

INDUSTRIAL AND SYSTEMS ENGINEERING



Counterfactual Explanations for Integer Optimization Problems

FELIX ENGELHARDT

RWTH Aachen University

JANNIS KURTZ AND Ş. İLKER BIRBIL

University of Amsterdam, Plantage Muidergracht 12, 1018 TV Amsterdam

TED K. RALPHS

Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA
18015

COR@L Technical Report 25T-026



Counterfactual Explanations for Integer Optimization Problems

FELIX ENGELHARDT¹, JANNIS KURTZ², Ş. İLKER BIRBİL², AND TED K. RALPHS³

¹RWTH Aachen University

²University of Amsterdam, Plantage Muidergracht 12, 1018 TV Amsterdam

³Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA 18015

October 13, 2025

Abstract

Counterfactual explanations (CEs) offer a human-understandable way to explain decisions by identifying specific changes to the input parameters of a *base* or *present* model that would lead to a desired change in the outcome. For optimization models, CEs have primarily been studied in limited contexts and little research has been done on CEs for general integer optimization problems. In this work, we address this gap. We first show that the general problem of constructing a CE is Σ_2^P -complete even for binary integer programs with just a single *mutable* constraint. Second, we propose solution algorithms for several of the most tractable special cases: (i) mutable objective parameters, (ii) a single mutable constraint, (iii) mutable right-hand-side, and (iv) all input parameters can be modified. We evaluate our approach using classical knapsack problem instances, focusing on cases with mutable constraint parameters. Our results show that our methods are capable of finding optimal CEs for small instances involving up to 40 items within a few hours. Additionally, we present experiments on the resource constrained shortest path problem.

Keywords: counterfactual explanations, integer optimization, complexity, constraint generation

1 Introduction

Automated decision-making increasingly shapes critical aspects of modern society, ranging from applications in medical scheduling and financial planning to supply chain optimization and disaster management. Many of these applications are modeled as integer optimization problems (IPs) and are solved by state-of-the-art methods, such as branch and cut. Despite their widespread use, the reasoning behind optimization-based decisions is often unclear to those affected. For instance, optimization algorithms are used for automated planning of surgeries in hospitals [CDB10, BBD⁺25]. A patient denied an earlier surgery date may ask for an explanation for this decision. In humanitarian logistics, optimal robust depot locations are calculated to quickly dispatch relief items to a region affected by a disaster [SWdF21]. Here, the properties of an optimal solution significantly depend on the choice of optimization model and uncertainty set. Officials from a certain region may ask for an explanation why their region did not receive a depot while another region did.

One established way to provide explanations for decisions made by automated systems is based on the paradigm of counterfactual explanations (CEs). A CE is a minimal change in model inputs that would lead to a desired change in outcome. For example, in the aforementioned depot location problem, a CE could be: “If the population in your region were 5% higher and the transportation time to region B were 3% smaller, it would be optimal to open a depot in your region.” Counterfactuals have become a cornerstone of explainable artificial intelligence that offers individuals understandable and actionable insights into automated predictions [VBH⁺24, MRIB22, MKR⁺24]. Their role in optimization problems—particularly those involving integrality constraints—has only begun to be explored. First works studied the concept of counterfactuals for binary linear optimization problems where only objective parameters are mutable [KB21, KSB21, KB23], drawing on techniques from constraint programming and inverse optimization. Later, [KIBd25] identified three types of counterfactuals of interest for optimization problems—weak, strong, and relative CEs—and analyzed them for linear optimization with continuous decision variables. In [LS24] the authors derive a heuristic for calculating weak CEs based on a penalty alternating direction method. Relative CEs have since been investigated for contextual stochastic optimization [RAF25], and CEs have been used to explain efficiency targets in data envelopment analysis [BRAM24] and to detect minimum number of constraints to make a problem infeasible in constraint programming [GOQ24]. Coming back to the case of integer optimization, however, the aforementioned three types of CEs remain largely unexplored, particularly when mutable parameters appear in the constraints.

Constructing CEs is closely related to a number of existing areas of research in optimization. The most obvious connection is to global sensitivity analysis, but rather than quantifying how outcomes vary under systematic parameter changes, CEs ask how parameters must be minimally changed to induce specified outcome changes. This perspective highlights the relationship to inverse optimization—indeed, one may view CE construction as a generalized form of inverse optimization [CMZ25, Wan13]. It also highlights connections to bilevel optimization: the model structures induced by weak CEs are similar to those in optimistic bilevel formulations [DKPVK15, Dem02, KLLS21], whereas strong CEs naturally correspond to pessimistic bilevel formulations [WTKR13]. Certain special cases connect directly to well-studied problem classes. The case in which only the objective function can be modified subsumes the most classical version of inverse optimization, in which one seeks an objective that makes a given solution optimal. The case in which only the right-hand-side can vary relates both to multiobjective optimization via the restricted value function and to the standard form of mixed integer bilevel linear optimization [Dem02]. Finally, the most general case considered here is equivalent to a bilevel optimization model in which the lower-level problem is a mixed integer bilinear optimization problem.

This paper fills a gap in the literature by studying weak and strong counterfactual explanations for integer linear optimization problems (ILPs) in which mutable parameters may appear both in the objective function and in the constraints. Our contributions are as follows:

- We present complexity results showing that constructing CEs in both the weak and strong cases is Σ_2^P -hard, even if the mutable parameters only appear in a single constraint.
- We propose solution algorithms to calculate weak and strong CEs for the case when only a single constraint may contain mutable parameters. We distinguish four different cases: the mutable parameters appear (i) only in the objective function, (ii) only in the constraints, (iii) only on the right-hand side of the constraints, or (iv) everywhere in the problem.
- We present computational experiments based on knapsack instances from `knplib`, showing that optimal CEs can be computed in a few hours for instances containing up to 40 items.

Additionally, we evaluate applications of CEs regarding resource constraint shortest path problems.

2 Definitions

In the following, we describe the basic definitions and the mathematical framework on which our study is based. To formally define the notion of a counterfactual explanation, we first need to define a so-called *present optimization problem* on which the current decision, which we refer to as the *present solution*, is based. In this paper, we consider the following form for the present problem:

$$\begin{aligned} \min \quad & \hat{c}^\top x \\ \text{s.t.} \quad & \hat{A}x \geq \hat{b}, \\ & x \in \mathcal{X}, \end{aligned} \tag{1}$$

where $\hat{c} \in \mathbb{Z}^n$ (the objective function), $\hat{b} \in \mathbb{Z}^n$ (the right-hand-side), and $\hat{A} \in \mathbb{Z}^{m \times n}$ (the constraint matrix) are the present values of the mutable parameters, i.e., those that can be modified to obtain a CE. The set $\mathcal{X} \subseteq \mathbb{Z}_+^n$ is the immutable feasible set, consisting of constraints that cannot be modified. In many applications, the set \mathcal{X} is the feasible set of a MILP, though for the algorithms described later, we only require the set \mathcal{X} be the feasible set of an optimization problem for which we have an appropriate black-box solver.

To reason about counterfactuals, we introduce two additional sets. The *avored solution space*, denoted by $\mathcal{D} \subseteq \mathbb{Z}_+^n$, consists of solutions that exhibit desirable properties and that we attempt to make optimal by modifying the mutable parameters to generate a CE. The *mutable parameter space*, denoted by $\mathcal{H} \subseteq \mathbb{Z}^n \times \mathbb{Z}^{m \times n} \times \mathbb{Z}^n$, are the triples (c, A, b) of values that the mutable parameters are allowed to take on in attempting to generate a CE. We assume that $\mathcal{D} \cap \mathcal{X} \neq \emptyset$ and that \mathcal{H} is bounded and contains $(\hat{c}, \hat{A}, \hat{b}) \in \mathcal{H}$.

The structure of \mathcal{H} depends on the nature of the application, but the following are the most typical ways in which the mutable parameter space would be specified:

1. Only objective function parameters are mutable: $\mathcal{H} = \mathcal{H}_c \times \{\hat{A}\} \times \{\hat{b}\}$
2. Only constraint parameters are mutable: $\mathcal{H} = \{\hat{c}\} \times \mathcal{H}_A \times \mathcal{H}_b$
3. Only the right-hand side parameters are mutable: $\mathcal{H} = \{\hat{c}\} \times \{\hat{H}_A\} \times \mathcal{H}_b$
4. All parameters are mutable: $\mathcal{H} = \mathcal{H}_c \times \mathcal{H}_A \times \mathcal{H}_b$

In principle, \mathcal{H} can be any (MILP-representable) set. In the remainder of the paper, we choose \mathcal{H} to be a discrete set for two reasons. First, in most applications, parameter values are naturally discrete. Even if we allow the parameters to take on rational values, we can still scale them to be integral without loss of generality. Second, allowing parameter values to be continuous can result in feasible sets that are open, which leads to difficulties in constructing CEs; see [KIBd25] for a discussion on open feasible regions.

We now define the problem of constructing weak and strong CEs. Given a present problem of the form (1), a favored solution space \mathcal{D} , and a mutable parameter space \mathcal{H} , a *weak CE* is any $(c, A, b) \in \mathcal{H}$ such that *some* optimal solution of the resulting modified instance lies in \mathcal{D} , i.e., fulfills the desired properties.

Definition 1. A *weak counterfactual explanation* for (1) is a point $(c, A, b) \in \mathcal{H}$ that satisfies the *weak CE condition*:

$$\mathcal{D} \cap \arg \min_{x \in \mathcal{X}: Ax \geq b} c^\top x \neq \emptyset. \tag{2}$$

In general, we are interested in CEs that modify the present problem as little as possible. To measure the amount of modification, we utilize a distance measure $\delta : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}_{\geq 0}$ and define the *cost* of a weak CE to be $\delta \left((c, A, b), (\hat{c}, \hat{A}, \hat{b}) \right)$. We call a weak CE with minimal cost an *optimal weak CE*. An optimal weak CE can be computed by solving the following optimistic bilevel problem.

$$\begin{aligned} \min_{c, A, b, x} \quad & \delta \left((c, A, b), (\hat{c}, \hat{A}, \hat{b}) \right) \\ \text{s.t.} \quad & x \in \arg \min_{x \in \mathcal{X}: Ax \geq b} c^\top x, \\ & x \in \mathcal{D}, \\ & (c, A, b) \in \mathcal{H}. \end{aligned} \tag{3}$$

When a problem has multiple optimal solutions, either due to its inherent structure or the characteristics of the solver used, a decision maker may want to ensure that *all* optimal solutions satisfy the desired properties specified by \mathcal{D} . This motivates the concept of a *strong CE*.

Definition 2. A *strong counterfactual explanation* for (1) is a point $(c, A, b) \in \mathcal{H}$ that satisfies the *strong CE condition*:

$$\arg \min_{x \in \mathcal{X}: Ax \geq b} c^\top x \subseteq \mathcal{D}. \tag{4}$$

Once again, for a given distance measure δ , the cost of a strong CE is then given by $\delta \left((c, A, b), (\hat{c}, \hat{A}, \hat{b}) \right)$. We call a strong CE with minimal cost δ an *optimal strong CE*. An optimal strong CE can be computed by solving the following pessimistic bilevel problem:

$$\begin{aligned} \min_{c, A, b, x} \quad & \delta \left((c, A, b), (\hat{c}, \hat{A}, \hat{b}) \right) \\ \text{s.t.} \quad & x \in \mathcal{D} \quad \forall x \in \arg \min_{x \in \mathcal{X}: Ax \geq b} c^\top x, \\ & (c, A, b) \in \mathcal{H}. \end{aligned} \tag{5}$$

To illustrate the definitions of weak and strong CEs, let us give a toy example.

Example 1. Consider the following instance of (1):

$$\begin{aligned} \min \quad & x_1 + 2x_2 + 2x_3 \\ \text{s.t.} \quad & x_1 + 3x_2 + 2x_3 \geq 3, \\ & x \in \{0, 1\}^3. \end{aligned}$$

The optimal solution is to select only $x_2 = 1$ and to set all other variables to zero. Assume that only the constraint parameters $\hat{a}_2 = 3$ and $\hat{a}_3 = 2$ are mutable, each in the range $\{0, 1, 2, 3, 4\}$. Here is the counterfactual question: “What is the minimal change in the mutable parameters we need to perform, such that a solution with $x_3 = 1$ is optimal, i.e., $\mathcal{D} = \{x \in \{0, 1\}^3 : x_3 = 1\}$?” Figure 1 shows the feasible regions for the weak and the strong CEs for this example.




Figure 1: Counterfactual regions for a (1) with right-hand side $b = 3$ and three item i_1, i_2, i_3 , the latter two of which have a mutable weight. The favored solution space consists of all all solutions that contain i_3 , i.e., $\mathcal{D}_1 = \{x \in \{0, 1\}^3 : x_3 = 1\}$. Green and orange regions represent strong and weak CE, respectively. Each region is labeled with the items that are part of an optimal solution and the optimum solution value. The black O marks the parameter values (\hat{A}) of the present problem.

In the remainder of the paper, we present both results establishing the computational complexity of the problem of constructing weak and strong CE, as well as algorithms for doing so. For reasons discussed later in the paper (see especially Remark 3), the algorithms presented assume $m = 1$ (we are only allowed to modify a single constraint). The complexity results, however, are general.

3 Complexity

In this section, we show that the problems of deciding whether a counterfactual explanation exists is complete for Σ_2^P in both the weak and strong cases. We use the oracle-based definition of the polynomial hierarchy, as formalized in [AB07], and the polynomial many-to-one notion of reduction introduced by [Kar72]. This means that to prove membership in Σ_2^P , we have to show that whenever a weak/strong CE exists, there is a formal certificate of this fact that can be verified in polynomial time, given an oracle for solving problems in NP. The assumption that we have access to an oracle for solving problems in NP means, in particular, that we assume that we have the ability to solve the decision version of an ILP in constant time, and hence, the optimization version of ILP in polynomial time.

3.1 Weak Counterfactual Explanations

We first consider the problem of determining whether there exists a weak CE with respect to a present optimization problem of the form (1), which we denote as WCE-FEAS. Formal complexity arguments require careful specification of the input and its (encoded) size. For this problem, the formal input is the triple $(\hat{c}, \hat{A}, \hat{b})$ and the sets \mathcal{X} , \mathcal{H} , and \mathcal{D} . To reason about complexity and input sizes, we must make assumptions about the sets \mathcal{X} , \mathcal{H} , and \mathcal{D} , and we assume here that these sets are the feasible regions of ILPs (the results should hold for any sets for which the problem of determining whether the set is non-empty is in NP). The size of the input is the sum of the sizes of the descriptions of the individual components.

When a weak CE exists, Equation (2) suggests both a certificate and a method of verifying it that requires solving an ILP. It remains to show that there must always exist such a certificate that is polynomial in the size of the input and that this certificate can be verified in polynomial time.

Theorem 1. WCE-FEAS is in Σ_2^p .

Proof. If a weak CE exists, this means that there must exist a point $x^* \in \mathbb{Z}^n$ and a triple $(c^*, A^*, b^*) \in \mathbb{Z}^n \times \mathbb{Z}^{m \times n} \times \mathbb{Z}^n$ such that $x^* \in \arg \min_{x \in \mathcal{X}: A^*x \geq b^*} (c^*)^\top x$ and that $x^* \in \mathcal{D}$. For $\tilde{c} := (c^*)^\top x^*$ we must then have that $(x^*, c^*, A^*, b^*) \in \mathcal{Q}$ where

$$\mathcal{Q} := \{(x, c, A, b) \in \mathcal{X} \cap \mathcal{D} \times \mathcal{H} : Ax \geq b, c^\top x = \tilde{c}\}$$

Since mixed-integer quadratic programming feasibility is in NP (see Theorem 1 in [PDM17]), there must exist $(\bar{x}, \bar{c}, \bar{A}, \bar{b}) \in \mathcal{Q}$ that has size polynomial in the original input data. We claim that $(\bar{x}, \bar{c}, \bar{A}, \bar{b}) \in \mathcal{Q}$ is then a certificate that can be checked in polynomial time, given an oracle for problems in NP.

The verification requires checking first whether $\bar{x} \in \mathcal{X} \cap \mathcal{D}$, which can be done in polynomial time under our assumptions. Second, we need to check that $\bar{x} \in \arg \min_{x \in \mathcal{X}: \bar{A}x \geq \bar{b}} \bar{c}^\top x$. This can also be done in polynomial time using the available NP oracle. \square

The following result then shows that WCE-FEAS is Σ_2^p -complete.

Theorem 2. WCE-FEAS is complete for Σ_2^p and cannot be approximated with a constant approximation factor in polynomial time, unless the polynomial hierarchy collapses, even when $\mathcal{H} = \{\hat{c}\} \times \mathcal{H}_a \times \{\hat{b}\}$ (for given $\hat{c} \in \mathbb{Z}^n$ and $\hat{b} \in \mathbb{Z}^m$) and $\mathcal{X} = \{0, 1\}^n$.

Proof. We reduce the unitary bilevel interdiction knapsack problem (UBIK) to WCE-FEAS. UBIK is defined as follows: Let $(\bar{c}, \bar{a}, \bar{b})$ be an instance of the binary knapsack problem with n items

$$\begin{aligned} & \max \bar{c}^\top x \\ & \text{s.t.} \quad \bar{a}^\top x \leq \bar{b}, \\ & \quad \quad x \in \{0, 1\}^n, \end{aligned}$$

where all entries of $(\bar{c}, \bar{a}, \bar{b})$ are positive and integer. Consider a follower, who desires a solution with objective value at least K , and a leader who can delete up to k items from the knapsack. Then, the UBIK consists of deciding if there exists a subset of items $I \subseteq [n]$ with $|I| \leq k$ such that for any feasible knapsack solution using only the items from $[n] \setminus I$, the objective value is strictly less than a given target K . It was shown in [BTCCH] that UBIK is Σ_2^p -complete. The UBIK problem is equivalent to an instance of WCE-FEAS where the present problem is

$$\begin{aligned} & \min -\bar{c}^\top x \\ & \text{s.t.} \quad -\bar{a}^\top x \geq -\bar{b}, \\ & \quad \quad x \in \{0, 1\}^n, \end{aligned}$$

which has been put into the required form by negating the input data, and the required sets are defined as

- $\mathcal{X} = \{0, 1\}^n$,
- $\mathcal{D} = \{x \in \{0, 1\}^n : -\bar{c}^\top x \geq -K + 1\}$, and

- $\mathcal{H} = \{\bar{c}\} \times \mathcal{H}_a \times \{\bar{b}\}$, where $\mathcal{H}_a = \{a \in \mathbb{R}^n : a_i = -\bar{a}_i - \delta_i \bar{b}, i \in [n], \sum_{i=1}^n \delta_i \leq k, \delta \in \{0, 1\}^n\}$.

The definition of \mathcal{D} encodes the condition that the optimal value is strictly less than K , while the definition of \mathcal{H} encodes that up to k items may be forced to have value zero in the solution (interdicted).

Now let $a \in \mathcal{H}_a$ be a weak CE. By construction, the maximum attainable objective of the knapsack problem with weights a is $K - 1$. Furthermore, at most k items were interdicted (since $a \in \mathcal{H}_a$) to achieve that, i.e., a YES instance for UBIK. Similarly, if the weak CE problem is infeasible this implies that for every set of at most size k of interdicted items there exists a KP solution that has objective value at least K . Hence, we showed that UBIK is a yes-instance iff the constructed weak CE problem is feasible.

Since UBIK is Σ_2^P -complete, the reduction shows that WCE-FEAS is Σ_2^P -hard. Together with Theorem 1, this implies that WCE-FEAS is Σ_2^P -complete, even under strong restrictions. Notably, the CE only requires solving a feasibility problem, which implies that no efficient approximation algorithms exist unless the polynomial hierarchy collapses. \square

Remark 1. For the (3) where only the right hand side b is mutable, it suffices to enumerate all relevant right hand side values of b . Assuming that $\mathcal{H}_b = \{b \in \mathbb{Z} : \underline{b} \leq b \leq \bar{b}\}$, we then have to solve the problem

$$\begin{aligned} \min_{b \in \mathbb{Z}, \underline{b} \leq b \leq \bar{b}} \delta \left((\hat{c}, \hat{A}, b), (\hat{c}, \hat{A}, \hat{b}) \right) \\ \text{s.t. } x \in \arg \min_{x \in \mathcal{X}: \hat{A}x \geq b} \hat{c}^\top x, \\ x \in \mathcal{D}. \end{aligned}$$

We can now iterate through all values of b , sorted by distance regarding $\delta \left((\hat{c}, \hat{A}, b), (\hat{c}, \hat{A}, \hat{b}) \right)$, and for each, we can check if the Condition (2) for (\hat{c}, \hat{A}, b) is true. The latter condition can be verified by checking

$$\min_{x \in \mathcal{X}: \hat{A}x \geq b} \hat{c}^\top x = \min_{x \in \mathcal{X} \cap \mathcal{D}: \hat{A}x \geq b} \hat{c}^\top x,$$

which involves solving two (potentially) NP-hard problems.

This remark leads to a solution algorithm when the present problem is a knapsack problem. If at the same time the favored solution space is given by a fixed number of knapsack constraints, then both problems we have to solve for every value of b can be solved in pseudopolynomial time, which leads to a pseudopolynomial algorithm. The exact computational complexity discussion follows.

Theorem 3. Let the underlying problem be a knapsack problem, i.e., $\mathcal{X} = \{0, 1\}^n$, $\hat{A} = \hat{a} \in \mathbb{Z}_+^n$ and

$$\mathcal{D} = \left\{ x \in \{0, 1\}^n : q_i^\top x \geq p_i, i = 1, \dots, T \right\}$$

with $q_i \in \mathbb{Z}_+^n$, $p_i \in \mathbb{Z}$ for all $i \in [T]$. Furthermore, assume $\mathcal{H}_b = \{b \in \mathbb{Z} : \underline{b} \leq b \leq \bar{b}\}$, then (3) can be solved in time $\mathcal{O} \left(n\bar{b}^2 \prod_{i=1}^T (p_i + 1) \right)$.

Proof. First, note that if $\underline{b} < 0$, then we can remove all negative values for b . Hence, in the worst case we have to check \bar{b} times whether b is a weak CE. For each value of b , we have to solve two multi-dimensional knapsack problems:

$$\min_{x \in \{0, 1\}^n: \hat{a}^\top x \geq b} \hat{c}^\top x \quad \text{and} \quad \min_{\substack{x \in \{0, 1\}^n: \hat{a}^\top x \geq b, \\ q_i^\top x \geq p_i \quad i=1, \dots, T}} \hat{c}^\top x.$$

The second problem can be solved by the classical pseudopolynomial algorithm for multi-dimensional knapsack problems in time $\mathcal{O}\left(nb \prod_{i=1}^T (p_i + 1)\right)$ (see [Fre04]) and this dominates the runtime of solving the classical knapsack problem on the left. Since $b \leq \bar{b}$ and we have to solve the latter problems \bar{b} times, the desired result follows. \square

3.2 Strong Counterfactual Explanations

In this section, we show complexity results and solution algorithms for the problem of determining whether there exists a strong CE, which we refer to as SCE-FEAS. The basic setup for SCE-FEAS is the same as for the case of WCE-FEAS, and the problem has the same form of input. The difference is in the verification problem, which is more involved in the case of a strong CE. Rather than needing to show that $\exists x \in \mathcal{D}, x \in \arg \min_{x \in \mathcal{X}: Ax \geq b} c^\top x$, we need to show that $\forall x \in \mathcal{D}, x \in \arg \min_{x \in \mathcal{X}: Ax \geq b} c^\top x$.

We first provide the following proposition showing that for a given point $(c, a, b) \in \mathcal{H}$, we can check if the point is a strong CE by solving three NP-hard optimization problems.

Proposition 1. *A parameter vector $(c, A, b) \in \mathcal{H}$ is a strong counterfactual explanation, if and only if*

$$\min_{x \in \mathcal{X}: Ax \geq b} c^\top x = \min_{x \in \mathcal{X} \cap \mathcal{D}: Ax \geq b} c^\top x \leq \min_{\substack{x \in \mathcal{X} \setminus \mathcal{D}: \\ Ax \geq b}} c^\top x - 1. \quad (6)$$

Proof. The first equation ensures that *some* solution from \mathcal{D} is indeed optimal for the parameters (c, a, b) , just as in the weak CE case. The second inequality ensures that *all* optimal solutions are contained in \mathcal{D} . The result follows directly from the definition of strong CE. \square

Note that the condition $x \notin \mathcal{D}$ used in the last of the three problems is not always possible to formulate compactly. If \mathcal{D} is defined by a single constraint, i.e., $\mathcal{D}_{\geq} = \{x \in \mathcal{X} : \sum_{i \in \mathcal{I}} \alpha_i x_i \leq \beta\}$, then the condition $x \notin \mathcal{D}$ is equivalent to $x \in \{x \in \mathcal{X} : \sum_{i \in \mathcal{I}} \alpha_i x_i > \beta\}$. However, if the description of \mathcal{D} is given by multiple constraints, a big-M formulation may need to be used to model $x \notin \mathcal{D}$. In the rest of the section, we assume that we have a polynomial size description for $x \notin \mathcal{D}$. Under this condition, we show that SCE-FEAS is in Σ_2^p by an argument similar to that for the weak CE, but with a different certificate. In the case that a strong CE exists, there are points $x_1^*, x_2^* \in \mathbb{Z}^n$ and a triple $(c^*, A^*, b^*) \in \mathbb{Z}^n \times \mathbb{Z}^{m \times n} \times \mathbb{Z}^n$ such that $x_1^* \in \arg \min_{x \in \mathcal{X}: A^* x \geq b^*} (c^*)^\top x$, $x_1^* \in \mathcal{D}$, and $x_2^* \in \arg \min_{x \in \mathcal{X}, x \notin \mathcal{D}: A^* x \geq b^*} (c^*)^\top x$, with $(c^*)^\top x_1^* \leq (c^*)^\top x_2^* - 1$. As in WCE-FEAS, we must show that a certificate of polynomial size, consisting of these elements, must exist and can be verified in polynomial time, given an oracle for problems in NP.

Theorem 4. *SCE-FEAS is in Σ_2^p .*

Proof. If a strong CE exists, this means that there must exist $x_1^*, x_2^* \in \mathbb{Z}^n$ and a triple $(c^*, A^*, b^*) \in \mathbb{Z}^n \times \mathbb{Z}^{m \times n} \times \mathbb{Z}^n$ such that $x_1^* \in \arg \min_{x \in \mathcal{X}: A^* x \geq b^*} (c^*)^\top x$, $x_1^* \in \mathcal{D}$, and $x_2^* \in \arg \min_{x \in \mathcal{X}, x \notin \mathcal{D}: A^* x \geq b^*} (c^*)^\top x$, with $(c^*)^\top x_1^* \leq (c^*)^\top x_2^* - 1$. For $\tilde{c}_1 := (c^*)^\top x_1^*$ we must then have

$$\begin{aligned} & (x_1^*, x_2^*, c^*, A^*, b^*) \in \mathcal{R} := \\ & \{(x_1, x_2, c, A, b) \in (\mathcal{X} \cap \mathcal{D}) \times \mathcal{X} \times \mathcal{H} : \\ & Ax_1 \geq b, Ax_2 \geq b, c^\top x_1 = \tilde{c}_1, \\ & c^\top x_1 \leq c^\top x_2 - 1\} \end{aligned}$$

Since MIQP feasibility is in NP, there must exist $(\bar{x}_1, \bar{x}_2, \bar{c}, \bar{A}, \bar{b}) \in \mathcal{R}$ that has size polynomial in the original input data (see Theorem 1 in [PDM17]). We claim that $(\bar{x}_1, \bar{x}_2, \bar{c}, \bar{A}, \bar{b}) \in \mathcal{R}$ is then a certificate that can be checked in polynomial time, given an oracle for problems in NP.

The verification requires checking first whether $\bar{x}_1 \in \mathcal{X} \cap \mathcal{D}$ and $\bar{x}_2 \in \mathcal{X}$, which can be done in polynomial time under our assumptions. Second, we need to check both that $\bar{x}_1 \in \arg \min_{x \in \mathcal{X}: \bar{A}x \geq \bar{b}} \bar{c}^\top x$ and whether $x_2^* \in \arg \min_{x \in \mathcal{X}, x \notin \mathcal{D}: \bar{A}x \geq \bar{b}} \bar{c}^\top x$. Both of these checks can also be done in polynomial time using the available NP oracle. \square

Theorem 5. SCE-FEAS is complete for Σ_2^P , even for $\mathcal{H} = \{\hat{c}\} \times \mathcal{H}_a \times \{\hat{b}\}$, where \mathcal{H}_a is a finite set, and $\mathcal{X} = \{0, 1\}^n$.

Proof. We use exactly the same reduction as in the proof of Theorem 2 with the same inputs, but we simply interpret those inputs as a description of an instance of SCE-FEAS instead of WCE-FEAS. \square

4 Algorithms

Building on our earlier observations regarding the structure of CE problems, we now present algorithms designed to compute both weak and strong CEs.

4.1 Weak Counterfactual Explanations

In the following, we develop three algorithms corresponding to the three cases of the mutable parameter space listed in Section 2.

Only objective function parameters are mutable. If \hat{a} and \hat{b} are fixed, the mutable parameter space is denoted by $\mathcal{H}_c \subseteq \mathbb{R}^n$. In this case, we can reformulate (3) as

$$\min_{x, c} \delta \left((c, \hat{a}, \hat{b}), (\hat{c}, \hat{a}, \hat{b}) \right) \tag{7a}$$

$$s.t. \quad c^\top x \leq c^\top y \quad \forall y \in \mathcal{X} : \hat{a}^\top y \geq \hat{b}, \tag{7b}$$

$$\hat{a}^\top x \geq \hat{b}, \tag{7c}$$

$$x \in \mathcal{D} \cap \mathcal{X}, \tag{7d}$$

$$c \in \mathcal{H}_c, \tag{7e}$$

where Constraints (7c) and (7d) ensure that an optimal solution x is feasible and lies in \mathcal{D} , while Constraints (7b) ensure that the solution x is optimal for the objective vector c . Note that the problem has one constraint for every feasible solution $y \in \mathcal{X} \cap \{y : \hat{a}^\top y \geq \hat{b}\}$, i.e., the formulation can be of exponential size. Hence, we iteratively generate the constraints by finding the most violating constraint via a separation problem in each iteration. This procedure is the same as in [KB23] applied to our problem setup. The whole procedure can be found in Algorithm 1. The correctness follows from the same argumentation as in [KB23].

Only constraint parameters are mutable. In this subsection, only the constraint parameters a, b are mutable, i.e., \hat{c} is fixed. The mutable parameter space is denoted as $\mathcal{H}_a \times \mathcal{H}_b$, or in short, $\mathcal{H}_{a,b} \subseteq \mathbb{Z}^{n+1}$.

Algorithm 1 Solve Problem (7)

Require: $\hat{a}, \hat{b}, \hat{c}, \mathcal{D}, \mathcal{X}, \mathcal{H}$

Ensure: Optimal solution c^* of (7).

set $\mathcal{Y} \leftarrow \emptyset$

optimal \leftarrow false

while not optimal **do**

 Calculate an optimal solution (\tilde{c}, \tilde{x}) of problem

$$\begin{aligned} & \min_{c,x} \delta \left((c, \hat{a}, \hat{b}), (\hat{c}, \hat{a}, \hat{b}) \right) \\ \text{s.t.} \quad & c^\top x \leq c^\top y \quad \forall y \in \mathcal{Y}, \\ & \hat{a}^\top x \geq \hat{b}, \\ & x \in \mathcal{D} \cap \mathcal{X}, \\ & c \in \mathcal{H}_c. \end{aligned}$$

 Calculate an optimal solution \tilde{y} of the separation problem

$$\begin{aligned} & \min \tilde{c}^\top y \\ \text{s.t.} \quad & \hat{a}^\top y \geq \hat{b}, \\ & y \in \mathcal{X}. \end{aligned}$$

if $\tilde{c}^\top \tilde{y} < \tilde{c}^\top \tilde{x}$ **then**

$\mathcal{Y} \leftarrow \mathcal{Y} \cup \{\tilde{y}\}$

else

 optimal \leftarrow true

\triangleright End While Loop

end if

end while

return \tilde{c}

We first define the smallest and the largest value an optimal solution $x \in \mathcal{D} \cap \mathcal{X}$ can have over all parameter configurations $(a, b) \in \mathcal{H}_{a,b}$. The smallest value is

$$c_{\min} := \min_{(a,b) \in \mathcal{H}_{\text{feas}}} \min_{\substack{a^\top x \geq b \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x \quad (8)$$

and the largest value is

$$c_{\max} := \max_{(a,b) \in \mathcal{H}_{\text{feas}}} \min_{\substack{a^\top x \geq b \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x, \quad (9)$$

where $\mathcal{H}_{\text{feas}} := \{(a, b) \in \mathcal{H}_{a,b} : \{x \in \mathcal{D} \cap \mathcal{X} : a^\top x \geq b\} \neq \emptyset\}$. The following lemma provides a simple way to calculate bounds for c_{\min} and c_{\max} .

Lemma 1. *It holds*

$$c_{\min} \geq \min_{\substack{a_{\max}^\top x \geq b_{\min} \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x =: \underline{c}$$

and

$$c_{\max} \leq \max_{\substack{a_{\max}^\top x \geq b_{\min} \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x =: \bar{c},$$

where $(a_{\max})_i := \max_{(a,b) \in \mathcal{H}_{a,b}} a_i$ and $b_{\min} := \min_{(a,b) \in \mathcal{H}_{a,b}} b$.

Proof. To prove the first result consider an arbitrary parameter configuration $(a, b) \in \mathcal{H}_{\text{feas}}$ and any feasible solution x^* , i.e., it holds $a^\top x^* \geq b$. Since $(\hat{a}, \hat{b}) \in \mathcal{H}_{a,b}$ such a point (a, b) always exists. Then this solution is feasible for $a_{\max}^\top x \geq b_{\min}$ as well, since

$$a_{\max}^\top x^* \geq a^\top x^* \geq b \geq b_{\min},$$

where the first inequality follows from $x^* \geq 0$ and the definition of a_{\max} . Hence, it holds for every $(a, b) \in \mathcal{H}_{\text{feas}}$ that

$$\min_{\substack{a^\top x \geq b \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x \geq \min_{\substack{a_{\max}^\top x \geq b_{\min} \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x,$$

which proves the first result. The second result follows from a similar argumentation. Applying the results above for all $(a, b) \in \mathcal{H}_{\text{feas}}$, we obtain

$$\min_{\substack{a^\top x \geq b \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x \leq \max_{\substack{a^\top x \geq b \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x \leq \max_{\substack{a_{\max}^\top x \geq b_{\min} \\ x \in \mathcal{D} \cap \mathcal{X}}} \hat{c}^\top x,$$

which proves the result. \square

Since \hat{c} and all solutions $x \in \mathcal{X}$ have only integer entries, the objective value of the best solution in $\mathcal{D} \cap \mathcal{X}$ for any $(a, b) \in \mathcal{H}_{\text{feas}}$ is contained in

$$\mathcal{C} := \{c_{\min}, c_{\min} + 1, \dots, c_{\max}\}. \quad (10)$$

The main idea of the algorithm developed in this section is to iterate over the values in \mathcal{C} and find the best parameter setting $(a, b) \in \mathcal{H}_{a,b}$ that ensures the existence of a optimal solution in \mathcal{D} with the corresponding optimal value from \mathcal{C} . We first show the following lemma.

Lemma 2. *For every $v \in \mathcal{C}$, the following problem is either infeasible or any optimal solution $(a^*(v), b^*(v))$ is a (not necessarily optimal) weak counterfactual explanation:*

$$\min_{a,b,x} \delta \left((\hat{c}, a, b), (\hat{c}, \hat{a}, \hat{b}) \right) \quad (11a)$$

$$s.t. \quad a^\top y \leq b - 1 \quad \forall y \in \mathcal{X} : \hat{c}^\top y \leq v - 1, \quad (11b)$$

$$\hat{c}^\top x = v, \quad (11c)$$

$$a^\top x \geq b, \quad (11d)$$

$$x \in \mathcal{D} \cap \mathcal{X}, \quad (11e)$$

$$(a, b) \in \mathcal{H}_{a,b}. \quad (11f)$$

Proof. Assume the problem is feasible. Since $\mathcal{H}_{a,b}$ is discrete and bounded the problem must have an optimal solution $(a^*(v), b^*(v), x^*)$. We have to show that this optimal solution is a counterfactual explanation, i.e., there exists a solution $\tilde{x} \in \mathcal{D}$ with

$$\tilde{x} \in \arg \min_{x \in \mathcal{X} : a^*(v)^\top x \geq b^*(v)} \hat{c}^\top x.$$

We show in the following that x^* fulfills the latter condition. First, Constraint (11e) ensures that $x^* \in \mathcal{D} \cap \mathcal{X}$. Constraint (11d) ensures that

$$a^*(v)^\top x^* \geq b^*(v)$$

holds. Second, Constraints (11b) and (11c) ensure that every solution in \mathcal{X} which has a better objective value than $\hat{c}^\top x^*$ is infeasible. Hence, x^* must be optimal for the problem, and $(a^*(v), b^*(v))$ is a weak counterfactual explanation. \square

The intuition behind Lemma 2 is as follows: The first set of constraints ensures that all solutions $y \in \mathcal{X}$ which are better than the considered optimal value v are infeasible for the solution $(a^*(v), b^*(v))$, i.e., the optimal value of the problem for the parameters $(\hat{c}, a^*(v), b^*(v))$ is at least v . The remaining constraints ensure that a solution $x \in \mathcal{D}$ exists which is feasible for the constraint parameters $(a^*(v), b^*(v))$ and has optimal value v . If such a solution exists, this solution must be optimal, and hence $(a^*(v), b^*(v))$ is a weak CE and a possible candidate for an optimal weak CE.

To solve Problem (11) effectively, a cut generation procedure can be performed, where the constraints for a solution y are separated iteratively. This procedure is shown in Algorithm 2. Note that if in Algorithm 2 an upper bound d^* for the cost of a CE is known, we use that bound to cut off non-improving solutions. From classical constraint generation theory the following theorem directly follows.

Theorem 6. *Algorithm 2 calculates an optimal solution of Problem (11).*

We now show that there must exist a value $v^* \in \mathcal{C}$ such that $(a^*(v^*), b^*(v^*))$ is an optimal weak counterfactual explanation.

Lemma 3. *There exists a $v^* \in \mathcal{C}$, such that any optimal solution $(a^*(v^*), b^*(v^*))$ of Problem (11) with $v = v^*$ is an optimal weak CE.*

Algorithm 2 Solve Problem (11)

Require: $v, \hat{a}, \hat{b}, \hat{c}, \mathcal{D}, \mathcal{X}, \mathcal{H}_{a,b}$, known upper bound d^* for (3)

Ensure: Optimal solution $(a^*(v), b^*(v))$ of (11).

set $\mathcal{Y} \leftarrow \emptyset$

optimal \leftarrow false

while not optimal **do**

 Calculate an optimal solution (\tilde{a}, \tilde{b}) of problem

$$\begin{aligned} & \min_{a,b,x} \delta\left((\hat{c}, a, b), (\hat{c}, \hat{a}, \hat{b})\right) \\ \text{s.t. } & a^\top y \leq b - 1 \quad \forall y \in \mathcal{Y}, \\ & \hat{c}^\top x = v, \\ & a^\top x \geq b, \\ & x \in \mathcal{D} \cap \mathcal{X}, \\ & \delta\left((\hat{c}, a, b), (\hat{c}, \hat{a}, \hat{b})\right) \leq d^*, \\ & (a, b) \in \mathcal{H}_{a,b}. \end{aligned}$$

If problem is infeasible, stop and return *infeasible*.

Otherwise, calculate an optimal solution \tilde{y} of the separation problem

$$\begin{aligned} & \max \tilde{a}^\top y \\ \text{s.t. } & \hat{c}^\top y \leq v - 1, \\ & y \in \mathcal{X}. \end{aligned}$$

if $\tilde{a}^\top \tilde{y} \geq \tilde{b}$ **then**

$\mathcal{Y} \leftarrow \mathcal{Y} \cup \{\tilde{y}\}$

else

 optimal \leftarrow true

end if

end while

return (\tilde{a}, \tilde{b})

▷ End While Loop

Proof. Let $(a^*, b^*) \in \mathcal{H}_{a,b}$ be an optimal weak CE, i.e., an optimal solution of Problem (3). Then there exists a solution $x^* \in \mathcal{D}$ which is optimal for the problem

$$\begin{aligned} v^* &:= \min \hat{c}^\top x \\ \text{s.t.} \quad &(a^*)^\top x \geq b^*, \\ &x \in \mathcal{X}. \end{aligned}$$

Since $(a^*, b^*) \in \mathcal{H}_{\text{feas}}$, $c_{\min} \leq v^* \leq c_{\max}$ must hold. And since $\hat{c} \in \mathbb{Z}^n$ and $\mathcal{X} \subset \mathbb{Z}_+^n$, v^* is integer and it must hold that $v^* \in \mathcal{C}$. Hence, (a^*, b^*, x^*) is feasible for Constraints (11c) – (11f) for $v = v^*$. Furthermore, from the optimality of x^* and the integrality of y and \hat{c}, a^*, b^* , it follows that

$$(a^*)^\top y \leq b^* - 1 \quad \forall y \in \mathcal{X} : \hat{c}^\top y \leq v^* - 1.$$

We can conclude that (a^*, b^*, x^*) is feasible for Problem (11) with $v = v^*$. Since every optimal solution of (11) is a weak CE by Lemma 2, it follows that $(a^*(v^*), b^*(v^*))$ must be an optimal weak CE. \square

The latter two lemmas indicate that we can solve Problem (3) by iterating over all values $v \in \mathcal{C}$ and solve the corresponding Problem (11). The best solution $(a^*(v), b^*(v))$ found over all these values must then be an optimal solution of (3). While this idea is feasible, it may happen that \mathcal{C} contains a large number of values, and solving (3) for each is too expensive. To cut off unpromising solutions, we derive a lower bound for the objective function values in the following lemma.

Lemma 4. *For a given $\bar{v} \in \mathbb{Z}$ the optimal value of the following problem is a lower bound for every feasible master problem (11) with integer $v \geq \bar{v}$:*

$$\begin{aligned} \min_{a,b} \quad &\delta \left((\hat{c}, a, b), (\hat{c}, \hat{a}, \hat{b}) \right) \\ \text{s.t.} \quad &a^\top y \leq b - 1 \quad \forall y \in \mathcal{X} : \hat{c}^\top y \leq \bar{v} - 1, \\ &(a, b) \in \mathcal{H}_{a,b}. \end{aligned} \tag{12}$$

Proof. Consider an arbitrary $v \geq \bar{v}$ for which (11) is feasible and a corresponding arbitrary feasible solution $(a(v), b(v), x(v))$ of (11). We show that $(a(v), b(v))$ is feasible for (12) to prove the result. Clearly, $(a(v), b(v)) \in \mathcal{H}_{a,b}$ is true. Furthermore, from Constraints (11b) it follows that

$$a(v)^\top y \leq b(v) - 1 \quad \forall y \in \mathcal{X} : \hat{c}^\top y \leq v - 1.$$

Since $\bar{v} \leq v$, $(a(v), b(v))$ must also be feasible for (12). This proves the result. \square

Problem (12) can be solved again by a constraint generation procedure, where we iteratively generate constraints for corresponding solutions $y \in \mathcal{X}$. An important property of Problem (12) is the following:

Corollary 1. *The lower bound provided by Lemma 4 is non-decreasing in \bar{v} .*

Proof. This follows immediately from the fact that for $v \geq \bar{v}$:

$$\{y \in \mathcal{X} : \hat{c}^\top y \leq \bar{v} - 1\} \subseteq \{y \in \mathcal{X} : \hat{c}^\top y \leq v - 1\}$$

\square

The resulting lower bounds can be used as a stopping criterion when iterating over all values $v \in \mathcal{C}$. More precisely, iterate through \mathcal{C} starting with $v = c_{\min}$ and increasing it. For each $v \in \mathcal{C}$, we can solve (11) to obtain a feasible weak CE (if the corresponding problem is feasible). We can store the best feasible weak CE $(a_{\text{best}}, b_{\text{best}})$ and the corresponding distance $d_{\text{best}} = \delta(a_{\text{best}}, \hat{a}) + \delta(b_{\text{best}}, \hat{b})$. At the same time, for the current v , we can calculate the lower bound (12) with $\bar{v} = v$. If this lower bound is at least as large as the best known value d_{best} , we can terminate early and $(a_{\text{best}}, b_{\text{best}})$ is an optimal weak CE. The whole procedure is presented in Algorithm 3.

Algorithm 3 Optimal Weak CE for Mutable Constraint Parameters.

Require: $\hat{a}, \hat{b}, \hat{c}, \mathcal{D}, \mathcal{H}_{a,b}$

Ensure: Optimal weak CE (a^*, b^*) .

Best known solution: $(a^*, b^*) = \emptyset$

Best known distance: $d^* = \infty$

Best known lower bound: $lb = -\infty$

Calculate bounds \underline{c}, \bar{c} as in Lemma 1.

for all $v = \underline{c}, \underline{c} + 1, \dots, \bar{c}$ **do**

Set lb to the optimal value of (12) with $\bar{v} = v$.

if $lb \geq d^*$ **then**

Stop and **return** (a^*, b^*) .

else

Calculate opt. solution $(a^*(v), b^*(v))$ of Problem (11) by Algorithm 2.

if $\delta(a^*(v), \hat{a}) + \delta(b^*(v), \hat{b}) < d^*$ **then**

$(a^*, b^*) \leftarrow (a^*(v), b^*(v))$

$d^* \leftarrow \delta(a^*(v), \hat{a}) + \delta(b^*(v), \hat{b})$

end if

end if

end for

return (a^*, b^*)

Remark 2. Before running Algorithm 3, we verify whether (\hat{a}, \hat{b}) is already a weak CE, which then must be the optimal one. This verification involves solving two deterministic integer optimization problems as described in Section 3.1. If the check is not successful, we run Algorithm 3.

Theorem 7. Algorithm 3 calculates an optimal weak CE for the case where only the parameters a and b are mutable.

Proof. The correctness follows from Lemma 1, 2 and 4 and Theorem 6. □

All parameters are mutable. In this section, all parameters in the objective and the constraint are mutable. In this case the weak CE problem (3) can be formulated as follows:

$$\begin{aligned}
& \min_{c,a,b,x} \delta\left((c, a, b), (\hat{c}, \hat{a}, \hat{b})\right) \\
& s.t. \quad c^\top x \leq c^\top y \quad \forall y \in \mathcal{X} : a^\top y \geq b, \\
& \quad \quad a^\top x \geq b, \\
& \quad \quad x \in \mathcal{D} \cap \mathcal{X}, \\
& \quad \quad (c, a, b) \in \mathcal{H}.
\end{aligned}$$

Since the number and type of constraints depend on the decision variables a and b , the algorithms from the previous sections are not applicable. However, (3) can be reformulated as

$$\begin{aligned} \min_{c,a,b,x,z} \quad & \delta((c, a, b), (\hat{c}, \hat{a}, \hat{b})) \\ \text{s.t.} \quad & c^\top x \leq c^\top y + Mz_y \quad \forall y \in \mathcal{X}, \end{aligned} \tag{13a}$$

$$a^\top y \geq b - Mz_y \quad \forall y \in \mathcal{X}, \tag{13b}$$

$$a^\top y \leq b - 1 + M(1 - z_y) \quad \forall y \in \mathcal{X}, \tag{13c}$$

$$x \in \mathcal{D} \cap \mathcal{X}, \tag{13d}$$

$$a^\top x \geq b, \tag{13e}$$

$$z_y \in \{0, 1\} \quad \forall y \in \mathcal{X}, \tag{13f}$$

$$(c, a, b) \in \mathcal{H}, \tag{13g}$$

where M is a sufficiently large value. Constraints (13b) and (13c) ensure that $z_y = 0$ if and only if y is feasible for the choice of constraint parameters (a, b) , and $z_y = 1$, otherwise. Then, Constraints (13a) ensure that the chosen solution x has the best objective value below all feasible solutions $y \in \mathcal{X}$, while for all infeasible $y \in \mathcal{X}$ the constraint is redundant due to the big-M value. Constraints (13d) and (13e) ensure that the solution x is feasible for the chosen parameters a, b and lies in \mathcal{D} .

For linear representable distance functions δ , formulation (13) is bilinear, and thus can be solved by state-of-the-art optimization solvers. However, it can contain an exponential number of constraints. Similar to the previous algorithms it can be solved by iteratively generating solutions $y \in \mathcal{X}$, corresponding variables z_y and corresponding constraints. The procedure is presented in Algorithm 4.

Theorem 8. *Algorithm 4 calculates an optimal weak CE.*

Proof. This follows analogously to Theorem 7. □

Remark 3. *We restricted the analysis of this section to the case of a single mutable constraint. The reason is that in the master problems (11) and (13) we model the condition that certain solutions y are infeasible. For multiple mutable constraints this would involve modeling the condition $Ax \not\geq b$ which involves the introduction of big-M constraints, leading to the equivalent formulation*

$$\begin{aligned} a_i^\top x &\leq b - 1 + M(1 - z_i), \quad i = 1, \dots, m, \\ \sum_{i=1}^m z_i &\geq 1, \\ z &\in \{0, 1\}^n. \end{aligned}$$

In our iterative framework, this results in generating a large set of big-M constraints and corresponding binary variables. Since the resulting methods are computationally intractable even for few dimensions, we refrain from presenting this case in detail. On the other hand, for the case of a single mutable constraint the condition $a^\top x \not\geq b$ can be easily formulated as $a^\top x \leq b - 1$.

4.2 Strong Counterfactual Explanations

As in the case of weak CE algorithms, we can develop three algorithms corresponding to the three cases of the mutable parameter space listed in Section 2.

Algorithm 4 Solve Problem (13)

Require: $\hat{a}, \hat{b}, \hat{c}, \mathcal{D}, \mathcal{X}, \mathcal{H}$

Ensure: Optimal solution (a^*, b^*, c^*) of (13).

set $\mathcal{Y} \leftarrow \emptyset$

optimal \leftarrow false

while not optimal **do**

 Calculate an optimal solution $(\tilde{a}, \tilde{b}, \tilde{c}, \tilde{x})$ of problem

$$\begin{aligned} \min_{c,a,b,x,z} \quad & \delta\left((c, a, b), (\hat{c}, \hat{a}, \hat{b})\right) \\ \text{s.t.} \quad & c^\top x \leq c^\top y + Mz_y \quad \forall y \in \mathcal{Y}, \\ & a^\top y \geq b - Mz_y \quad \forall y \in \mathcal{Y}, \\ & a^\top y \leq b - 1 + M(1 - z_y) \quad \forall y \in \mathcal{Y}, \\ & x \in \mathcal{D} \cap \mathcal{X}, \\ & a^\top x \geq b, \\ & z_y \in \{0, 1\} \quad \forall y \in \mathcal{Y}, \\ & (c, a, b) \in \mathcal{H}. \end{aligned}$$

 Calculate an optimal solution \tilde{y} of the separation problem

$$\begin{aligned} \min \quad & \tilde{c}^\top y \\ \text{s.t.} \quad & \tilde{a}^\top y \geq \tilde{b}, \\ & y \in \mathcal{X}. \end{aligned}$$

if $\tilde{c}^\top \tilde{y} < \tilde{c}^\top \tilde{x}$ **then**

$\mathcal{Y} \leftarrow \mathcal{Y} \cup \{\tilde{y}\}$

else

 optimal \leftarrow true

end if

end while

return $(\tilde{a}, \tilde{b}, \tilde{c})$

▷ End While Loop

Only objective function parameters are mutable. In this subsection, only the objective parameters are mutable, i.e., \hat{a} and \hat{b} are fixed. The mutable parameter space is denoted as $\mathcal{H}_c \subset \mathbb{R}^n$. In this case, we can reformulate (5) as

$$\min_{x,c} \delta \left((c, \hat{a}, \hat{b}), (\hat{c}, \hat{a}, \hat{b}) \right) \quad (14a)$$

$$s.t. \quad c^\top x \leq c^\top y - 1 \quad \forall y \in \mathcal{X} \setminus \mathcal{D} : \hat{a}^\top y \geq \hat{b}, \quad (14b)$$

$$\hat{a}^\top x \geq \hat{b}, \quad (14c)$$

$$x \in \mathcal{D} \cap \mathcal{X}, \quad (14d)$$

$$c \in \mathcal{H}_c, \quad (14e)$$

where Constraints (14c) and (14d) ensure that a solution x from \mathcal{D} is feasible, while Constraints (14b) ensure that any feasible solution in $\mathcal{X} \setminus \mathcal{D}$ cannot be optimal. It follows that for an optimal solution c^* of (14) all optimal solutions must be contained in \mathcal{D} . Note that (14) can have one constraint for every feasible solution $y \in \mathcal{X} \setminus \mathcal{D}$, of which there can be exponentially many. Hence, similar to Algorithm 1 we iteratively generate the constraints by finding the most violating constraint via the separation problem

$$\begin{aligned} & \min c^\top y \\ s.t. & \quad \hat{a}^\top y \geq \hat{b}, \\ & \quad y \in \mathcal{X} \setminus \mathcal{D}. \end{aligned}$$

Note that the essential difference to the weak CE algorithm is the constraint $y \in \mathcal{X} \setminus \mathcal{D}$ in the separation problem. For simple cases of \mathcal{D} (e.g., if \mathcal{D} is described by a single constraint or by variable fixations), this condition can be modeled by linear constraints. However, in general this condition may involve the use of big-M constraints, similarly as described in Remark 3.

Only constraint parameters are mutable. The ideas for the weak CEs can be extended to the strong CE case. We assume again that \hat{c} only has integer entries and define the set \mathcal{C} as in (10). Analogously to Lemma 2, we can verify the following result: For every $v \in \mathcal{C}$, the following problem is either infeasible or an optimal solution thereof is a (not necessarily optimal) strong counterfactual explanation:

$$\min_{a,b,x} \delta \left((\hat{c}, a, b), (\hat{c}, \hat{a}, \hat{b}) \right) \quad (15a)$$

$$s.t. \quad a^\top y \leq b - 1 \quad \forall y \in \mathcal{X} \setminus \mathcal{D} : \hat{c}^\top y \leq v, \quad (15b)$$

$$\hat{c}^\top x = v, \quad (15c)$$

$$a^\top x \geq b, \quad (15d)$$

$$x \in \mathcal{D} \cap \mathcal{X}, \quad (15e)$$

$$(a, b) \in \mathcal{H}_{a,b}. \quad (15f)$$

If the strong CE problem is feasible, Constraints (15c) – (15e) ensure that there exists a solution $x \in \mathcal{D}$ which is feasible and this solution has objective value v . Constraints (15b) ensure that every solution, which is not in \mathcal{D} and has objective value at least as good as v , is not feasible. Hence, feasible solutions outside of \mathcal{D} are not optimal.

The same result of Lemma 3 holds for the strong CE version, i.e., there exists a $v^* \in \mathcal{C}$, such that any optimal solution of (15) is an optimal solution of (5). Likewise, similar to the lower bound

in Lemma 4, it is easy to see that the following problem provides a lower bound for every problem (15) with integer $v \geq \bar{v}$:

$$\min_{a,b} \delta \left((\hat{c}, a, b), (\hat{c}, \hat{a}, \hat{b}) \right) \quad (16a)$$

$$s.t. \quad a^\top y \leq b - 1 \quad \forall y \in \mathcal{X} \setminus \mathcal{D} : \hat{c}^\top y \leq \bar{v}, \quad (16b)$$

$$(a, b) \in \mathcal{H}_{a,b}. \quad (16c)$$

Hence Algorithm 2 and Algorithm 3 can be directly adapted for the strong CE case by replacing each subproblem with the adapted version presented above. The details of the corresponding algorithms are presented in the Appendix as Algorithm ?? and Algorithm ??.

All parameters are mutable. In this section all parameters in the objective and the constraints are mutable. In this case, the strong CE problem (5) can be reformulated as

$$\min_{c,a,b,x} \delta \left((c, a, b), (\hat{c}, \hat{a}, \hat{b}) \right) \quad (17a)$$

$$s.t. \quad c^\top x \leq c^\top y - 1 \quad \forall y \in \mathcal{X} \setminus \mathcal{D} : a^\top y \geq b, \quad (17b)$$

$$a^\top x \geq b, \quad (17c)$$

$$x \in \mathcal{D} \cap \mathcal{X}, \quad (17d)$$

$$(c, a, b) \in \mathcal{H}. \quad (17e)$$

The main difference to the weak CE case is that (17b) ensures that all solutions which are not in \mathcal{D} cannot be optimal. As for the weak CEs, we can then reformulate Problem (17) as

$$\min_{c,a,b,x,z} \delta \left((c, a, b), (\hat{c}, \hat{a}, \hat{b}) \right) \quad (18a)$$

$$s.t. \quad c^\top x \leq c^\top y - 1 + Mz_y \quad \forall y \in \mathcal{X} \setminus \mathcal{D}, \quad (18b)$$

$$a^\top y \geq b - Mz_y \quad \forall y \in \mathcal{X} \setminus \mathcal{D}, \quad (18c)$$

$$a^\top y \leq b - 1 + M(1 - z_y) \quad \forall y \in \mathcal{X} \setminus \mathcal{D}, \quad (18d)$$

$$x \in \mathcal{D} \cap \mathcal{X}, \quad (18e)$$

$$a^\top x \geq b, \quad (18f)$$

$$z_y \in \{0, 1\} \quad \forall y \in \mathcal{X}, \quad (18g)$$

$$(c, a, b) \in \mathcal{H}, \quad (18h)$$

where M is a large enough big-M value. Constraints (18c) and (18d) ensure that $z_y = 0$ if and only if y is feasible for the choice of constraint parameters (a, b) and $z_y = 1$ otherwise. Hence, Constraints (18b) ensure that the chosen solution x has the best objective value below all feasible solutions $y \in \mathcal{X}$ while for all infeasible $y \in \mathcal{X}$ the constraint is redundant due to the big-M value. Constraints (18e) and (18f) ensure that the selected solution x is feasible for the chosen parameters a, b and lies in \mathcal{D} .

Analogously to the weak CEs, formulation (18) is bilinear and can be solved by state-of-the-art optimization solvers. However, it can contain an exponential number of constraints. Similar to the previous algorithms it can be solved by iteratively generating solutions $y \in \mathcal{X}$, corresponding variables z_y and corresponding constraints. The procedure is presented in Algorithm ?? in the Appendix.

5 Computational Study

In this section, we perform numerical experiments on a set of test problems to evaluate the performance of the proposed algorithms in obtaining both weak and strong CEs.

5.1 Setup of Knapsack Instances

Data. We tested the algorithms on knapsack benchmark instances from [kplib](https://github.com/likr/kplib)¹, which we interpreted as cover instances to be consistent with the notation throughout this work. The instances were originally designed according to the procedure outlined in Kellerer, Pferschy & Pisinger [KPP04]. We used instances with a data range of 1000² and instance sizes in [10, 20, 30, 40], which were generated by slicing larger instances. Since it is well-known that the practical hardness of knapsack problems depends on the instance structure [Pis05, SMCM21], e.g., if items value/weight ratios are correlated, the problems tend to be harder to solve in practice [Pis05], we use both 20 uncorrelated and strongly correlated instances each. In total, we consider $4 \times 2 \times 20 = 160$ testing instances per favored solution space and CE type.

Favored solution spaces. Let $\mathcal{I}^+, \mathcal{I}^- \subset \{1, \dots, n\}$ with $\mathcal{I}^+ \cap \mathcal{I}^- = \emptyset$ be index sets of variables that are to be fixed to 1 or 0, respectively. Then, we define two sets for *positive and negative fixations* as $\mathcal{D}^+ = \{x \in \{0, 1\}^n : x_i = 1 \ \forall i \in \mathcal{I}^+\}$ and $\mathcal{D}^- = \{x \in \{0, 1\}^n : x_i = 0 \ \forall i \in \mathcal{I}^-\}$ respectively. In our computational experiments, the number of elements to be fixed is set to the following fraction:

$$\left\lceil 0.1 \cdot \frac{b}{\frac{1}{n} \sum_{i \in \bar{n}} a_i} \right\rceil.$$

Intuitively, this means 10% of the items in an expected cover solution are fixed to some value. In each case, the present cover problem was first solved. Then, the favored elements were randomly selected among the items that were not part of the present solution.

We also define a favored solution space \mathcal{D}^\geq for given parameters $\alpha \in \{0, 1\}^n$ and $\beta \in \mathbb{N}$:

$$\mathcal{D}^\geq = \{x \in \{0, 1\}^n : \sum_{i \in \mathcal{I}} \alpha_i x_i \geq \beta\}.$$

For our computational experiments, 10% of all items are randomly selected to be part of the constraint and $\beta = 1$, i.e., at least one item from the constraint set must be selected.

For each of the three types of favored solution spaces, we need to be able to take the complement of \mathcal{D} to model the constraint $y \in \mathcal{X} \setminus \mathcal{D}$ in the separation problems for the strong CEs. The complement of \mathcal{D}^\geq is given by the knapsack constrained set

$$\{x \in \{0, 1\}^n : \sum_{i \in \mathcal{I}} \alpha_i x_i \leq \beta - 1\}.$$

Similarly, the complements of \mathcal{D}^+ and \mathcal{D}^- are

$$\{x \in \{0, 1\}^n : \sum_{i \in \mathcal{I}^+} x_i \leq |\mathcal{I}^+| - 1\},$$

¹See <https://github.com/likr/kplib>.

²Note that Algorithms 3, 4, ?? and ?? are sensitive to changes in the data range by design, i.e., changes in the data range will lead to a proportional change in the number of subproblems, and thus the runtime.

$$\{x \in \{0, 1\}^n : \sum_{i \in \mathcal{I}^-} x_i \geq 1\},$$

respectively.

Mutable parameter space. For the mutable parameter spaces, we consider only settings where parameters can be changed homogeneously by a flat percentage of 5%. For the distance δ we use the sum of absolute deviations, i.e., the ℓ_1 -norm, as the objective function. Thus, in total we get $160 \times 3 \times 2 = 960$ different instances.

Implementation. We implemented all proposed algorithms in Python using Gurobi 12.0.0 as a solver. All data and code are made available upon request. All computations were ran on RWTH Aachen Universities’ High-Performance Computing (HPC) CLAIX on Intel Xeon 8468 Sapphire nodes with a single 2.1 GHz core and 5GB RAM each. We used a maximum runtime of 10h.

5.2 Computational Results on Knapsack Instances

Existence and cost of counterfactual explanations. In total, we found an optimal counterfactual for 668 instances among 960. Of the remainder, in 10 instances of size 30 and in 32 instances of size 40 the runtime of 10h was insufficient to prove optimality or infeasibility. A further 250 instances were flagged as infeasible, generally due to finding an infeasible lower bound subproblem. Of the 250 infeasible instances, only four were strongly correlated, and the remainder were among the uncorrelated instances.

Finding a feasible solution required on average a relative change of $0.25 \pm 0.48\%$ in weights, with a range of $0\% - 3\%$. However, note that for many uncorrelated instances, we did not find a feasible solution. Thus, for those instances, a larger mutable parameter space might be necessary. Note that larger parameter spaces imply larger ranges of \mathcal{C} as well. In our study, up to 2590 different potential objective values were considered per instance with an average of 1160 ± 597 values. This excludes the effect of the lower bound proposed in Lemma 4 shown in the next subsection.

Effect of lower bound. We find that the lower bound reduces the number of iterations required to solve an instance to optimality by more than 50%. Figures 2 and 3 illustrate this for strong and weak CEs, respectively.

We also observe that good solutions are centered around the halfway point of 50%, which tends to be close to the original objective value. Thus heuristically searching for small changes to the weights that lead to solutions close to the original objective function value appears to be a promising search strategy. However, this sacrifices the lower bound from Lemma 4 that significantly accelerates the exact approach. The observations above apply for both weak and strong CEs equally.

Scaling of runtime. As to be expected, runtime increases significantly with instance size. Figures 4a and 4b illustrate this. The number of cuts required scales equally with instance size, as can be seen in Figures 4c and 4d. In summary, CE can be computed using the algorithms provided. However, runtime and memory demands scale exponentially with instance size, which is to be expected for a Σ_2^P -complete problem.

We can also illustrate this using primal-dual gaps. For that we restrict the data displayed to the 1st and 2nd item from each size, i.e 10, 20, 30, and 40. This is shown in Figure 5.

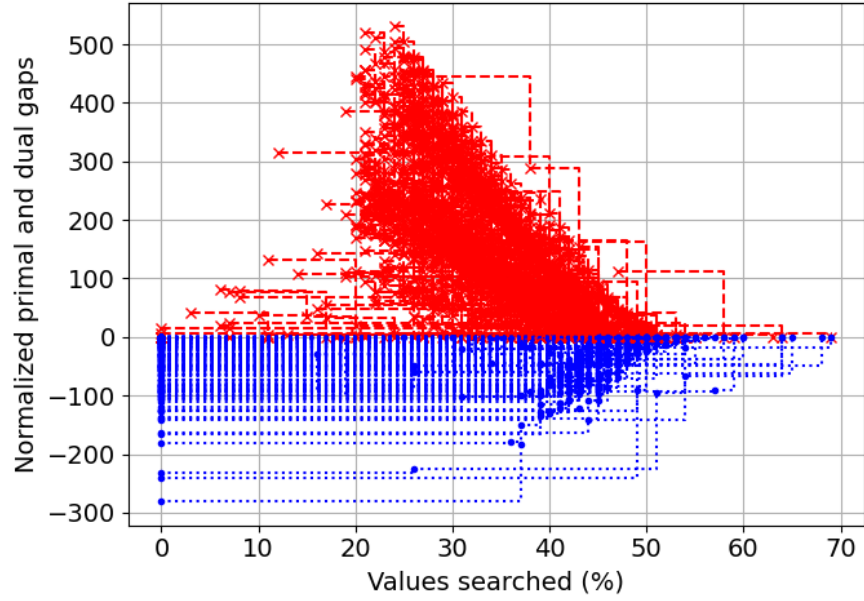


Figure 2: Progress towards optimal solution for 480 strong CE instances of size 10, 20, 30, 40 that were solved to optimality. Red are primal and blue are dual bounds. The x -axis is normalized to the total number of values in \mathcal{C} . The y -axis is normalized around each instance's optimum.

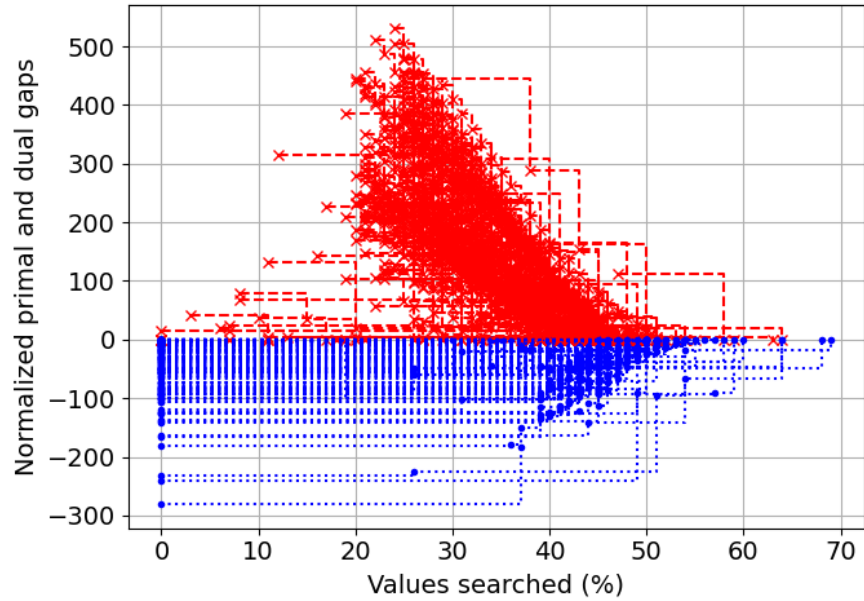


Figure 3: Progress towards solution for 480 weak CE instances of size 10, 20, 30, 40 that were solved to optimality. Red are primal and blue are dual bounds. The x -axis is normalized to the total number of values in \mathcal{C} . The y -axis is normalized around each instance's optimum.

Deriving insights from counterfactuals. Computing CEs can also enable us to quantify (partial) decision values in terms of their contribution to an objective. Although strongly correlated instances are generally considered to be harder, the fact that all items have a similar cost/weight ratio means small changes in weights are sufficient to add/remove an element from a solution.

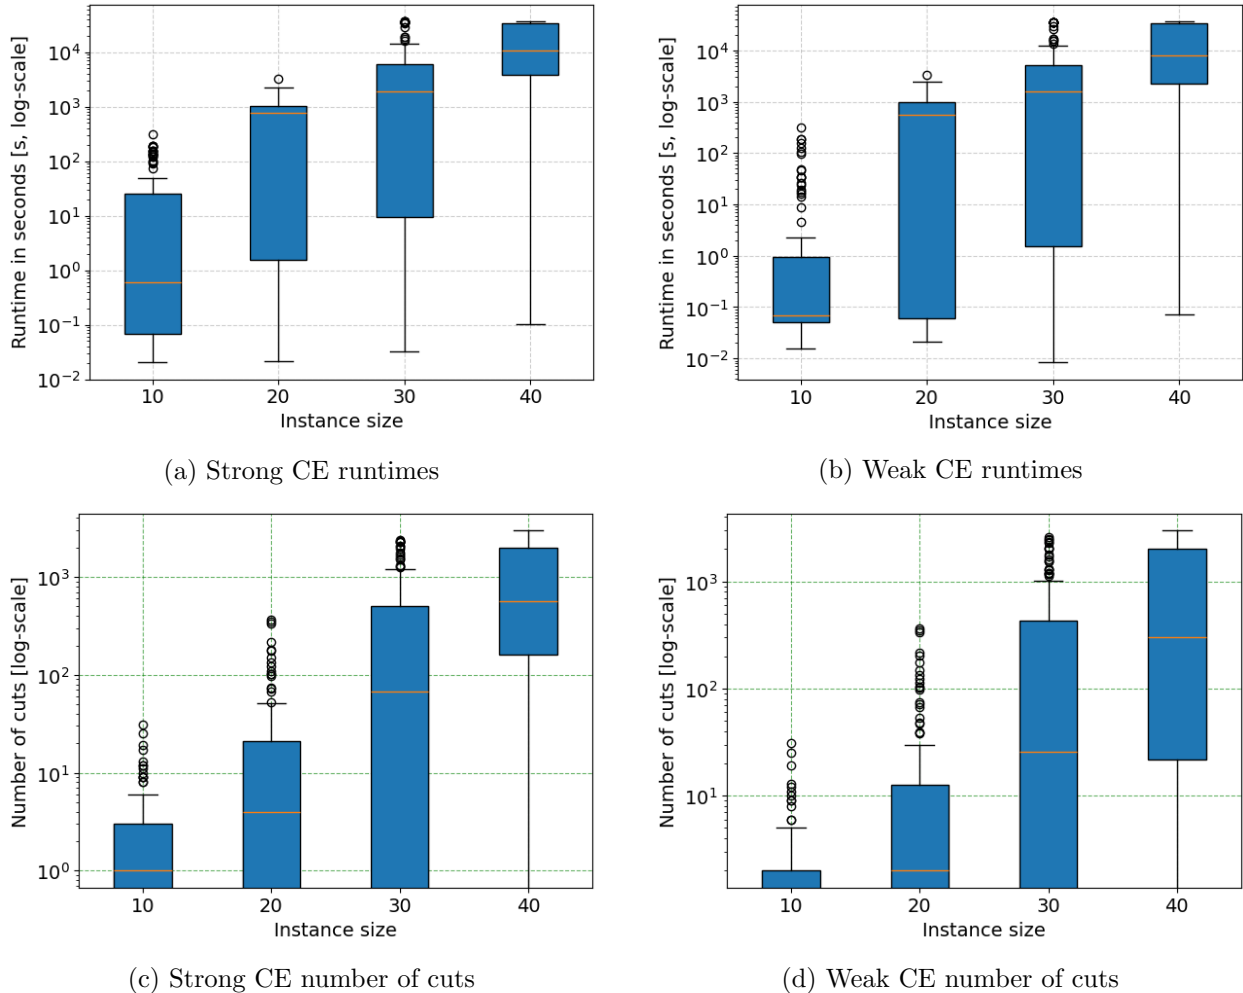


Figure 4: Boxplot comparison of runtimes (top row) and total number of generated cuts (bottom row) for 480 strong and weak CE instances with sizes 10, 20, 30, and 40. Instances not solved to optimality are assigned 36000s (10h).

As an example, consider CEs for the cover instance in Table 1.

Table 1: Exemplary strongly correlated cover instance. For a demand/capacity of 2358, an optimal cover picks items 1, 2, and 7 for a total objective of 2701.

Item	0	1	2	3	4	5	6	7	8	9
Weight	135	848	764	256	496	450	652	789	94	29
Costs	235	948	864	356	596	550	752	889	194	129

We compute a strong CE for every item that is not currently part of the optimal solution, with a positive fixation for each. More precisely, for each item i that is not part of the optimal solution we calculate the minimal parameter change needed such that item i is included in all optimal solutions. The results are given in Table 2.

We observe different types of modifications: To enforce item 0, the incumbent solution is made less attractive, thus item 7 is then substituted by items 0 and 6. For items 4,5 and 6, we see




Figure 5: Progress towards solution for 64/96 strong CE instances solved to optimality. Blue instances are size 40, red instances are size 30, yellow and green correspond to sizes 20, 10. The x -axis is normalized to the total number of values in \mathcal{C} . The y -axis is normalized around each instance’s optimum.

Table 2: Comparison of CEs for different items in a strongly correlated instance sorted by cost. The cost/weight ratio is based on the initial weights, items [1, 2, 7] have a ratio of 0.89/0.88/0.88 respectively.

Item	CE cost	$\frac{\text{cost}}{\text{weight}}$ ratio	Changes to weights	New solution
0	44	0.57	1:-9, 2:-5, 7:-30	[0, 1, 2, 6]
4	44	0.83	1:-41, 2:-3	[2, 4, 5, 6]
5	44	0.81	1:-42, 7:-2	[2, 4, 5, 6]
6	44	0.87	1:-42, 7:-2	[2, 4, 5, 6]
8	51	0.48	2:-5, 4:-2, 6:+1, 7:-39, 8:+4	[1, 2, 6, 8]
3	64	0.72	2:-38, 5:+21, 7:-6	[1, 3, 5, 7]
9	65	0.22	2:+42, 3:+22, 9:+1	[1, 2, 6, 9]

that CEs themselves contain alternatives, e.g., to ensure either item is part of an optimal solution multiple equivalent changes of weights lead to the new solution of [2, 4, 5, 6]. Item 8 is significantly less attractive, enforcing it to be part of a solution requires a higher CE cost, even more so for items 3 and 9. Finally, we observe that simple metrics such as cost/weight are insufficient to capture modeling dynamics – item 9 is very good in terms of cost/weight ratio, but still unattractive for optimal solutions. Thus, CEs allow for more differentiated evaluation of (partial) solutions.

5.3 An Application to the Resource Constrained Shortest Path Problem

In many cases, even considering few mutable constraints allows us to derive insights into an instance. Next, we examine the following setting: The task is to route traffic through a given graph, where using an arc comes with a toll. Given a limited toll budget this problem can be modeled as a

Resource Constrained Shortest Path (RCSP). A CE in the toll parameters asks now how we have to change the toll such that an optimal path would use a certain edge or not.

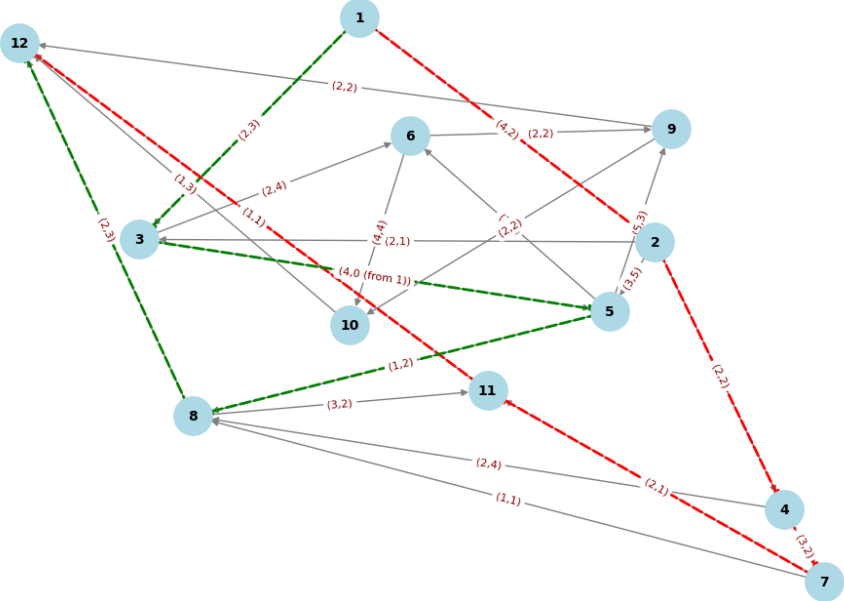


Figure 6: CE for RCSP problem. Edge values indicate (length, toll). The red path is the original shortest path for a toll budget of 8; the least change in toll that leads to (1,2) no longer being part of an optimal solution is given by lowering the toll of (3,4) by one unit, resulting in the green RCSP.

Figure 6 illustrates a counterfactual in such a setting for a toy example. If we assume a toll budget of 8 units, initially, the red path is the shortest path, with total length 12. Assume the arc (1,2) goes through the city center and we would like to avoid routing traffic through there whenever possible. Our question would then be: What is the smallest change so that the arc (1,2) is no longer part of a shortest path? The result is given by the green path, which has total length 9 and requires a single change to the data, i.e., lowering the toll of arc (3,5) by one unit. In our example, this would mean that if we wanted to avoid routing traffic through the city center (1,2), the best way to do so is to waive the toll on (3,5) making it more attractive to bypass the city center.

In terms of mathematics, note that the RCSP contains multiple flow conservation constraints. However, since those are immutable, they do not affect our algorithm. Furthermore, we can bound the c values by finding the optimal RCSP for each feasible value of b , i.e., the toll budget.

Solving RCSP pricing problems. Finally, we test our algorithm on a set of benchmark instances from literature. For that, we use the SPPRCLIB data set provided by Jepsen et al. [JPS08]. The instances represent pricing problems from solving capacitated vehicle routing problems with column generation, see [FLL⁺06]. As such, they contain negative node weights, as well as a resource constraint on the nodes. In Appendix ??, we provide a ILP formulation of such problems. We restrict ourselves to the case that source and sink are the first and last indexed node, respectively, and round node weights to the nearest thousand and scaled down accordingly. The instances are complete graphs with 45 – 262 nodes. We investigate what would be the least change to the weights in the resource constraint by a maximum of 10% to exclude the current first node from the shortest

path. In the context of a pricing problem, this would imply computing a least change to the duals that would have led to a different column being added.

Of 44 instances, we found a weak CE for 22. Of these 22 instances, for 5 the original solution was already a counterfactual explanation. Finding a CE took an average of 15 ± 19 s and a maximum of 63s. In about half the cases, the new RSCPs only have small alterations, for the other half, a completely new path is chosen. We note that we can thus compute CEs also for larger and more complex problems, provided the mutable part is limited. In the context of RSCP, this might allow insights into reoccurring patterns or the lack thereof in column generation.

6 Conclusion

In this work, we studied weak and strong counterfactual explanations (CEs) for integer (linear) optimization problems. We showed that in the general case, computing a counterfactual increases complexity by one level in the polynomial hierarchy, that is, NP-complete problems have a counterfactual that is Σ_2^P -complete, which is equivalent to general bilevel integer linear programming. We present algorithms to calculate counterfactual explanations for integer linear problems and implemented a tractable version thereof, which we evaluated on cover problem instances. We found that the proposed approach works well for instances with up to 40 items.

There are several open research directions to extend our work. Further research could consider and exploit specific (combinatorial) structures of the present problem to detect situations where calculating CEs is on a lower rank of the polynomial hierarchy and computationally less demanding. Developing heuristics also presents another promising avenue for further research. These methods can deliver computationally efficient solutions that closely approximate the optimal results. Beyond their standalone utility, these heuristic solutions can also serve as effective starting points for the exact algorithms that we proposed here.

7 Code and Data Disclosure

The code and data to support the numerical experiments in this paper can be found at URL.

Acknowledgements

Computations were performed with computing resources granted by RWTH Aachen University High Performance computing. We thank Yosuke Onoue, who build and hosts the `klib` knapsack instance library. We thank Christina Büsing and Hector Geffner for enabling this research collaboration, as well as Christoph Grüne for insightful discussions on complexity theory.

References

- [AB07] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2007.
- [BBD⁺25] Jeroen Beliën, Sally Brailsford, Erik Demeulemeester, Derya Demirtas, Erwin W. Hans, and Paul Harper. Fifty years of operational research applied to healthcare. *European Journal of Operational Research*, 326(2):189 – 206, 2025.

- [BRAM24] Peter Bogetoft, Jasone Ramírez-Ayerbe, and Dolores Romero Morales. Counterfactual analysis and target setting in benchmarking. *European Journal of Operational Research*, 315(3):1083–1095, 2024.
- [BTCCH] Alberto Boggio Tomasaz, Margarida Carvalho, Roberto Cordone, and Pierre Hosteins. On the Completeness of Several Fortification-Interdiction Games in the Polynomial Hierarchy. *Mathematics of Operations Research*.
- [CDB10] Brecht Cardoen, Erik Demeulemeester, and Jeroen Beliën. Operating room planning and scheduling: A literature review. *European Journal of Operational Research*, 201(3):921–932, 2010.
- [CMZ25] Timothy CY Chan, Rafid Mahmood, and Ian Yihang Zhu. Inverse optimization: Theory and applications. *Operations Research*, 73(2):1046–1074, 2025.
- [Dem02] Stephan Dempe. *Foundations of Bilevel Programming*. Springer, 2002.
- [DKPVK15] Stephan Dempe, Vyacheslav Kalashnikov, Gerardo A Pérez-Valdés, and Nataliya Kalashnykova. Bilevel Programming Problems. *Energy Systems. Springer, Berlin*, 10:978–3, 2015.
- [FLL⁺06] Ricardo Fukasawa, Humberto Longo, Jens Lygaard, Marcus Poggi de Aragão, Marcelo Reis, Eduardo Uchoa, and Renato F. Werneck. Robust Branch-and-Cut-and-Price for the Capacitated Vehicle Routing Problem. *Mathematical Programming*, 106(3):491–511, 2006.
- [Fre04] Arnaud Freville. The multidimensional 0–1 knapsack problem: An overview. *European Journal of Operational Research*, 155(1):1–21, 2004.
- [GOQ24] Sharmi Dev Gupta, Barry O’Sullivan, and Luis Quesada. Counterfactual explanation through constraint relaxation. In *2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 396–403. IEEE, 2024.
- [JPS08] Mads Kehlet Jepsen, Bjørn Petersen, and Simon Spoorendonk. *A Branch-and-Cut Algorithm for the Elementary Shortest Path Problem with a Capacity Constraint*. Number 08-01 in Technical Report. Department of Computer Science, University of Copenhagen, 2008.
- [Kar72] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.
- [KB21] Anton Korikov and J Christopher Beck. Counterfactual Explanations via Inverse Constraint Programming. In *27th International Conference on Principles and Practice of Constraint Programming (CP 2021)*, pages 35–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.
- [KB23] Anton Korikov and J Christopher Beck. Objective-Based Counterfactual Explanations for Linear Discrete Optimization. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 18–34. Springer, 2023.
- [KIBd25] Jannis Kurtz, Ş. İlker Birbil, and Dick den Hertog. Counterfactual explanations for linear optimization. *European Journal of Operational Research*, 2025.

- [KLLS21] Thomas Kleinert, Martine Labbé, Ivana Ljubić, and Martin Schmidt. A Survey on Mixed-Integer Programming Techniques in Bilevel Optimization. *EURO Journal on Computational Optimization*, 9:100007, 2021.
- [KPP04] Hans Kellerer, Ulrich Pferschy, and David Pisinger. Exact Solution of the Knapsack Problem. In *Knapsack Problems*, pages 117–160. Springer, Berlin, Heidelberg, 2004.
- [KSB21] Anton Korikov, Alexander Shleyfman, and J Christopher Beck. Counterfactual Explanations for Optimization-Based Decisions in the Context of the GDPR. In *ICAPS 2021 Workshop on Explainable AI Planning*, 2021.
- [LS24] Henri Lefebvre and Martin Schmidt. Computing Counterfactual Explanations for Linear Optimization: A New Class of Bilevel Models and a Tailored Penalty Alternating Direction Method. Technical report, 2024. Published: 2024-12-19, Updated: 2025-01-08.
- [MKR⁺24] Donato Maragno, Jannis Kurtz, Tabea E. Röber, Rob Goedhart, Ş. İlker Birbil, and Dick den Hertog. Finding Regions of Counterfactual Explanations via Robust Optimization. *INFORMS Journal on Computing*, 36(5):1316–1334, 2024.
- [MRIB22] Donato Maragno, Tabea E. Röber, and Ş. İlker Birbil. Counterfactual Explanations Using Optimization With Constraint Learning. In *OPT 2022: Optimization for Machine Learning (NeurIPS 2022 Workshop)*, 2022.
- [PDM17] Alberto Del Pia, Santanu S Dey, and Marco Molinaro. Mixed-integer quadratic programming is in NP. *Mathematical Programming*, 162(1):225–240, 2017.
- [Pis05] David Pisinger. Where are the hard knapsack problems? *Computers & Operations Research*, 32(9):2271–2284, 2005.
- [RAF25] Jasone Ramírez-Ayerbe and Emma Frejinger. Relative Explanations for Contextual Problems with Endogenous Uncertainty: An Application to Competitive Facility Location. *arXiv preprint*, 2025.
- [SMCM21] Kate Smith-Miles, Jeffrey Christiansen, and Mario Andrés Muñoz. Revisiting where are the hard knapsack problems? via Instance Space Analysis. *Computers & Operations Research*, 128:105184, 2021.
- [SWdF21] VF Stienen, JC Wagenaar, Dick den Hertog, and HA Fleuren. Optimal depot locations for humanitarian logistics service providers using robust optimization. *Omega*, 104:102494, 2021.
- [VBH⁺24] Sahil Verma, Varich Boonsanong, Minh Hoang, Keegan Hines, John Dickerson, and Chirag Shah. Counterfactual Explanations and Algorithmic Recourses for Machine Learning: A Review. *ACM Comput. Surv.*, 56(12), October 2024.
- [Wan13] Lizhi Wang. Branch-and-bound algorithms for the partial inverse mixed integer linear programming problem. *Journal of Global Optimization*, 55(3):491–506, March 2013.
- [WTKR13] Wolfram Wiesemann, Angelos Tsoukalas, Polyxeni-Margarita Kleniati, and Berç Rustem. Pessimistic Bilevel Optimization. *SIAM Journal on Optimization*, 23(1):353–380, 2013.