Algorithm 2 beat all other alternatives by large margins across the whole spectrum, showing that jointly optimizing the decision policy and the counterfactual explanations offer clear additional gains; (ii) the counterfactual explanations found by Algorithms 1 and 2 provide higher utility gains as the number of feature values increases and thus the search space of counterfactual explanations becomes larger; and, (iii) a small number of counterfactual explanations is enough to provide significant gains in terms of utility with respect to the optimal decision policy without counterfactual explanations.

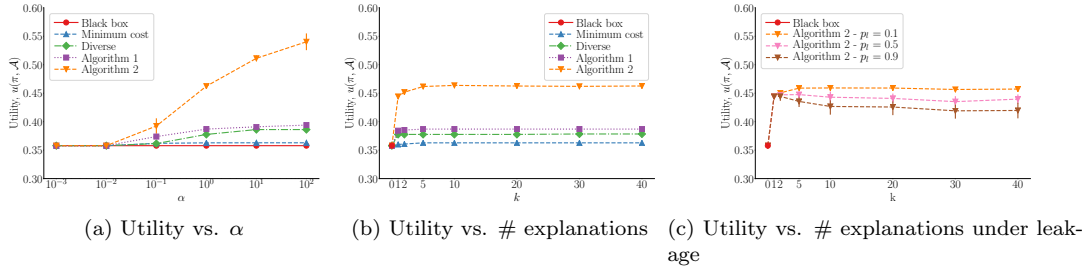


Figure 4: Results on LendingClub data. Panels (a) and (b) show the utility achieved by five types of decision policies and counterfactual explanations against the value of the parameter α , which controls how difficult it is to change feature values, and the number of counterfactual explanations k , respectively. Panel (c) shows the utility achieved by the same six types of decision policies and counterfactual explanations against the number of counterfactual explanations k for several values of the leakage probability p_l . In Panel (a), we set $k = 20$ and in Panels (b) and (c), we set $\alpha = 1$. In all panels, it is $m = 200$ and we repeat each experiment involving randomization 20 times.

Figure 3(c) shows the average cost individuals had to pay to change from their initial features to the feature value of the counterfactual explanation they receive. As one may have expected, the results show that, under the counterfactual explanations of minimum cost (Minimum cost and Diverse), the individuals invest less effort to change their initial features and the effort drops as the number of counterfactual explanations increases. In contrast, our methods incentivize the individuals to achieve the highest self-improvement, particularly when we jointly optimize the decision policy and the counterfactual explanations.

5.2 Experiments on Real Data

Experimental setup. We use a publicly available dataset¹¹ with information about all accepted and rejected loan applications in LendingClub during the 2007-2018 period. For each application, the dataset contains various demographic features about the applicant and for each accepted application it contains the current loan status (*e.g.*, Current, Late, Fully Paid), the latest payment information and the FICO scores.

Here, we follow the same pre-processing steps as in Tabibian et al. (2020), which we recall in what follows, for completeness. We use information about accepted applications to train a decision tree classifier (DTC) with $m = 200$ leaves. This classifier predicts whether an applicant fully pays a loan ($y = 1$) or defaults/has a charged-off debt ($y = 0$) on the basis of a set of *raw* features, *i.e.*, the loan amount, employment length, state, debt to income ratio, zip code and credit score (more information about raw features can be found in Appendix A.1). The classifier achieves a 90% accuracy, as estimated using 3-fold cross-validation. Then, for each (accepted or rejected) application, its unidimensional feature \mathbf{x} is set to be the leaf of the DTC where the application is mapped into, given the raw features, and the conditional probability $P(y = 1 | \mathbf{x})$ is approximated using the prediction of the DTC for the corresponding leaf. Then, we compute the cost between feature values $c(\mathbf{x}_i, \mathbf{x}_j)$ by comparing the raw features of the applications mapped to each leaf of the DTC, where we multiply the cost by a scaling factor $1/\alpha$ representing the hardness of changing features (Appendix A.2 provides more information on the computation of $c(\mathbf{x}_i, \mathbf{x}_j)$). Moreover, we compute the cost of giving a loan $\gamma = 0.42$ to the 70-th percentile of the $P(y = 1 | \mathbf{x})$ value for applicants who default/has a charged-off debt. Finally, we compare the utility achieved by our explanation methods with the same baselines we used on synthetic data.

Results. First, we just compare the utility achieved by each of the decision policies and counterfactual explanations for several values of the parameter α and number of counterfactual explanations k . Figures 4(a, b) summarizes the results, which show that, as α increases and the cost of changing to feature values with higher outcome values decreases, the competitive advantage by jointly optimizing the decision policy and the counterfactual explanations (Algorithm 2) grows significantly. Moreover,

¹¹<https://www.kaggle.com/wordsforthewise/lending-club/version/3>

similarly as in the experiments on synthetic data, just a small number of counterfactual explanations is sufficient to provide significant gains in terms of utility with respect to the optimal decision policy without counterfactual explanations.

Next, we challenge the assumption that individuals do not share the counterfactual explanations they receive with other individuals with different feature values. To this end, we assume that, given the set of counterfactual explanations \mathcal{A} found by Algorithm 2, individuals with initial feature value \mathbf{x} receive the counterfactual explanation $\mathcal{E}(\mathbf{x}) \in \mathcal{A}$ given by Eq. 4 and, with probability p_l , they also receive an additional explanation $\mathcal{E}'(\mathbf{x})$ picked at random from \mathcal{A} and they follow the counterfactual explanation who benefits them the most. Figure 4(c) summarizes the results for several values of p_l and number of counterfactual explanations, which show that, while providing counterfactual explanations is always better, if the leakage probability p_l is large, we are better off providing a small number of counterfactual explanations.

6 Conclusions

In this paper, we have studied the problem of finding the decision policies and counterfactual explanations that maximize utility in a setting in which individuals who are subject to the decisions taken by the policies use the counterfactual explanations they receive to invest effort strategically. Given a pre-defined policy, we have first shown that finding the optimal set of counterfactual explanations is a hard computational problem but presents some favorable properties, which allow the standard greedy algorithm to enjoy approximation guarantees. Additionally, we have also shown that, given a set of counterfactual explanations, the optimal decision policy is deterministic and can be found in polynomial time. Finally, building on these results, we have shown that, perhaps surprisingly, the problem of jointly finding both the optimal policy and set of counterfactual explanations reduces to maximizing a non-monotone submodular function, a problem that can be solved using a randomized greedy algorithm with approximation guarantees.

By uncovering a previously unexplored connection between strategic machine learning and interpretable machine learning, our work opens up many interesting directions for future work. For example, we have adopted a specific type of mechanism to provide counterfactual explanations (*i.e.*, one feature value per individual using a Stackelberg formulation). A natural next step would be to extend our analysis to other types of mechanisms fitting a variety of real-world applications. Moreover, we have assumed that the cost individuals pay to change features is given. However, our algorithms would be more effective if we develop a methodology to reliably estimate the cost function from real observational (or interventional) data. In our work, we have assumed that features take discrete values and individuals who are subject to the decisions do not share information between them. It would be interesting to lift these assumptions, extend our analysis to real-valued feature values, and develop decision policies and counterfactual explanations that are robust to information sharing between individuals (refer to Figure 4(c)). Finally, by assuming that $P(y|\mathbf{x})$ does not change after individuals best respond, we are implicitly assuming that the features \mathbf{x} are causal. However, in practice, this assumption is likely to be violated, as recently noted by Miller et al. (2019). It would be worth exploring the use of counterfactual explanations to distinguish between noncausal and causal features.

References

- Michael Brückner and Tobias Scheffer. Stackelberg games for adversarial prediction problems. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 547–555, 2011.
- Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1433–1452. SIAM, 2014.
- Gruia Calinescu, Chandra Chekuri, Martin Pal, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

- Supriyo Chakraborty, Richard Tomsett, Ramya Raghavendra, Daniel Harborne, Moustafa Alzantot, Federico Cerutti, Mani Srivastava, Alun Preece, Simon Julier, Raghuveer M Rao, et al. Interpretability of deep learning models: a survey of results. In *2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/IOP/SCI)*, pages 1–6. IEEE, 2017.
- S. Coate and G. Loury. Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 1993.
- Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 797–806, 2017.
- Nilesh Dalvi, Pedro Domingos, Sumit Sanghai, and Deepak Verma. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 99–108, 2004.
- Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 55–70, 2018.
- Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- R. Fryer and G. Loury. Valuing diversity. *Journal of Political Economy*, 2013.
- David Gunning and David W Aha. Darpa’s explainable artificial intelligence program. *AI Magazine*, 40(2):44–58, 2019.
- Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pages 111–122, 2016.
- Dorit S Hochbaum and Anu Pathria. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics (NRL)*, 45(6):615–627, 1998.
- L. Hu and Y. Chen. A short-term intervention for long-term fairness in the labor market. In *WWW*, 2018.
- Lily Hu, Nicole Immorlica, and Jennifer Wortman Vaughan. The disparate effects of strategic manipulation. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 259–268, 2019.
- Amir-Hossein Karimi, Gilles Barthe, Borja Belle, and Isabel Valera. Model-agnostic counterfactual explanations for consequential decisions. *arXiv preprint arXiv:1905.11190*, 2019.
- Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- Niki Kilbertus, Manuel Gomez-Rodriguez, Bernhard Schölkopf, Krikamol Muandet, and Isabel Valera. Fair decisions despite imperfect predictions. In *AISTATS*, 2019.
- Jon Kleinberg and Manish Raghavan. How do classifiers induce agents to invest effort strategically? In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 825–844, 2019.
- Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. Human decisions and machine predictions. *The quarterly journal of economics*, 133(1):237–293, 2018.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

- Zachary C Lipton. The mythos of model interpretability. *Queue*, 16(3):31–57, 2018.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*, 2017.
- John Miller, Smitha Milli, and Moritz Hardt. Strategic adaptation to classifiers: A causal perspective. *arXiv preprint arXiv:1910.10362*, 2019.
- Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 230–239, 2019.
- Shira Mitchell, Eric Potash, Solon Barocas, Alexander D’Amour, and Kristian Lum. Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions. *arXiv preprint arXiv:1811.07867*, 2018.
- Ramaravind K. Mothilal, Amit Sharma, and Chenhao Tan. Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020.
- W James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences*, 116(44):22071–22080, 2019.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical programming*, 14(1):265–294, 1978.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you? explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016.
- Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- Chris Russell. Efficient search for diverse coherent explanations. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 20–28, 2019.
- Roberto Solis-Oba. Approximation algorithms for the k-median problem. In *Efficient Approximation and Online Algorithms*, pages 292–320. Springer, 2006.
- Behzad Tabibian, Stratis Tsirtsis, Moein Khajehnejad, Adish Singla, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. Optimal decision making under strategic behavior. *Arxiv:1905.09239*, 2020.
- Gabriele Tolomei, Fabrizio Silvestri, Andrew Haines, and Mounia Lalmas. Interpretable predictions of tree-based ensembles via actionable feature tweaking. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 465–474, 2017.
- Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 10–19, 2019.
- Isabel Valera, Adish Singla, and Manuel Gomez Rodriguez. Enhancing the accuracy and fairness of human decision making. In *Advances in Neural Information Processing Systems*, pages 1769–1778, 2018.
- Paul Voigt and Axel Von dem Bussche. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed.*, Cham: Springer International Publishing, 2017.
- Sandra Wachter, Brent Mittelstadt, and Luciano Floridi. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*, 7(2):76–99, 2017a.
- Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017b.
- Adrian Weller. Challenges for transparency. 2017.

A Additional details on the experiments on real data

A.1 Raw features

The decision tree classifier (DTC) predicts whether an applicant fully pays a loan ($y = 1$) or defaults/charged off ($y = 0$) based on the following raw features:

- Loan Amount: The amount that the applicant initially requested.
- Employment Length: How long the applicant has been employed.
- State: The state where the applicant lives inside the United States of America.
- Debt to Income Ratio: The ratio between the applicant’s financial debts and her average income.
- Zip Code: The zip code of the applicant’s residential address.
- FICO Score: The applicant’s FICO score, which is a credit score based on consumer credit files. The FICO scores are in the range of 300-850 and the average of the high and low range for the FICO score of each applicant has been used for this study.

A.2 Computation of the cost between feature values

To compute the cost between feature values $c(\mathbf{x}_i, \mathbf{x}_j)$, we compare the raw features (refer to Appendix A.1) of the applications mapped to each leaf of the decision tree classifier (DTC). More specifically, for each leaf \mathbf{x}_i of the DTC, we consider the concatenation of the state code and partial zip code data (which we called state-zip), the debt to income ratio and the employment length of each applicant mapped into the leaf \mathbf{x}_i . Then, given a pair of leafs \mathbf{x}_i and \mathbf{x}_j , we compute the cost $c(\mathbf{x}_i, \mathbf{x}_j)$ as the sum of three terms:

- The ratio of unique values of state-zip for applicants mapped into leaf \mathbf{x}_j that no applicant mapped into leaf \mathbf{x}_i has. This term is weighted by 0.25.
- The difference between the average debt to income ratios of applicants mapped into leaf \mathbf{x}_i and leaf \mathbf{x}_j , where we set all negative differences to zero, we normalize the resulting differences by the 50-th percentile of debt to income values, and we cap the values to be equal or less than 1. This follows the intuition that it is more costly to move from a leaf with a higher value of average debt to income ratio to a leaf with a lower value. This term is weighted by 0.5.
- The difference between the average employment lengths of applicants mapped into leaf \mathbf{x}_j and leaf \mathbf{x}_i , where we set all negative differences to zero, we normalize the resulting differences to lie between 0 and 1. This follows the intuition that it is more costly to move from a leaf with a lower value of employment length to a leaf with a higher value. This term is weighted by 0.25.