

Inverse Mixed Integer Optimization: Polyhedral Insights and Trust Region Methods

Merve Bodur, Timothy C. Y. Chan, Ian Yihang Zhu

Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario M5S 3G8, Canada

Inverse optimization – determining parameters of an optimization problem that render a given solution optimal – has received increasing attention in recent years. While significant inverse optimization literature exists for convex optimization problems, there have been few advances for discrete problems, despite the ubiquity of applications that fundamentally rely on discrete decision-making. In this paper, we present a new set of theoretical insights and algorithms for the general class of inverse mixed integer linear optimization problems. Specifically, a general characterization of optimality conditions is established and leveraged to design new cutting plane solution algorithms. Through an extensive set of computational experiments, we show that our methods provide substantial improvements over existing methods in solving the largest and most difficult instances to date.

Key words: Inverse Optimization, Mixed Integer Programming, Cutting Planes, Inverse-Feasibility, Decomposition Methods, Trust Regions

1. Introduction

Inverse optimization – inferring unobserved parameters of a (*forward*) optimization problem that render a given (*forward-feasible*) solution optimal – has received increasing attention in recent years. Most of the literature has focused on inverse optimization for convex forward optimization problems, allowing the bi-level inverse problem to be reformulated as a single-level convex problem using the Karush-Kuhn-Tucker conditions. Examples include linear (e.g., Zhang and Liu 1996, Ahuja and Orlin 2001), multiobjective linear (e.g. Naghavi et al. 2019), conic (e.g. Iyengar and Kang 2005) and general convex optimization problems (e.g. Zhang and Xu 2010, Zhang et al. 2010). These models have been applied across a wide range of application domains (Chan et al. 2021) and have also been extended methodologically in several modern directions including estimating model parameters from multiple observed solutions while considering concepts such as statistical consistency (Aswani et al. 2018) or distributional robustness (Esfahani et al. 2018).

In contrast, there have been very few advances in inverse optimization for discrete forward problems, despite the ubiquity of methodological research in discrete optimization and real-world applications that have discrete decisions. Part of the challenge may stem from the fact that the

optimality conditions for discrete problems like mixed integer linear optimization (MILO) generally do not lead to computationally tractable solution algorithms like they do for convex problems. For example, Schaefer (2009) and Lamperski and Schaefer (2015) demonstrate how superadditive duality can be used to develop exact reformulations of bi-level inverse integer and mixed-integer optimization problems, respectively. However, the resulting single-level formulation is an exponentially large linear program that is intractable beyond problems with few variables.

The other main idea in the literature for a general-purpose inverse MILO solution method is a cutting plane algorithm. First proposed in Wang (2009), the idea is to decompose the inverse MILO problem into a master problem and a subproblem. The former is a relaxation of the inverse MILO problem that provides candidates for parameters to be inferred, whereas the latter identifies *extreme points* of the convex hull of the forward-feasible region that generate cuts to send back to the master problem. While it was shown that small problems could be solved efficiently, this approach does not scale well to larger problems. The computational cost of computing each extreme point is large in general, since it involves solving an instance of the forward MILO problem. The only advance to this method since its development was a proposed heuristic that parallelizes the computation of extreme points (Duan and Wang 2011).

In this paper, we develop a new cutting plane framework for solving inverse MILO problems. We begin by providing a more general characterization of the optimality conditions of the inverse MILO problem relative to previous literature. We demonstrate that extreme points are sufficient but not necessary for characterizing inverse optimality, and that more parsimonious representations of the problem exist. Using this insight, we develop a new cutting plane method where the cut generation subroutine efficiently identifies interior forward-feasible points using trust regions. These cuts come at significantly lower computational cost compared to cuts generated from extreme points. In addition, we observe that inverse MILO problems can be solved with many fewer cuts compared to the classical cutting plane algorithm.

Our specific contributions are:

1. We present a novel characterization of optimality conditions for inverse MILO problems by introducing the concept of *generator sets* (Section 2). Generator sets characterize the complete family of reformulations of the MILO problem that preserve the feasible region of the inverse optimization problem. Our main result provides necessary and sufficient conditions for any set to be a generator set. We also show that there exist “small” generator sets, which are attractive from a computational perspective.
2. Motivated by our theoretical insights, we propose a new family of cutting plane methods that restrict the feasible region of the MILO problem using trust regions (Section 3). We also propose computational enhancements that further speed up the algorithm (Section 4). Our algorithms generalize the classical cutting plane approach to solving inverse MILO problems.

3. We propose an extension of the inverse MILO problem to a setting with multiple input data points and illustrate how our cutting plane algorithm can be easily extended to solve this problem (Section 5).
4. Through a comprehensive numerical study over instances drawn from the MIPLIB 2017 benchmark library (Gleixner et al. 2019), we demonstrate that our new cutting plane algorithm significantly outperforms the state-of-the-art algorithm by solving the largest and most difficult problem instances in the literature to date (Sections 6 and 7). We observe that our new cut generation methods can both substantially speed up cut generation and reduce the number of cuts required.

1.1. Related literature

Below, we review relevant inverse optimization models and solution methods, as well as the main application areas in which these models are considered. A more general overview and survey of inverse optimization can be found in Chan et al. (2021).

1.1.1. Models and methods. The classical inverse optimization problem, which is to infer a set of model parameters that render a given feasible solution optimal, has been studied over a wide range of problem settings. The parameters that are inferred can be the cost vector or the constraints. While there are a few studies that focus on estimating constraint parameters (Güler and Hamacher 2010, Birge et al. 2017, Chan and Kaw 2020), the vast majority of papers focus on estimating the cost vector. Most early works focused on deriving algorithms to solve inverse models for particular forward problems (e.g., shortest path, minimum spanning tree), while more recent works consider general-purpose methods for broader classes of forward problems (e.g., mixed integer optimization). Heuberger (2004) provides an overview of the early works whereas the literature discussed in the introduction are examples of more recent works. In this paper, we focus on estimating the cost vector for general mixed integer forward problems, particularly in the form of the classical inverse model.

Inverse optimization problems with “noisy” data have also been studied (e.g. Troutt et al. 2006, Keshavarz et al. 2011, Chan et al. 2014, 2019). These problems are characterized by the fact that there does not exist a set of (non-trivial) parameters that render a given solution, or a set of solutions, optimal. For these problems, the cost vector is estimated by minimizing various notions of suboptimality (e.g. Bertsimas et al. 2015, Aswani et al. 2018, Esfahani et al. 2018, Babier et al. 2021). Although we focus predominantly on the classical inverse model in the paper, we show that the insights and methods developed can be easily extended to a model for multiple “noisy” data points that resembles the models found for inverse convex optimization.

Finally, we borrow the term “trust region” from the general optimization community, a term which broadly describes the restriction of a particular search space. The use of trust region concepts have appeared in a range of different domains such as nonlinear optimization and, more relevant to our paper, decomposition algorithms. Within the latter domain, trust regions have been applied to the master problems of decomposition models for stochastic programming, and are shown to help generate stronger cuts in this context (Linderoth and Wright 2003, Zverovich et al. 2012, Rahmaniyan et al. 2017). In contrast, we apply trust regions to the cut generation problem in our decomposition framework, which leads to improvements in both strength of cuts *and* cut generation time over a wide array of problem structures.

1.1.2. Applications. Inverse optimization has been studied in a wide variety of applications to estimate latent parameters and infer subjective preferences using observed decision data. Examples of such applications can be found in energy markets (Ruiz et al. 2013, Saez-Gallego et al. 2016, Birge et al. 2017), healthcare (Erkin et al. 2010, Chan et al. 2014, Babier et al. 2020), finance (Bertsimas et al. 2012, Utz et al. 2014, Yu et al. 2020), and transportation (Chow and Recker 2012, Chow et al. 2014, Xu et al. 2018, Zhang et al. 2018), where inverse optimization is used to provide insights into market structure, treatment design, risk aversion and route-choice preferences, respectively. This literature makes use of both classical inverse models and inverse models for “noisy” decision data.

Inverse optimization is also prevalent in the bi-level optimization literature, especially in regards to pricing and incentive design problems. Recent examples of such applications include the design of tolls in traffic and commodity transportation networks (Marcotte et al. 2009, Brotcorne et al. 2011, Esfandeh et al. 2016, Kuiteing et al. 2017, 2018), price schedules and carbon taxes in energy systems (Zhou et al. 2011, Afşar et al. 2021), and profit-sharing mechanisms in carrier alliances (Agarwal and Ergun 2010, Houghtalen et al. 2011). Here, the classical inverse optimization model appears as a “subproblem” within the solution process used to solve these pricing problems. Specifically, a “master problem” generates a decision “target” (e.g., a sustainable routing decision) while an inverse model computes a set of objective perturbations (e.g., tolls on roads) for which the target becomes an optimal decision for an underlying decision-maker (e.g., a transportation company). The master problem can also propose new targets if the inverse subproblem is infeasible (Afşar et al. 2021).

Despite the ubiquity of decision-making models in these application areas that fundamentally rely on making discrete choices, existing literature has focused primarily on inverse optimization for continuous forward problems. The few exceptions to date are in energy planning (Zhou et al. 2011), transportation (Chow and Recker 2012), and sustainability (Turner and Chan 2013). The first two applications were formulated as general inverse MILO problems and applied the classical cutting plane algorithm (Wang 2009). The sustainability application, an inverse knapsack problem, was reformulated using superadditive duality and solved by restricting focus to linear functions.

1.2. Notation

Throughout the paper, vectors and matrices are written in bold, while sets are defined using calligraphic letters. Subscripts denote specific elements of a vector, whereas superscripts denote different vectors. We use $(\cdot)^\top$ to denote the transpose operator. For any set \mathcal{S} , $\text{conv}(\mathcal{S})$ is its convex hull, $\text{ext}(\mathcal{S})$ is its set of extreme points, and $\mathcal{E}(\mathcal{S}) = \text{ext}(\text{conv}(\mathcal{S}))$.

2. Structure of Inverse MILO Problems

In this section, we study the structure of inverse MILO problems. We provide a novel characterization of inverse-feasibility over general non-convex forward optimization problems. We do this through the definition of a *generator set*, which characterizes the complete family of reformulations of the forward optimization problem that preserve the feasible region of the inverse optimization problem. In turn, we show how certain feasible region definitions are amenable to more efficient solution methods for the inverse optimization problem. Proofs of our results are provided in Appendix A.1 in the Online Supplement.

2.1. Problem definition

Our inverse optimization model is based on the following MILO problem, known as the *forward problem*:

$$\begin{aligned} \mathbf{FP}(\mathbf{c}, \mathcal{X}) : & \underset{\mathbf{x}}{\text{minimize}} && \mathbf{c}^\top \mathbf{x} \\ & \text{subject to} && \mathbf{x} \in \mathcal{X} := \{\mathbf{Ax} \geq \mathbf{b}, \mathbf{x} \in \mathbb{Z}^{n-q} \times \mathbb{R}^q\}. \end{aligned} \quad (1)$$

Let $\mathcal{F}(\mathbf{c}, \mathcal{X})$ be the optimal solution set of $\mathbf{FP}(\mathbf{c}, \mathcal{X})$. Elements of \mathcal{X} and $\mathcal{F}(\mathbf{c}, \mathcal{X})$ are called *forward-feasible* and *forward-optimal*, respectively. For a given forward-feasible solution $\hat{\mathbf{x}} \in \mathcal{X}$, let $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$ define the *inverse-feasible region*, which is the set of all cost vectors $\mathbf{c} \in \mathbb{R}^n$ that render the solution $\hat{\mathbf{x}}$ optimal, i.e.,

$$\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) := \{\mathbf{c} \in \mathbb{R}^n \mid \hat{\mathbf{x}} \in \mathcal{F}(\mathbf{c}, \mathcal{X})\}. \quad (2)$$

The inverse optimization problem is defined as the following optimization model,

$$\begin{aligned} \mathbf{IO}(\mathbf{c}^0, \hat{\mathbf{x}}, \mathcal{X}) : & \underset{\mathbf{c} \in \mathcal{P}}{\text{minimize}} && g(\mathbf{c}) \\ & \text{subject to} && \mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}), \end{aligned} \quad (3)$$

which seeks to find a cost vector \mathbf{c} that minimizes an objective function $g(\mathbf{c})$, is inverse-feasible, and satisfies any context-specific constraints $\mathbf{c} \in \mathcal{P}$. A common choice of $g(\mathbf{c})$ is

$$g(\mathbf{c}) := \|\mathbf{c} - \mathbf{c}^0\|_1.$$

The inverse optimization model with this objective seeks to find feasible cost vectors \mathbf{c} that minimize the deviation to a given reference cost vector \mathbf{c}^0 , as measured by the Manhattan distance.

Practically, this objective function is relevant in many applications (see Section 1.1 for details). First, in estimation problems, \mathbf{c}^0 can represent an initial estimate of the cost vector, and given an observed decision $\hat{\mathbf{x}}$, the cost estimates can be refined (minimally) to be compatible with $\hat{\mathbf{x}}$. For example, \mathbf{c}^0 may represent an estimate of network link costs used to describe route preferences (e.g., where initial estimates are based on road distances), and an observed routing decision $\hat{\mathbf{x}}$ can be used to refine these estimates (Burton and Toint 1992, Ahuja and Orlin 2001, Chen et al. 2021). Second, in bi-level optimization problems, \mathbf{c}^0 represents the known cost vector of a decision-maker, and a set of objective perturbations to \mathbf{c}^0 can be designed and prescribed such that a decision target $\hat{\mathbf{x}}$ becomes an optimal choice for the decision-maker. For example, the elements of \mathbf{c}^0 may represent the cost of producing energy using different resources, in which case the vector $\mathbf{c} - \mathbf{c}^0$ would describe the value of carbon taxes or subsidies that must be imposed on each resource for a low-emission production plan $\hat{\mathbf{x}}$ to become optimal for an energy producer (Zhou et al. 2011, Rathore and Jakhar 2021). In these settings, there may be context-specific constraints on \mathbf{c} , such as non-negativity ($\mathcal{P} = \mathbb{R}_+^n$), integrality ($\mathcal{P} = \mathbb{Z}^n$), variable bounds ($\mathcal{P} = [\mathbf{l}, \mathbf{u}]^n$), or a transformation into a multi-objective problem ($\mathcal{P} = \{\sum_{i=1}^m \alpha_i \mathbf{v}^i, \alpha_i \geq 0, i = 1, \dots, m\}$ where $\mathbf{v}^1, \dots, \mathbf{v}^m$ are given objectives and $\boldsymbol{\alpha}$ is a decision vector of objective weights).

From this point on, we present our inverse optimization model as

$$\begin{aligned} \mathbf{IO}(\mathbf{c}^0, \hat{\mathbf{x}}, \mathcal{X}) : & \underset{\mathbf{c} \in \mathcal{P}}{\text{minimize}} && \|\mathbf{c} - \mathbf{c}^0\|_1 \\ & \text{subject to} && \mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}). \end{aligned} \tag{4}$$

We make this modeling choice because the model is well-motivated and because it is the formulation that is studied in all previous works on inverse mixed integer optimization (Wang 2009, Schaefer 2009, Duan and Wang 2011, Lamperski and Schaefer 2015, Bulut and Ralphs 2015). Nonetheless, we emphasize that this modeling choice comes without loss of generality. The theoretical insights and methodological contributions presented in this paper focus exclusively on difficulties posed by the inverse-feasibility constraint (equation (2)), and are thus directly applicable for any inverse MILO problem in the form of model (3). We also assume without loss of generality that the inverse optimization problem is feasible. If it is not, then the methods proposed in this paper can prove its infeasibility.

In the remainder of this section (Sections 2.2 and 2.3), we focus on the core difficulty of solving the inverse MILO problem: characterizing the inverse-feasible region $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$. In contrast to convex optimization problems, where this condition can be transformed into convex constraints using the Karush-Kuhn-Tucker conditions, determining whether a given $\hat{\mathbf{x}}$ is optimal with respect to some \mathbf{c} , let alone the \mathbf{c} that optimizes the objective function of the inverse problem, requires solving the forward MILO problem, which is NP-hard in general. Furthermore, attempting to use inverse linear

optimization models to approximate inverse MILO models by ignoring the integrality constraints of the forward-feasible region can result in arbitrarily bad approximations. This is further discussed and illustrated in Appendix A.2 of the Online Supplement.

2.2. Preliminary remarks

We first note three basic properties of the inverse-feasible region, then describe it using a set of linear constraints.

REMARK 1. *The inverse-feasible region is nonempty, since $\mathbf{0} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$.*

REMARK 2. *The inverse-feasible region contains only the zero vector, i.e., $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) = \{\mathbf{0}\}$, if and only if \mathcal{X} is full dimensional and $\hat{\mathbf{x}}$ is not on the boundary of $\text{conv}(\mathcal{X})$. If there exists at least one equality constraint or one binary variable used in the description of \mathcal{X} , then $\mathcal{C}(\mathbf{x}, \mathcal{X}) \neq \{\mathbf{0}\}$ for all $\mathbf{x} \in \mathcal{X}$. Similarly, if there exists at least one constraint that is binding at $\hat{\mathbf{x}}$, including variable bounds, then $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) \neq \{\mathbf{0}\}$.*

REMARK 3. *While the forward-feasible region \mathcal{X} does not have to be bounded, the inverse-feasible region $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$ will only contain objective vectors for which the forward problem is bounded. In particular, if \mathbf{c}^{rec} is a nonzero vector in the recession cone of $\text{conv}(\mathcal{X})$, then we have $\mathbf{c}^{rec} \notin \mathcal{C}$. The reason is that $\hat{\mathbf{x}}$ is assumed to be forward-optimal with a finite objective value (under an unobserved objective function) and $\inf\{(\mathbf{c}^{rec})^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X}\} = -\infty$. Thus $\hat{\mathbf{x}} \notin \arg \min\{(\mathbf{c}^{rec})^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X}\}$. In other words, no such \mathbf{c}^{rec} can make $\hat{\mathbf{x}}$ forward-optimal.*

A natural way to explicitly formulate model (4) is to re-write $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$ using a potentially infinite set of linear constraints

$$\mathbf{c}^\top (\hat{\mathbf{x}} - \mathbf{x}) \leq 0, \quad \forall \mathbf{x} \in \mathcal{X}, \quad (5)$$

suggesting that a cutting plane method can be used to solve $\mathbf{IO}(\mathbf{c}^0, \hat{\mathbf{x}}, \mathcal{X})$. Indeed, Wang (2009) proposed a simple cutting plane algorithm to solve the inverse MILO problem in a finite (but exponential, in general) number of iterations by noting that it is sufficient to replace \mathcal{X} in equation (5) with the extreme points of its convex hull.

REMARK 4. *(Wang 2009) Model (4) is equivalent to*

$$\underset{\mathbf{c} \in \mathcal{P}}{\text{minimize}} \quad \|\mathbf{c} - \mathbf{c}^0\|_1 \quad (6a)$$

$$\text{subject to} \quad \mathbf{c}^\top (\hat{\mathbf{x}} - \mathbf{x}^j) \leq 0, \quad \forall \mathbf{x}^j \in \mathcal{E}(\mathcal{X}). \quad (6b)$$

The cutting plane algorithm proposed in Wang (2009) starts without constraints (6b), and iteratively adds cuts of this form through the computation of new extreme points \mathbf{x}^j . The efficiency of this approach depends heavily on the difficulty of generating cuts, which amounts to solving the

forward MILO problem, as well as the number of extreme points of the convex hull of \mathcal{X} (Bulut and Ralphs 2015). With the exception of parallelizing the search for extreme points (Duan and Wang 2011), the algorithm proposed in Wang (2009) remains the state-of-the-art in solving inverse MILO problems.

2.3. Optimality conditions

The main insight from Wang (2009) is that replacing constraints (5) with (6b) preserves the set of optimal solutions to the inverse optimization problem, i.e., $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{E}(\mathcal{X})) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$. In this subsection, we characterize the full family of sets \mathcal{G} such that $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{G}) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$.

DEFINITION 1. Given \mathcal{X} , any set $\mathcal{G} \subseteq \mathbb{R}^n$ for which

$$\mathcal{C}(\hat{\mathbf{x}}, \mathcal{G}) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) \quad (7)$$

is a *generator set* for \mathcal{X} . If $\mathcal{G} \subseteq \mathcal{X}$, then \mathcal{G} is a *forward-feasible generator set*.

Examples of generator sets include $\text{conv}(\mathcal{X})$ and $\mathcal{E}(\mathcal{X})$, with the latter being a forward-feasible generator set.

Next, we demonstrate there may exist a large family of such sets for \mathcal{X} . First, note that constraint set (5) defines a polyhedral cone. Second, each half-space $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) \leq 0$ induced by a feasible point $\mathbf{x} \in \mathcal{X}$ is not uniquely determined by $\hat{\mathbf{x}} - \mathbf{x}$, but rather by any multiple of the vector $\hat{\mathbf{x}} - \mathbf{x}$. In other words, the half-space induced by \mathbf{x} is the same as the one induced by $\hat{\mathbf{x}} + \lambda(\mathbf{x} - \hat{\mathbf{x}})$ for any $\lambda > 0$. The same cut can thus be induced by infinitely many other points of \mathbb{R}^n .

These observations are illustrated in Figure 1. The extreme points $\{\mathbf{x}^1, \dots, \mathbf{x}^7\}$ in Figure 1a induce half spaces for $\mathbf{c} \in \mathbb{R}^n$ and the intersection of these half-spaces is the cone shown in Figure 1b. However, the same cone can be formed using the (blue) diamond points, all of which are *interior* forward-feasible points, plus \mathbf{x}^1 , in Figure 1a.

The last observation above is formalized in the following lemma. Let $\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x}) := \hat{\mathbf{x}} + \lambda(\mathbf{x} - \hat{\mathbf{x}})$, $\lambda > 0$, define a point along the ray from $\hat{\mathbf{x}}$ to \mathbf{x} .

LEMMA 1. Let $\mathcal{E}(\mathcal{X}) = \{\mathbf{x}^1, \dots, \mathbf{x}^N\}$ and $\bar{\mathcal{E}}(\mathcal{X}) = \{\mathbf{y}(\lambda_1, \hat{\mathbf{x}}, \mathbf{x}^1), \dots, \mathbf{y}(\lambda_N, \hat{\mathbf{x}}, \mathbf{x}^N)\}$ for arbitrary positive scalars $\lambda_i > 0, i = 1, \dots, N$. Then,

$$\mathcal{C}(\hat{\mathbf{x}}, \mathcal{E}(\mathcal{X})) = \mathcal{C}(\hat{\mathbf{x}}, \bar{\mathcal{E}}(\mathcal{X})).$$

Lemma 1 shows that the set of extreme points $\mathcal{E}(\mathcal{X})$ is not necessary to characterize $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$ and can instead be replaced with a set of interior and/or exterior points. This observation can be generalized with a necessary and sufficient condition for any collection of points to be a generator set. Let $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ denote the polyhedral cone pointed at $\hat{\mathbf{x}}$ generated by $\{\mathbf{y}(1, \hat{\mathbf{x}}, \mathbf{x})\}_{\mathbf{x} \in \mathcal{X}}$.

THEOREM 1. A set $\mathcal{G} \subseteq \mathbb{R}^n$ is a generator set if and only if $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) = \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$.




Figure 1 (a) The convex hull of a forward-feasible region $\mathcal{X} \subseteq \mathbb{Z}^2$ defined on a two-dimensional lattice is shaded with the points in $\mathcal{E}(\mathcal{X})$ marked. (b) The half-space in the cost vector space defined by each extreme point, and the inverse-feasible region lying at the intersection of half-spaces shaded.

While Lemma 1 shows that generator sets can be constructed without using any extreme points, Theorem 1 goes further and shows that far fewer points than the number of extreme points may be sufficient to construct a generator set. As long as the set of extreme rays of $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ are the same as those of $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$, then \mathcal{G} is a generator set – all other points in \mathcal{G} that do not form an extreme ray are redundant in defining the inverse-feasible region. Referring back to Figure 1a, this result implies that any two points $\mathbf{y}(\lambda_1, \hat{\mathbf{x}}, \mathbf{x}^1)$ and $\mathbf{y}(\lambda_7, \hat{\mathbf{x}}, \mathbf{x}^7)$ with $\lambda_1, \lambda_7 > 0$ are sufficient to define the generator set in that example. In Figure 1b, the inverse-feasible region is exactly the intersection of the half-spaces determined by $\hat{\mathbf{x}} - \mathbf{x}^1$ and $\hat{\mathbf{x}} - \mathbf{x}^7$.

Note that for any forward-feasible point $\mathbf{x} \in \mathcal{X}$, the ray $\{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}_{\lambda \geq 0}$ is in $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$. Thus, $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G}) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ for any collection of forward-feasible points $\mathcal{G} \subseteq \mathcal{X}$. Then, the set $\mathcal{G} \subseteq \mathcal{X}$ is a forward-feasible generator set if $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$, or equivalently if $\mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. This is shown by the following corollary.

COROLLARY 1. *A set $\mathcal{G} \subseteq \mathbb{R}^n$ is a forward-feasible generator set if and only if $\mathcal{G} \subseteq \mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$.*

Analyzing forward-feasible generator sets can be particularly informative for the design of cutting plane algorithms. For example, inverse MILO problems can be solved by the cutting plane algorithm presented in Wang (2009) because $\mathcal{E}(\mathcal{X})$ is a forward-feasible generator set. Below, we present an equivalent characterization of forward-feasible generator sets that does not rely on the notion of rays. This result provides useful intuition to aid in designing a new cutting plane algorithm. Let $\mathcal{B}(\epsilon, \hat{\mathbf{x}})$ denote a closed ball of radius $\epsilon > 0$ around $\hat{\mathbf{x}}$.

THEOREM 2. A set $\mathcal{G} \subseteq \mathbb{R}^n$ is a forward-feasible generator set if and only if $\mathcal{G} \subseteq \mathcal{X}$ and there exists an $\epsilon > 0$ such that

$$\mathcal{B}(\epsilon, \hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X}) \subseteq \text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\}). \quad (8)$$

Theorem 2 states that a set $\mathcal{G} \subseteq \mathcal{X}$ is a forward-feasible generator set if and only if the convex hull of the set of points in \mathcal{G} and $\{\hat{\mathbf{x}}\}$ contains an epsilon ball around $\hat{\mathbf{x}}$ intersected with the convex hull of \mathcal{X} . Going back to Figure 1a, $\mathcal{G} = \{\mathbf{x}^1, \mathbf{x}^7\}$ is a forward-feasible generator set since the convex hull of $\{\mathbf{x}^1, \mathbf{x}^7, \hat{\mathbf{x}}\}$ contains the part of an epsilon ball around $\hat{\mathbf{x}}$ inside $\text{conv}(\mathcal{X})$, which is shaded in red. Considering the ray $\{\mathbf{y}(\lambda_1, \hat{\mathbf{x}}, \mathbf{x}^1)\}_{\lambda_1 \geq 0}$, only \mathbf{x}^1 is a forward-feasible point, which means that any forward-feasible generator set must include \mathbf{x}^1 . However, there are multiple forward-feasible points along the ray $\{\mathbf{y}(\lambda_7, \hat{\mathbf{x}}, \mathbf{x}^7)\}_{\lambda_7 \geq 0}$, which means that any of those feasible points, namely (5,10), (6,10) or (7,10), can be part of the forward-feasible generator set.

The two key takeaways from Theorem 2 are that a forward-feasible generator set can: (i) consist of many fewer than $|\mathcal{E}(\mathcal{X})|$ points, and (ii) be generated from non-extreme points that lie much closer to $\hat{\mathbf{x}}$.

3. Cutting Plane Algorithm with Trust Regions

The theoretical insights from the previous section indicate that a new approach to generating cuts may yield a more effective solution algorithm. Instead of generating cuts by identifying extreme points of the convex hull of the forward-feasible region \mathcal{X} , we generate cuts by identifying interior points of \mathcal{X} , which we accomplish using *trust regions*. More specifically, we intersect \mathcal{X} with a trust region, and generate cuts at interior points of \mathcal{X} obtained as extreme points of the convex hull of this restriction of \mathcal{X} .

Let $\mathcal{T}(\hat{\mathbf{x}}, p)$ be a trust region around the point $\hat{\mathbf{x}}$ of size $p \geq 1$, defined as

$$\mathcal{T}(\hat{\mathbf{x}}, p) := \{\mathbf{y} \in \mathbb{R}^n \mid \|\hat{\mathbf{x}} - \mathbf{y}\|_1 \leq p\}. \quad (9)$$

We simplify notation to \mathcal{T} when the parameter p is not the focus of discussion.

Guided by Theorem 2, we define trust regions to be *centered at $\hat{\mathbf{x}}$* . The trust region, when imposed onto \mathcal{X} , represents an attempt to capture a forward-feasible generator set. The 1-norm in equation (9) is chosen because the trust region has a polynomial number of extreme points. When intersected with \mathcal{X} , the resulting restricted forward-feasible region $\mathcal{T} \cap \mathcal{X}$ itself is more likely to have a smaller number of extreme points compared to \mathcal{X} . This can reduce the number of forward-feasible points that need to be identified in order to build a generator set. Furthermore, the extreme points of this smaller, restricted region may also be computationally easier to identify. These considerations will be discussed in detail below.

3.1. Model decomposition and cutting plane framework

The inverse MILO problem is decomposed into a master problem, describing a relaxation of the inverse model (4), and a cut generation subroutine that is used to iteratively tighten the master problem. The framework of our cutting plane algorithm is given in Algorithm 1. A global information set \mathcal{J} is used to pass information between successive iterations of the subroutine, which specifically includes the trust region and outer-loop index for our purposes. It is initialized with a trust region of size p^0 centered at $\hat{\mathbf{x}}$, and the iteration index $i = 0$.

Algorithm 1 A general cutting plane algorithm for inverse MILO

Input: An inverse MILO problem instance $(\mathbf{c}^0, \hat{\mathbf{x}}, \mathcal{X})$, initial trust region size p^0

Output: An inverse-optimal solution \mathbf{c}^*

- 1: Initialize $i = 0$, $\mathcal{J}^0 = (\mathcal{T}(\hat{\mathbf{x}}, p^0), i)$, $\tilde{\mathcal{X}}^i = \emptyset$, $\tilde{\mathbf{c}}^i = \mathbf{c}^0$
 - 2: Run **SUBROUTINE** $(\tilde{\mathbf{c}}^i, \hat{\mathbf{x}}, \mathcal{X}, \mathcal{J}^i)$, let $\tilde{\mathbf{x}}^i$ and \mathcal{J}^{i+1} be its output
 - 3: **while** $\tilde{\mathbf{x}}^i \neq \hat{\mathbf{x}}$ **do**
 - 4: $i \leftarrow i + 1$
 - 5: $\tilde{\mathcal{X}}^i \leftarrow \tilde{\mathcal{X}}^{i-1} \cup \{\tilde{\mathbf{x}}^{i-1}\}$
 - 6: Solve **MP** $(\hat{\mathbf{x}}, \tilde{\mathcal{X}}^i)$, let $\tilde{\mathbf{c}}^i$ be its optimal solution
 - 7: Run **SUBROUTINE** $(\tilde{\mathbf{c}}^i, \hat{\mathbf{x}}, \mathcal{X}, \mathcal{J}^i)$, let $\tilde{\mathbf{x}}^i$ and \mathcal{J}^{i+1} be its output
 - 8: **return** $\mathbf{c}^* = \tilde{\mathbf{c}}^i$
-

The master problem is defined below, where $\tilde{\mathcal{X}} \subseteq \mathcal{X}$ denotes a finite set of forward-feasible points.

$$\mathbf{MP}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}): \underset{\mathbf{c} \in \mathcal{P}}{\text{minimize}} \quad \|\mathbf{c} - \mathbf{c}^0\|_1 \tag{10a}$$

$$\text{subject to} \quad \mathbf{c}^\top (\hat{\mathbf{x}} - \mathbf{x}) \leq 0, \quad \forall \mathbf{x} \in \tilde{\mathcal{X}}, \tag{10b}$$

$$\mathbf{y}^\top \mathbf{A} = \mathbf{c}^\top, \tag{10c}$$

$$\mathbf{y} \geq \mathbf{0}. \tag{10d}$$

Constraints (10b) ensure that $\hat{\mathbf{x}}$ is optimal with respect to \mathbf{c} over $\tilde{\mathcal{X}}$. As mentioned in Remark 3, the inverse-feasible region $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$ only contains cost vectors for which the forward optimization problem is bounded. However, due to its relaxed nature, the master problem can propose cost vectors that lead to unbounded forward problems in the subroutine. This can be prevented by constraints (10c) and (10d), as proposed in Wang (2009), which are not needed if the forward-feasible region is bounded. Let $\tilde{\mathbf{c}}$ denote an optimal solution to the master problem **MP** $(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$.

3.2. Cut generation subroutine

Given $\tilde{\mathbf{c}}$, the cut generation subroutine either finds a feasible point $\tilde{\mathbf{x}} \in \mathcal{X}$ satisfying $\tilde{\mathbf{c}}^\top(\hat{\mathbf{x}} - \tilde{\mathbf{x}}) > 0$ (i.e., $\tilde{\mathbf{c}} \notin \mathcal{C}(\hat{\mathbf{x}}, \tilde{\mathcal{X}} \cup \{\tilde{\mathbf{x}}\})$), which generates a cut, or returns $\hat{\mathbf{x}}$, verifying that $\tilde{\mathbf{c}} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$. Note that the verification in the latter case requires solving the forward problem $\mathbf{FP}(\tilde{\mathbf{c}}, \mathcal{X})$, which means the cut generation subroutine is at least as hard.

The cut generation subroutine is presented in Subroutine 1. It consists of four key components: the cut generation problem (Step 3) and three functions to modify the trust region (Steps 2, 5 and 8), which we describe next.

Subroutine 1 Cut generation subroutine

Input: A candidate objective $\tilde{\mathbf{c}}$, forward-feasible point $\hat{\mathbf{x}}$, forward-feasible region \mathcal{X} , information set $\mathcal{J}^{\text{in}} = \{\mathcal{T}, i\}$

Output: A forward-feasible point $\tilde{\mathbf{x}}$ and updated information set \mathcal{J}^{out}

- 1: Initialization: $k = 1, \mathcal{T}^0 = \mathcal{T}$.
 - 2: $\mathcal{T}^k \leftarrow \text{REMOVE}(\mathcal{T}^{k-1})$
 - 3: Solve $\mathbf{FP}(\tilde{\mathbf{c}}, \mathcal{T}^k \cap \mathcal{X})$, let $\tilde{\mathbf{x}}$ be the solution
 - 4: **if** $\tilde{\mathbf{c}}^\top(\hat{\mathbf{x}} - \tilde{\mathbf{x}}) > 0$ **then**
 - 5: $\mathcal{T}^* \leftarrow \text{SAVE}(\mathcal{T}^k)$,
 - 6: **Return** $\tilde{\mathbf{x}}, \mathcal{J}^{\text{out}} = \{\mathcal{T}^*, i + 1\}$
 - 7: **else if** $\mathcal{T}^k \neq \mathbb{R}^n$ **then**
 - 8: $\mathcal{T}^{k+1} \leftarrow \text{UPDATE}(\mathcal{T}^k)$
 - 9: $k \leftarrow k + 1$
 - 10: Return to Step 2
 - 11: **else**
 - 12: **Return** $\hat{\mathbf{x}}, \mathcal{J}^{\text{in}}$
-

3.2.1. Cut generation problem (Step 3). Given a candidate objective $\tilde{\mathbf{c}}$ and a trust region \mathcal{T} , the cut generation problem is the forward problem solved over $\mathcal{T} \cap \mathcal{X}$ (i.e., $\mathbf{FP}(\tilde{\mathbf{c}}, \mathcal{T} \cap \mathcal{X})$), which we refer to as a *subregion* of \mathcal{X} . We note that the choice of 1-norm in the definition of employed trust regions ensures that the cut generation problem remains linear.

If the optimal solution of the cut generation problem $\tilde{\mathbf{x}}$ satisfies the condition $\tilde{\mathbf{c}}^\top(\hat{\mathbf{x}} - \tilde{\mathbf{x}}) > 0$, then $\tilde{\mathbf{x}}$ is returned; we call such an $\tilde{\mathbf{x}}$ a *violated forward-feasible point*. If no cut is found, then there are two possibilities depending on whether a trust region was used. If no trust region was used, i.e., $\mathcal{T}^k = \mathbb{R}^n$, then $\hat{\mathbf{x}}$ is returned and Algorithm 1 terminates with $\tilde{\mathbf{c}}$. Otherwise, $\tilde{\mathbf{c}}$ renders

$\hat{\mathbf{x}}$ optimal over a subregion of \mathcal{X} , but not necessarily over \mathcal{X} . Thus, larger subregions need to be considered before it can be verified that no additional violated cuts exist, which leads to a trust region update.

3.2.2. UPDATE function (Step 8). When a subregion $\mathcal{T}(\hat{\mathbf{x}}, p) \cap \mathcal{X}$ has been exhausted of violated cuts, we increase the size p of the trust region to δp for some $\delta > 1$:

$$\text{UPDATE}(\mathcal{T}(\hat{\mathbf{x}}, p)) = \mathcal{T}(\hat{\mathbf{x}}, \delta p). \quad (11)$$

For example, setting $\delta = 2$ doubles the size of the trust region each time UPDATE is called. There is a natural trade-off to be considered when choosing δ . If δ is too small, then the trust regions grow slowly and may lead to the generation of subregions with no violated cuts. This will result in time-consuming, “empty” calls to the cut generation problem, i.e., calls in which the cut generation problem is solved but is unable to generate any violated forward-feasible point. On the other hand, if δ is too large, the computational advantages of solving the cut generation problem over a smaller subregion with fewer extreme points may quickly be lost.

3.2.3. REMOVE function (Step 2). The REMOVE function periodically removes the existing trust region (by setting $\mathcal{T} = \mathbb{R}^n$) at regular intervals in both outer (i.e., master) iterations i and inner (i.e., subroutine) iterations k . In particular, for fixed index values $i^* > 0$ and $k^* > 0$,

$$\text{REMOVE}(\mathcal{T}) = \begin{cases} \mathbb{R}^n, & \text{if } i \in \{i^*, 2i^*, \dots\} \text{ or } k = k^*, \\ \mathcal{T}, & \text{otherwise.} \end{cases} \quad (12)$$

In other words, the cut generation problem is solved over \mathcal{X} every i^* iterations in the outer loop and when the inner loop reaches the iteration limit of k^* . The REMOVE function guarantees finite convergence of Algorithm 1, irrespective of how the UPDATE function is defined, because an extreme point of \mathcal{X} is returned at least every i^* iterations in a finite number of (at most k^*) attempts to verify that a candidate objective $\tilde{\mathbf{c}}^i$ is inverse-feasible when no more violated cuts exist.

Note that the classical cutting plane algorithm proposed in Wang (2009) is a special case of our trust region-based approach where $i^* = k^* = 1$.

3.2.4. SAVE function (Step 5). The SAVE function is called when a forward-feasible point yielding a violated cut is found by the cut generation problem. If such a point is found during an iteration where the trust region is \mathbb{R}^n (i.e., when REMOVE is called), the SAVE function saves the previous trust region. On the other hand, if the REMOVE function did not remove the trust region during the current iteration, then the current trust region is saved. This ensures that there always exists a “non-trivial” trust region ($\mathcal{T}(\hat{\mathbf{x}}, p) < \infty$) saved.

Mathematically, the SAVE function is defined as:

$$\text{SAVE}(\mathcal{T}^k) = \begin{cases} \mathcal{T}^{k-1} & \text{if } \mathcal{T}^k = \mathbb{R}^n, \\ \mathcal{T}^k & \text{otherwise.} \end{cases}$$

3.2.5. Discussion and example. Trust regions can reduce the difficulty of solving the cut generation problem by reducing the size of the forward-feasible region. However, they can also lead to the generation of *stronger* cuts.

We use Figure 2a to illustrate this latter point. A trust region of size one ($p = 1$) is imposed on a two-dimensional integer forward-feasible region, creating a subregion containing two extreme points (indicated using blue diamonds), excluding $\hat{\mathbf{x}}$. Those two extreme points of the subregion have the same cut generation potential as the five extreme points of \mathcal{X} labeled in gray. More specifically, the set of cuts generated by the former set of extreme points can dominate those generated by the latter. Once this trust region is exhausted, it must be updated. In Figure 2a,




Figure 2 Two examples of when a trust region is applied on a two-dimensional integer feasible region. The convex hull of \mathcal{X} and $\mathcal{X} \cap \mathcal{T}$ is shaded in light gray and dark gray, respectively. The shaded cone at the right bottom corner of each figure represents $-\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X} \cap \mathcal{T})$ while the darker shaded cone represents $-\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$, a subset of $-\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X} \cap \mathcal{T})$.

increasing the size of the trust region by a single unit would cover $\tilde{\mathbf{x}}$, the only additional point needed to build a generator set. The trust region approach guarantees that at most three cuts need to be generated, or equivalently three forward-feasible points need to be identified, to solve the inverse MILO problem. On the other hand, the cutting plane algorithm without trust regions, which computes extreme points of \mathcal{X} , may require generating many more cuts.

The REMOVE function periodically removes trust regions and provides opportunities for the algorithm to compute points outside of the existing trust region. We consider another example in Figure 2b with a trust region of $\mathcal{T}(\hat{\mathbf{x}}, 2)$ over a different two-dimensional integer forward-feasible region. Assuming the subregion $\mathcal{X} \cap \mathcal{T}(\hat{\mathbf{x}}, 2)$ has been exhausted of cuts, removal of this trust region would lead to the computation of $\tilde{\mathbf{x}}$ with certainty, as it is the only extreme point of $\text{conv}(\mathcal{X})$ that can still generate a valid cut. This would terminate the algorithm at the next step. On the other hand, if the size of the trust region was increased incrementally each time without removing the trust region, many unnecessary interior points would be computed before a sufficiently large

trust region captures $\tilde{\mathbf{x}}$. Periodic removal of trust regions thus helps to reduce the potential of being overly-conservative, and when used in conjunction with trust regions create an effective cut generation subroutine.

4. Algorithm Enhancements

In this section, we propose two enhancements to Subroutine 1 that can be implemented independently or together. These two enhancements are introduced to improve cut generation speed, particularly during iterations when the subregion is large.

4.1. Early-stop heuristic for the cut generation problem

As defined in Subroutine 1, the cut generation problem is solved to optimality in each iteration. However, this can be time consuming, particularly during iterations where the trust region is large or removed entirely. Here, we introduce a solution time threshold τ . If the cut generation problem is not solved by τ , but at least one violated forward-feasible point has been found, the feasible point of maximum violation is returned. If no violated forward-feasible point has been found by τ , the cut generation problem will return the first such point found after τ . This enhancement can reduce the cut generation time and can be effective for two reasons: (i) Solving the cut generation problem may not necessarily provide the best cuts anyways, as discussed in Section 2, and (ii) the solver may have already found an optimal violated forward-feasible point in the cut generation problem but has not computed a bound strong enough to prove its optimality. When a standard MILO solver, such as CPLEX and Gurobi, is used to solve the forward problem in Subroutine 1, the early-stop heuristic can be implemented using a callback function.

4.2. Stochastic dimensionality reduction of the trust region

This enhancement is motivated by the observation that a violated forward-feasible point generated within a trust region may have many components with the same value as the corresponding components in $\hat{\mathbf{x}}$. For example, if \mathcal{X} contains only integer variables, then all forward-feasible points within a trust region $\mathcal{T}(\hat{\mathbf{x}}, p)$ around $\hat{\mathbf{x}}$ will differ from $\hat{\mathbf{x}}$ by at most p components. By explicitly restricting *which* subset of components can differ from the values of $\hat{\mathbf{x}}$, we can decrease the dimensionality and hence the size of the cut generation problem.

Let $\mathcal{S} \subset \mathcal{N} := \{1, \dots, n\}$ and

$$\mathcal{T}_{\mathcal{S}}(\hat{\mathbf{x}}, p) := \{\mathbf{y} \in \mathbb{R}^n \mid \mathbf{y} \in \mathcal{T}(\hat{\mathbf{x}}, p), y_j = \hat{x}_j \forall j \notin \mathcal{S}\} \quad (13)$$

define a lower dimensional trust region containing points that can differ from $\hat{\mathbf{x}}$ only in the indices contained in \mathcal{S} .

We extend the previous cut generation subroutine to generate lower dimensional trust regions by randomly selecting $\mathcal{S} \subset \mathcal{N}$, as shown in Subroutine 2. At a high level, the main difference between

Subroutine 2 Cut generation subroutine with dimensionality reduction

Input: A candidate objective $\tilde{\mathbf{c}}$, forward-feasible point $\hat{\mathbf{x}}$, collection of points $\tilde{\mathcal{X}}$, information set $\mathcal{J}^{\text{in}} = \{\mathcal{T}, i, h\}$

Output: A forward-feasible point $\tilde{\mathbf{x}}$ and updated information set \mathcal{J}^{out}

- 1: Initialization: $k = 1, \mathcal{T}^0 = \mathcal{T}$.
 - 2: $\mathcal{T}^k \leftarrow \text{REMOVE}(\mathcal{T}^{k-1})$
 - 3: Solve **FP**($\tilde{\mathbf{c}}, \mathcal{T}^k \cap \mathcal{X}$), let $\tilde{\mathbf{x}}$ be the solution
 - 4: **if** $\tilde{\mathbf{c}}^\top(\hat{\mathbf{x}} - \tilde{\mathbf{x}}) > 0$ **then**
 - 5: $\mathcal{T}^* \leftarrow \text{SAVE}(\mathcal{T}^k)$
 - 6: **if** $\mathcal{T}^k \neq \mathbb{R}^n$ **then**
 - 7: $h \leftarrow 0$
 - 8: **Return** $\tilde{\mathbf{x}}, \mathcal{J}^{\text{out}} = \{\mathcal{T}^*, i + 1, h\}$
 - 9: **else if** $\mathcal{T}^k \neq \mathbb{R}^n$ **then**
 - 10: $h \leftarrow h + 1$
 - 11: $\mathcal{T}^{k+1} \leftarrow \text{S-UPDATE}(\mathcal{T}^k)$
 - 12: **if** $h = h^* + 1$ **then**
 - 13: $h \leftarrow 0$
 - 14: $k \leftarrow k + 1$
 - 15: Return to Step 2
 - 16: **else**
 - 17: **Return** $\hat{\mathbf{x}}, \mathcal{J}^{\text{in}}$
-

Subroutine 1 and Subroutine 2 is that when the trust region size (p) is increased, Subroutine 2 first computes a series of low dimensional trust regions of the increased size, exhausts the corresponding subregions of violated forward-feasible points, and then considers the full dimensional trust region. More specifically, Subroutine 2 considers the full dimensional trust region $\mathcal{T}(\hat{\mathbf{x}}, p)$ only after h^* low dimensional trust regions of size p have *consecutively* failed to generate a cut. The index h tracks the number of consecutive “empty” low dimensional trust regions (step 10), and is reset to zero when a cut is successfully generated within a trust region (step 7).

The trust region updating function is modified accordingly to reflect this process, and is slightly more involved than the original UPDATE function in equation (11). To generate a low dimensional trust region, we randomly sample a set $\mathcal{S} \subset \mathcal{N}$. Let s denote a value less than $|\mathcal{N}|$, and let $\text{RAND}(s)$ denote a function that generates a random subset of \mathcal{N} of size s . The value s , which represents the cardinality of the set \mathcal{S} , can be either a fixed, predetermined value or a function of the trust region

size, i.e., $s(p)$. For now, we assume the latter, and will discuss further in the following paragraph. The new stochastic (indicated by S-) update function becomes:

$$\text{S-UPDATE}(\mathcal{T}_{\mathcal{S}}(\hat{\mathbf{x}}, p)) = \begin{cases} \mathcal{T}(\hat{\mathbf{x}}, p), & \text{if } h = h^*, \\ \mathcal{T}_{\text{RAND}(s(p))}(\hat{\mathbf{x}}, p), & \text{if } h < h^*, \\ \mathcal{T}_{\text{RAND}(s(\delta p))}(\hat{\mathbf{x}}, \delta p), & \text{if } h = h^* + 1. \end{cases} \quad (14)$$

where δ , the trust region growth rate, is the same as in the original UPDATE function. When $h < h^*$, a new low dimensional trust region is generated. When $h = h^*$, the full dimensional trust region is generated instead. Finally, when the full dimensional trust region fails to identify any violated forward-feasible point, i.e., $h = h^* + 1$, the S-UPDATE function generates a new low dimensional trust region of an increased size.

We conclude by discussing the choice of the function $s(p)$ in equation (14), which defines the cardinality of the set \mathcal{S} to be sampled given that the trust region is of size p . While $s(p)$ can be defined simply to return a fixed value (less than $|\mathcal{N}|$) for any input p , we define $s(p)$ to help offset the reduction in computational efficiency associated with larger trust regions. In particular, we define $s(p)$ to be a decreasing function of the trust region size p , i.e., we increasingly reduce the dimensionality of the trust region as its size grows. Our specific choice of the function $s(p)$ is:

$$s(p) = \max\{\lfloor (1 - \kappa(p - 1))n \rfloor, \lfloor qn \rfloor\}. \quad (15)$$

In equation (15), the parameter $\kappa \in (0, 1)$ denotes the dimensionality reduction rate: when the size of the trust region increases by 1 unit, the dimension of the trust region decreases by κ . The parameter $q \in (0, 1)$ specifies an upper bound on the dimensionality reduction, so that the dimension of the trust region will be no less than qn .

5. Extension to Multi-Point Inverse MILO

In this section, we present an extension of our base inverse MILO model (4). While the literature surrounding inverse MILO problems has almost exclusively focused on model (4), the literature on inverse convex optimization problems has expanded in the direction of data-driven parameter estimation, where multiple solutions can be used as input. In light of this, we propose a *multi-point* inverse MILO model, which takes multiple input solutions from potentially different forward-feasible regions to generate a cost vector. We show how our trust region-based cutting plane algorithms can be directly extended to solve this model.

5.1. Problem description

Let \mathcal{D} be the index set of D data points, $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_D$, which are feasible for their respective feasible regions $\mathcal{X}_1, \dots, \mathcal{X}_D$. The following model is a natural extension of model (4) to the multi-point case:

$$\underset{\mathbf{c} \in \mathcal{P}}{\text{minimize}} \quad \|\mathbf{c} - \mathbf{c}^0\|_1 \quad (16a)$$

$$\text{subject to} \quad \mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}_d, \mathcal{X}_d), \quad \forall d \in \mathcal{D}. \quad (16b)$$

The data-driven inverse optimization literature has also considered “noisy” data cases for continuous forward problems, where it is not possible to compute a single vector \mathbf{c} that is inverse-feasible for each and every decision input. In the light of this, we also propose the following extension of model (16), where we introduce an inverse-feasible cost vector $\bar{\mathbf{c}}_d$ for each data point $\hat{\mathbf{x}}_d \in \mathcal{X}_d$:

$$\underset{\mathbf{c} \in \mathcal{P}, \{\bar{\mathbf{c}}_d\}_{d \in \mathcal{D}}}{\text{minimize}} \quad \|\mathbf{c} - \mathbf{c}^0\|_1 + \lambda \sum_{d \in \mathcal{D}} \|\bar{\mathbf{c}}_d - \mathbf{c}\|_1 \quad (17a)$$

$$\text{subject to} \quad \bar{\mathbf{c}}_d \in \mathcal{C}(\hat{\mathbf{x}}_d, \mathcal{X}_d), \quad \forall d \in \mathcal{D}. \quad (17b)$$

The regularization term λ represents a trade-off between making the cost vector \mathbf{c} close to the prior \mathbf{c}^0 and to the set of inverse-feasible cost vectors $\{\bar{\mathbf{c}}_d\}_{d \in \mathcal{D}}$.

5.2. Extended multi-point cutting plane algorithm

We now show how a natural extension of Algorithm 1 can be used to solve models (16) and (17). The key observation is that each feasible region \mathcal{X}_d in (16b) and (17b) can be replaced with a forward-feasible generator set $\mathcal{G}_d \subseteq \mathcal{X}_d$. So, the previous trust region ideas are applicable here. The cutting plane algorithm for solving model (17) is provided in Algorithm 2. The algorithm can be easily simplified to solve model (16) by ignoring all appearances of the set $\{\bar{\mathbf{c}}_d\}_{d \in \mathcal{D}}$ and replacing all appearances of terms $\bar{\mathbf{c}}_d$ with \mathbf{c} .

The master problem defines a relaxation of model (17), where $\tilde{\mathcal{X}}_d \subseteq \mathcal{X}_d$ is a finite set for each d :

$$\mathbf{MP}(\mathbf{c}^0, \{\hat{\mathbf{x}}_d, \mathcal{X}_d\}_{d \in \mathcal{D}}) : \underset{\mathbf{c} \in \mathcal{P}, \{\bar{\mathbf{c}}_d\}_{d \in \mathcal{D}}}{\text{minimize}} \quad \sum_{d \in \mathcal{D}} \|\mathbf{c} - \mathbf{c}^0\|_1 + \lambda \sum_{d \in \mathcal{D}} \|\bar{\mathbf{c}}_d - \mathbf{c}\|_1 \quad (18a)$$

$$\text{subject to} \quad \bar{\mathbf{c}}_d^\top (\hat{\mathbf{x}}_d - \mathbf{x}_d) \leq 0, \quad \forall \mathbf{x}_d \in \tilde{\mathcal{X}}_d, \quad d \in \mathcal{D}. \quad (18b)$$

With multiple data points, there is a cut generation problem associated with each $d \in \mathcal{D}$. Thus, there is a new issue that must be considered, specifically that of a breadth-first versus depth-first search: whether the master problem is called after cuts for each d are generated versus after the first cut for a single d . Let $V^* \in \{1, \dots, D\}$ be a parameter that denotes the minimum number of violated forward-feasible points, if any, that must be found until the master problem is called. The larger V^* is, the more the algorithm behaves “breadth-first”. In the extreme case where $V^* = D$, the master problem is called only when the master problem candidate solution is considered for every d . This may reduce the number of master problem calls, but at the expense of potentially solving many unnecessary cut generation problems that cannot identify violated cuts. For example, this may happen when a generator set has been identified for several forward-feasible regions already but not all of them. In the other extreme where $V^* = 1$, the master problem is called as soon as a violated forward-feasible point is found for any $d \in \mathcal{D}$. In the algorithm, count_v (count_e) denotes the number of forward-feasible regions for which a violated point has (has not) been found. The algorithm terminates when $\text{count}_e = D$, i.e., there does not exist a single violated forward-feasible point in any forward-feasible region.

Algorithm 2 Cutting plane algorithm for multi-point inverse MILO**Input:** An inverse MILO problem instance $(\{\hat{\mathbf{x}}_d, \mathcal{X}_d\}_{d \in \mathcal{D}})$, initial trust region size p^0 **Output:** An inverse-optimal solution $(\mathbf{c}^*, \{\bar{\mathbf{c}}_d^*\}_{d \in \mathcal{D}})$

```

1: Initialize  $i = 0, \mathcal{J}_d^0 = (\mathcal{T}(\hat{\mathbf{x}}_k, p^0), i), \tilde{\mathcal{X}}_d^i = \emptyset, \text{count}_e = 0, \text{count}_v = 0.$ 
2: while  $\text{count}_e < |\mathcal{D}|$  do
3:      $\text{count}_e \leftarrow 0$ 
4:      $\text{count}_v \leftarrow 0$ 
5:     Solve MP $(\{\hat{\mathbf{x}}_d, \tilde{\mathcal{X}}_d^i\}_{d \in \mathcal{D}})$ , let  $\tilde{\mathbf{c}}^i, \{\tilde{\mathbf{c}}_d^i\}_{d \in \mathcal{D}}$  be its optimal solution
6:     while  $(\text{count}_v < V^*)$  and  $(\text{count}_e + \text{count}_v < |\mathcal{D}|)$  do
7:         for  $d = 1, \dots, D$  do
8:             Solve SUBROUTINE $(\tilde{\mathbf{c}}_d^i, \hat{\mathbf{x}}_d, \mathcal{X}_d, \mathcal{J}_d^i)$ , let  $\tilde{\mathbf{x}}_d^{i+1}$  and  $\mathcal{J}_d^{i+1}$  be its output
9:             if  $\tilde{\mathbf{x}}_d^{i+1} = \hat{\mathbf{x}}_d$  then
10:                  $\text{count}_e \leftarrow \text{count}_e + 1$ 
11:                  $\tilde{\mathcal{X}}_d^{i+1} \leftarrow \tilde{\mathcal{X}}_d^i$ 
12:             else
13:                  $\text{count}_v \leftarrow \text{count}_v + 1$ 
14:                  $\tilde{\mathcal{X}}_d^{i+1} \leftarrow \tilde{\mathcal{X}}_d^i \cup \{\tilde{\mathbf{x}}_d^{i+1}\}$ 
15:          $i \leftarrow i + 1$ 
16: return  $\mathbf{c}^* = \tilde{\mathbf{c}}^i, \bar{\mathbf{c}}_d^* = \tilde{\mathbf{c}}_d^i \forall d \in \mathcal{D}$ 

```

6. Experimental Design

In this section, we discuss the setup of our computational experiments, which include the generation of test instances, the different subroutine variants and enhancements considered, and the parameterization of each subroutine. Our computational experiments are focused on inverse MILO problems of the form presented in model (4) with $\mathcal{P} = \mathbb{R}^n$. We made this modeling choice to focus the presentation and discussion on the key insights gained from applying our new cutting plane algorithms over various problem structures and forward-feasible regions. This model is also studied in all previous computational studies of inverse MILO (Wang 2009, Schaefer 2009, Duan and Wang 2011). To generate comprehensive insights into the performance of our algorithms, we apply them to a diverse set of problem instances, much larger than any set considered in the literature to date.

6.1. Test bank generation

To generate a comprehensive set of inverse MILO problems over many different problem structures, we draw from the MIPLIB 2017 benchmark library (Gleixner et al. 2019), which includes 240 MIPLIB problems. We use a subset of these problems to generate a test bank of inverse MILO

problems as follows. We first consider all MIPLIB problems with fewer than 12,000 variables and constraints. This reduces our bank to 125 total MIPLIB problems. For each problem, we attempt to generate a feasible solution by solving the problem with a randomly generated cost vector. If the problem is unbounded or cannot be solved within 10 minutes using Gurobi 8.1.0, we try again with a different randomly generated cost vector. If this procedure cannot generate three feasible solutions within ten attempts, the MIPLIB problem is dropped from consideration. Otherwise, three feasible points are generated for each MIPLIB problem. This approach results in 73 MIPLIB problems being included in our test bank, each with three feasible points, totalling 219 inverse MILO problem instances. Each instance is constructed by using one of the feasible solutions as $\hat{\mathbf{x}}$, the corresponding original cost vector of the MIPLIB problem (i.e., not the random one used to generate $\hat{\mathbf{x}}$) as \mathbf{c}^0 and the corresponding set of MIPLIB problem constraints as \mathcal{X} . The 3 instances constructed per MIPLIB problem are labeled using the MIPLIB name and a suffix of t1, t2, and t3 (see Appendix A.4 in the Online Supplement).

Of the 73 included MIPLIB problems, 9 have only binary variables, 42 have binary and continuous variables, 14 have binary and integer variables, and 8 have binary, integer and continuous variables. The number of variables in these problems ranges from 34 to 11,700 with an average of 3,185. The number of constraints ranges from 4 to 10,900 with an average of 2,315. These problems represent applications such as scheduling, production, shipment, assignment, set covering and bin packing.

6.2. Subroutine variants

To solve these inverse MILO instances, we implement and compare five different cutting plane algorithms, described below:

- **CP.** This is the classical cutting plane approach proposed in Wang (2009), which serves as the benchmark algorithm.
- **CP-ES.** This is a simple extension of the CP algorithm in which we embed the early-stop heuristic presented in Section 4.1. We choose a time threshold of 5 seconds, at which time the most violated forward-feasible solution is returned ($\tau = 5$). If no violated points have been found by 5 seconds, the subroutine returns the first one found after the 5-second threshold has been exceeded.
- **CPTR.** This is the cutting plane algorithm with trust regions, defined in Subroutine 1. The subroutine initializes with a trust region of size one ($p^0 = 1$ in equation (9)). The trust region doubles in size each time the trust region becomes redundant ($\delta = 2$ in equation (11)). The trust region is removed every 10th cut that is generated and when no violated forward-feasible points can be found within the trust region ($i^* = 10$ and $k^* = 2$ in equation (12)).
- **CPTR-ES.** This is an extension of the CPTR algorithm in which we embed the early-stop heuristic with a threshold of 5 seconds ($\tau = 5$), similar to the CP-ES algorithm.

- **CPTR-ES-DR.** This considers the CPTR algorithm with both the early-stop heuristic and the stochastic dimensionality reduction enhancement. We implement a dimensionality reduction rate of 3% ($\kappa = 0.03$ in equation (15)). The dimension of the trust region will never be lower than 80% of n ($q = 0.8$). Finally, we revert to the full dimensional trust region when ten consecutive low-dimensional trust regions are found to have no violated forward-feasible points ($h^* = 10$).

All experiments were coded in Python 3.7 and optimization problems are solved using Gurobi 8.1.0 parameterized with a single thread. The experiments were conducted on a Intel Core i7-4790 processor at 3.60 GHz on a Windows 8.1 Pro. A time constraint of one hour was set for all instances.

7. Numerical Results

In this section, we present the computational results using the algorithms outlined in Section 6. We begin with a high-level overview of the results, summarized by Figures 3 and 4. We then provide an in-depth study of the advantages and potential trade-offs that exist when using trust regions and the additional enhancements (Section 7.1). In particular, we compare cut generation speed and cut strength across the various algorithms. We also conduct a sensitivity analysis, examining how the results change when several baseline parameters are modified (Section 7.2). We conclude with a brief summary of the main takeaways from our numerical results (Section 7.3). Results for each and every instance can be found in Appendix A.4 in the Online Supplement.




Figure 3 Performance profiles of different cutting-plane algorithms; 66, 99, 107, 134 and 136 instances, out of 215 total instances, were solved by the CP, CP-ES, CPTR, CPTR-ES and CPTR-ES-DR algorithms respectively over a one-hour time window.

From the performance profiles in Figure 3, we observe that CPTR-ES and CPTR-ES-DR solve significantly more instances than CPTR or CP-ES, which in turn solve many more than CP, the baseline algorithm. The two best algorithms are also significantly faster than the others. For instance, they can each solve the same number of instances as CP (66 instances) and CP-ES (107 instances) in only 2.5% and 12% of the required time, respectively. The main takeaway from these high-level results is that either trust regions or the early-stop heuristic can lead to dramatic improvements in the solution time, compared to CP. When trust regions and early-stopping are used together, the improvements are even more significant.

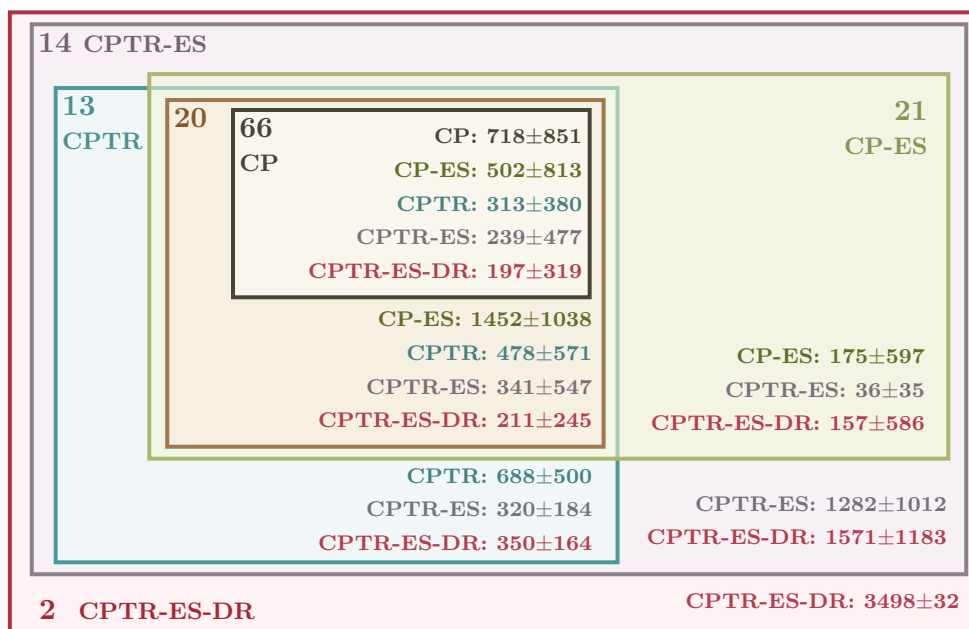


Figure 4 Summary of computational results over solved instances. Single numbers denote the number of solved instances in the difference sets. Solution times in seconds are given in the form of average \pm standard deviation.

Figure 4 presents a more granular breakdown of the solved instances. Notably, the results exhibit a distinct *nested* structure in which each additional enhancement can solve strictly more instances than without the enhancement. For example, all 66 instances solved by CP are also solved by CPTR and CP-ES, which solve an additional 33 and 41 instances, respectively. CPTR-ES solves 134 instances, including all 120 instances solved by at least one of CP-ES or CPTR. Finally, CPTR-ES-DR solves 2 more instances beyond those solved by CPTR-ES. Within each group of solved instances, CPTR-ES and CPTR-ES-DR have the fastest solution times. An example of algorithm performance over a particular instance is shown in Appendix A.3 in the Online Supplement.

Before we proceed to a more detailed discussion of these results, we acknowledge that 83 of the instances were not solved by any of the algorithms, including our best-performing ones. Ultimately, the need to solve many MILO problems is an inherent limitation of all these cutting plane

algorithms. The MILO problems are used to generate cuts and validate optimality of a candidate cost vector. As these are MIPLIB problems, many of the MILO problems that need to be solved are time consuming relative to our chosen time limit. While trust regions can make certain MILO problems easier to solve, the algorithms eventually need to solve the full MILO problem to verify optimality. And we observe that many of the unsolved instances are precisely the ones in which MILO solution times increase dramatically when larger trust regions are considered.

7.1. Analysis and discussion of results

In this section, we analyze how the addition of trust regions and other enhancements lead to the observed performance results. The analysis is presented in the following order: (i) examining the effects of the early-stop heuristic on CP, (ii) analyzing the gains from adding trust regions to both CP and CP-ES, (iii) examining the effects of adding dimensionality reduction on CPTR-ES.

The two main factors that determine overall solution times are cut generation speed and strength of cuts. Solution times are reduced when both factors are improved together, or when the improvement in one eclipses any potential loss in the other. To facilitate the following discussion, we use Figure 5 to highlight the average cut generation speed over different trust region sizes and enhancements. In each of the three following subsections, we draw from Figure 5 and additional information about the strength of cuts (tailored for each section) to analyze the results. In this subsection, all values appearing in figures are plotted on a log scale.




Figure 5 The average per-instance cut generation times within different trust region sizes and with enhancements. Trust region sizes correspond to p in equation (9) and (13). Each data point is the average cut generation time within a trust region size for a particular instance. $p = 64+$ denotes all finite trust region sizes of $p \geq 64$, whereas No TR denotes no trust region (i.e., $p = \infty$).

7.1.1. Early-stop heuristic. Cut generation times within CP can be very high, as evident from Figure 5. The median of the average per-instance cut generation times exceeds 200 seconds, and CP fails to generate any cut within the time limit in over 25% of the instances. When the early-stop heuristic is applied, the median is reduced to 7 seconds, close to the early-stop τ value of 5 seconds. CP-ES is thus capable of generating many more cuts, as described next.

The iteration count for instances solved by CP-ES is shown in Figure 6. First, the ability to generate cuts faster allows CP-ES to generate many more cuts, which enables it to solve many more instances. Second, we observe that compared to the instances solved by CP, the early-stop heuristic does not significantly increase the iteration count. Part of this reason is that in a number of instances, CP and CP-ES are equivalent because the early-stop feature is never used; many of the instances that can be solved by CP are naturally ones in which cuts can be computed quickly. Nonetheless, the results in general suggest that cuts generated by the early-stop heuristic with $\tau = 5$ are not significantly weaker.




Figure 6 Iteration count and instance solvability when early-stop heuristic is added to CP.

Note that lowering the value of τ will not necessarily result in faster cut generation. For example, there exist many instances in which the average cut generation time is higher than τ (as observed in Figure 5), i.e., it takes longer than τ seconds on average to find any violated forward-feasible point. Cut generation times can also vary significantly depending on the given cost vector. In particular, many instances with an average cut generation time less than τ may still have a large number of iterations in which cut generation times are higher than τ .

7.1.2. Trust regions. Cut generation times can be substantially reduced using trust regions, as observed in Figure 5. For example, average cut generation time within trust regions of size $p \leq 8$ can be orders of magnitudes lower than without trust regions. Like the early-stop heuristic, this

reduction in cut generation time allows more cuts to be generated when instances are unsolved by CP, and the improved cut generation speed leads to many more instances being solved, as seen in Figure 7a. Of course, as trust region sizes increase, we observe that the cut generation times also tend to increase.

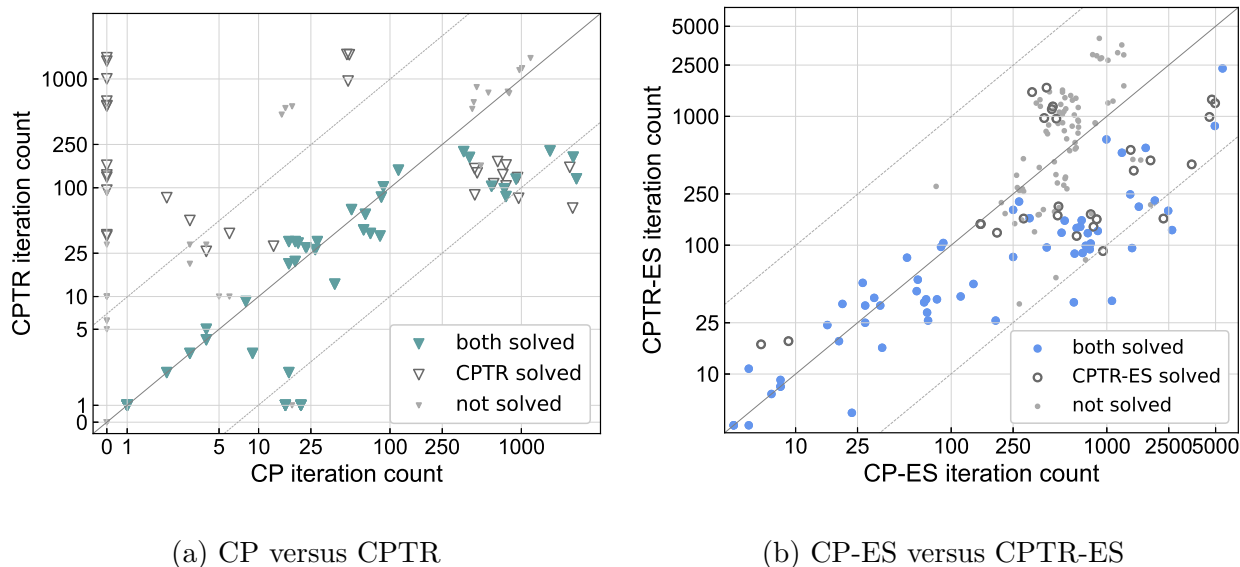


Figure 7 Iteration count and instance solvability after adding trust regions to the CP and CP-ES algorithms.

For the instances solved by CP, CPTR solves those same instances in significantly fewer iterations. This observation highlights that the cuts computed by CPTR are stronger. For example, the instances that require over 250 cuts from CP require around an order of magnitude fewer cuts with CPTR. Comparing CPTR-ES to CP-ES, the reduction in iterations is similar, as seen in Figure 7b. In general, the more cuts CP and CP-ES require to solve an instance, the greater the impact of trust regions on reducing iteration count.

The reduction in cut generation time and the increase in strength of cuts make trust regions very effective. Nonetheless, the magnitude of improvement in these two factors varies across instances. For example, reduction in cut generation time depends on how many cuts are computed within trust regions of smaller sizes. Secondly, a trust region of a particular size may lead to much faster cut generation in one instance than when the same trust region is applied to a different instance. To understand this phenomenon, it is worth noting that trust regions are, by definition, an additional set of constraints on the forward optimization problem. In some MIPLIB instances, the addition of trust regions of any size can “simplify” the forward-feasible region and lead to faster cut generation, whereas in other instances large trust regions may make the problem even more difficult to solve.

In Table 1 we show the results of adding trust regions to CP-ES over a subset of instances. The results are partitioned into three categories: ones that benefit mainly from improved cut strength,

Instance	CP-ES		CPTR-ES		Change (%)		
	num.it.	time (s)	num.it.	time (s)	num.it.	c.g.time	time
<i>Instances that benefit from reduced iteration count</i>							
mzzv42z.t3	945	>3600	90	285	-90	-18	-92
mzzv11.t2	732	3181	99	404	-86	-7	-87
drayage-25-23.t2	875	2657	129	321	-85	-19	-88
csched008.t1	23	260	5	64	-78	-2	-75
sp150x300d.t2	2035	181	222	29	-89	46	-84
<i>Instances that benefit from reduced cut generation time</i>							
seymour1.t2	396	>3600	970	263	145	-97	-93
seymour.t1	443	>3600	1136	449	156	-95	-88
roi2alpha3n4.t2	197	>3600	125	524	-37	-77	-85
neos-3083819-nubu.t3	250	1837	188	343	-25	-75	-81
ran14x18-disj-8.t1	994	726	662	96	-33	-80	-87
<i>Instances that benefit from both factors</i>							
neos-4954672-berkel.t1	1909	>3600	455	258	-76	-70	-93
50v-10.t2	4567	>3600	990	121	-78	-85	-97
drayage-25-23.t1	861	>3600	159	270	-82	-60	-92
drayage-100-23.t1	788	>3600	174	274	-78	-66	-92
csched008.t3	193	2768	26	130	-87	-66	-95
neos5.t3	81	2760	38	0	-53	-100	-100
glass-sc.t3	640	2074	136	146	-79	-67	-93
ran14x18-disj-8.t2	1777	1289	569	95	-68	-77	-93
neos-4338804-snowy.t2	1079	652	37	8	-97	-65	-99
mik-250-20-75-4.t1	2626	624	131	7	-95	-78	-99

Table 1 A subset of instances which illustrate the two key benefits of trust regions. The number of iterations (num.it.) and total time in seconds are provided, as well as the percentage in change of CPTR-ES over CP-ES in the number of iterations, cut generation time (c.g.time), and total solution time.

ones that benefit mainly from reduced cut generation time, and ones that benefit from both. While the results shown are for some of the better performing instances, they show performance over many different MIPLIB problem structures and paint a representative picture of the overall results. Specifically, in the majority of instances, average cut generation times are reduced, which alone is sufficient in reducing solution time. When coupled with improved cut strength, even larger reductions in solution time can be expected. In instances where average cut generation time is not reduced for the reasons described in the previous paragraph, reduced solution times can still be expected as a result of reduced iteration count. Finally, we note that the comparison is made over CP-ES, which already performs significantly better than CP, the baseline algorithm.

7.1.3. Dimensionality reduction. Figure 8 shows the iteration count and overall performance over the 134 instances solved by both the CPTR-ES and CPTR-ES-DR algorithms.

The first observation is that with very few exceptions, CPTR-ES-DR requires more iterations to solve each instance. This illustrates that considering lower dimensional trust regions generally leads to weaker cuts. The second observation illustrates the distinct trade-off that exists between




Figure 8 Comparison between the CPTR-ES and CPTR-ES-DR over instances that are solved by both.

cut generation speed and cut strength: when CPTR-ES-DR takes too many additional iterations, depicted by points being far to the left of the diagonal in Figure 8, CPTR-ES generally performs better. When CPTR-ES and CPTR-ES-DR take a similar number of iterations, i.e., points close to the diagonal, CPTR-ES-DR performs better. While CPTR-ES-DR does lead to more iterations in general, if this increase is modest, the reduction in cut generation time will lead to better overall performance. On the other hand, if the increase in iteration count is too large, CPTR-ES will generally perform better. The parameters of the dimensionality reduction enhancement can be tuned to balance this trade-off, which we discuss in the next subsection.

7.2. Sensitivity analysis

In the previous section, we analyzed results for a fixed set of baseline parameters to illustrate that performance gains can be achieved over a wide variety of problem structures without tuning. In this section, we consider several modifications of the baseline parameter values and show that even minor algorithm tuning can lead to dramatic gains. We consider modifications to most parameter values, as shown in Table 2, and report the change in performance over the baseline values. For simplicity, we consider each parameter modification independently, holding all other parameters fixed to their baseline values. We choose a convenience sample of the “_t1” instances solved by both the CPTR-ES and CPTR-ES-DR algorithms using baseline parameters, resulting in a total of 42 instances.

The percentage change in total iteration count and solution time for CPTR-ES and CPTR-ES-DR, relative to the baseline parameter values, are shown in Figure 9. The results highlight that, on average, the baseline parameter values perform quite well relative to the modified parameter values. However, we also show that it is possible to achieve significantly better performance if we

Algorithm Parameters	TR				ES	DR		
	p^0	δ	k^*	i^*	$\tau(s)$	κ	h^*	q
Baseline Values	1	2	2	10	5	0.03	10	0.8
		1.5		5	1	0.01	5	0.7
		4		25	20	0.05	15	0.9

Table 2 List of considered algorithm parameters for the CPTR algorithm and the two enhancements.

could identify the best parameter values for each instance independently (markers denoted “best”). For example, for CPTR-ES, if we could choose the best values for all three parameters for all instances, we could further reduce solution times by over 33%. Similarly, for CPTR-ES-DR, it is possible to achieve an 11% reduction in solution time with the best parameter values for each instance. While it is not possible to know the best parameters for each instance in advance, the results demonstrate that our default parameters tend to work well and that some tuning could lead to further improvements.

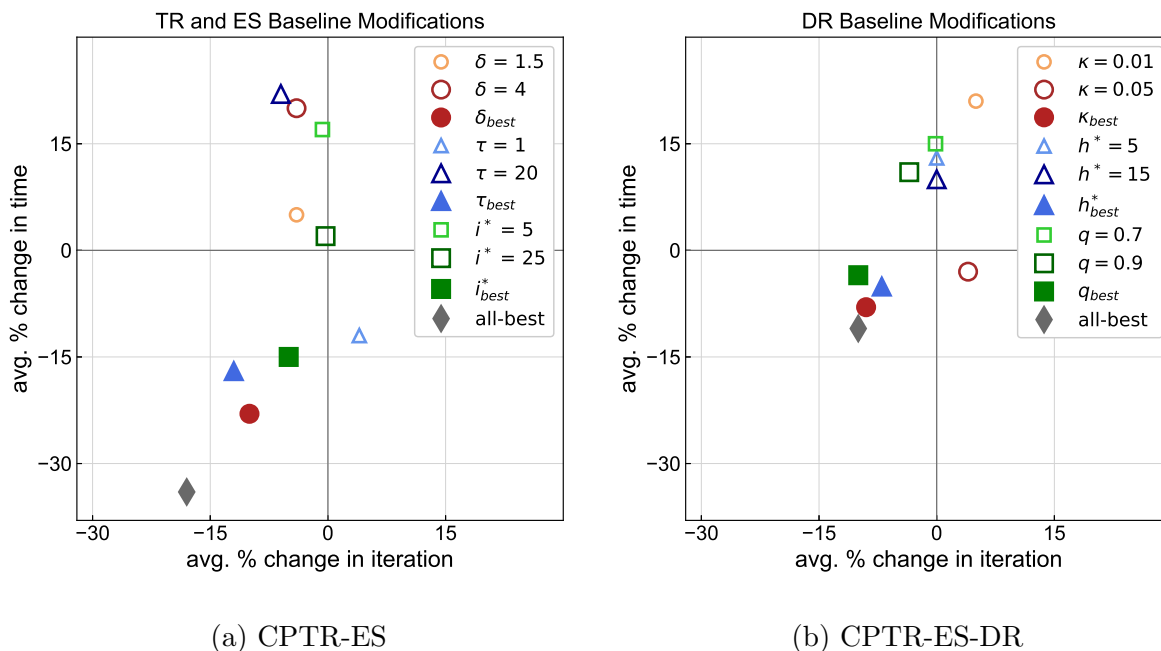


Figure 9 Modification of parameter values for the CPTR-ES and CPTR-ES-DR. Baseline values are $\delta = 2, \tau = 5, i^* = 10, \tau = 5, \kappa = 0.03, h^* = 10, q = 0.8$. The “best” subscript denotes the results when the best of the three (including baseline) parameter values is chosen for each instance, whereas the “all-best” label corresponds to the case where the best of all parameter values is chosen.

Comparing Figures 9a and 9b, it appears that solution time is more sensitive to the parameters for trust regions and the early-stop heuristic than the parameters for the dimensionality reduction

enhancement. This result is intuitive since trust regions and the early-stop heuristic yield much larger improvements when added to any algorithm (see Figure 4).

Based on the previous results, we derive a few general insights into parameter value choices for the trust regions and early stop heuristic:

- **Trust region growth rate (δ).** A significant gain in total solution time is observed when the trust region growth rate is high, i.e., $\delta = 4$. When the size of trust regions are increased too rapidly, the computational gains from solving smaller cut generation problems are generally reduced.
- **Frequency of trust region removal (i^*).** The total solution time is higher when $i^* = 5$ than when $i^* = 10$ or $i^* = 20$. This observation suggests that removing trust regions too frequently may reduce the computational gains of faster cut generation, especially when trust regions are small.
- **Early-stop value (τ).** Cuts may be weaker when τ is decreased. For example, $\tau = 1$ is the only parameter modification that leads to higher total iteration count. This is consistent with our theory, which suggests that the gain in cut strength when using trust regions comes from the computation of the extreme points of the corresponding subregion. When τ is low, an interior point of the subregion may be returned instead of an extreme point. Nonetheless, we observe through the progression of $\tau = 20$, $\tau = 5$ (baseline) and $\tau = 1$ that the improvement in cut generation speed generally outweighs the increase in iteration count from a lower τ value.

7.3. Takeaways from numerical results

The three main takeaways from our numerical results are: i) trust regions improve the strength of cuts and reduce cut generation time, particularly when trust regions are small, ii) the early-stop heuristic further reduces cut generation time, whereas the dimensionality reduction can help but comes with a potential cost of reducing the strength of cuts, iii) default parameter settings seem to work well, but additional tuning of algorithm parameters can lead to significant improvements in solution time.

8. Conclusion

In this paper, we develop a novel class of cutting plane algorithms for solving inverse MILO problems. Our approach of using trust regions to speed up computation is simple but effective, and is well-supported by our insights on the optimality conditions of inverse MILO problems. Through extensive computational experiments, we demonstrate that our cut generation techniques and enhancements are highly effective across a large range of inverse MILO problems.

References

- Afşar, Sezin, Luce Brotcorne, Patrice Marcotte, Gilles Savard. 2021. Revenue optimization in energy networks involving self-scheduled demand and a smart grid. *Computers & Operations Research* **134** 105366.
- Agarwal, Richa, Özlem Ergun. 2010. Network design and allocation mechanisms for carrier alliances in liner shipping. *Operations research* **58**(6) 1726–1742.
- Ahuja, Ravindra K, James B Orlin. 2001. Inverse optimization. *Operations Research* **49**(5) 771–783.
- Aswani, Anil, Zuo-Jun Shen, Auyon Siddiq. 2018. Inverse optimization with noisy data. *Operations Research* **66**(3) 870–892.
- Babier, Aaron, Timothy CY Chan, Taewoo Lee, Rafid Mahmood, Daria Terekhov. 2021. An ensemble learning framework for model fitting and evaluation in inverse linear optimization. *Informs Journal on Optimization* **3**(2) 119–138.
- Babier, Aaron, Rafid Mahmood, Andrea L McNiven, Adam Diamant, Timothy CY Chan. 2020. Knowledge-based automated planning with three-dimensional generative adversarial networks. *Medical Physics* **47**(2) 297–306.
- Bertsimas, Dimitris, Vishal Gupta, Ioannis Ch Paschalidis. 2012. Inverse optimization: A new perspective on the black-litterman model. *Operations research* **60**(6) 1389–1403.
- Bertsimas, Dimitris, Vishal Gupta, Ioannis Ch Paschalidis. 2015. Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming* **153**(2) 595–633.
- Birge, John R, Ali Hortaçsu, J Michael Pavlin. 2017. Inverse optimization for the recovery of market structure from market outcomes: An application to the miso electricity market. *Operations Research* **65**(4) 837–855.
- Brotcorne, Luce, Fabien Cirinei, Patrice Marcotte, Gilles Savard. 2011. An exact algorithm for the network pricing problem. *Discrete Optimization* **8**(2) 246–258.
- Bulut, Aykut, Ted K Ralphs. 2015. On the complexity of inverse mixed integer linear optimization. Tech. rep., Tech. rep. COR@ L Laboratory Report 15T-001-R3, Lehigh University.
- Burton, Didier, Ph L Toint. 1992. On an instance of the inverse shortest paths problem. *Mathematical Programming* **53**(1-3) 45–61.
- Chan, Timothy CY, Tim Craig, Taewoo Lee, Michael B Sharpe. 2014. Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research* **62**(3) 680–695.
- Chan, Timothy CY, Neal Kaw. 2020. Inverse optimization for the recovery of constraint parameters. *European Journal of Operational Research* **282**(2) 415–427.
- Chan, Timothy CY, Taewoo Lee, Daria Terekhov. 2019. Inverse optimization: Closed-form solutions, geometry, and goodness of fit. *Management Science* **65**(3) 1115–1135.
- Chan, Timothy CY, Rafid Mahmood, Ian Yihang Zhu. 2021. Inverse optimization: Theory and applications. *arXiv preprint arXiv:2109.03920* .

- Chen, Lu, Yuyi Chen, André Langevin. 2021. An inverse optimization approach for a capacitated vehicle routing problem. *European Journal of Operational Research* .
- Chow, Joseph YJ, Will W Recker. 2012. Inverse optimization with endogenous arrival time constraints to calibrate the household activity pattern problem. *Transportation Research Part B: Methodological* **46**(3) 463–479.
- Chow, Joseph YJ, Stephen G Ritchie, Kyungsoo Jeong. 2014. Nonlinear inverse optimization for parameter estimation of commodity-vehicle-decoupled freight assignment. *Transportation Research Part E: Logistics and Transportation Review* **67** 71–91.
- Duan, Zhaoyang, Lizhi Wang. 2011. Heuristic algorithms for the inverse mixed integer linear programming problem. *Journal of Global Optimization* **51**(3) 463–471.
- Erkin, Zeynep, Matthew D Bailey, Lisa M Maillart, Andrew J Schaefer, Mark S Roberts. 2010. Eliciting patients’ revealed preferences: an inverse markov decision process approach. *Decision Analysis* **7**(4) 358–365.
- Esfahani, Peyman Mohajerin, Soroosh Shafieezadeh-Abadeh, Grani A Hanasusanto, Daniel Kuhn. 2018. Data-driven inverse optimization with imperfect information. *Mathematical Programming* **167**(1) 191–234.
- Esfandeh, Tolou, Changhyun Kwon, Rajan Batta. 2016. Regulating hazardous materials transportation by dual toll pricing. *Transportation Research Part B: Methodological* **83** 20–35.
- Gleixner, Ambros, Gregor Hendel, Gerald Gamrath, Tobias Achterberg, Michael Bastubbe, Timo Berthold, Philipp M Christophel, Kati Jarck, Thorsten Koch, Jeff Linderoth, et al. 2019. Miplib 2017: Data-driven compilation of the 6th mixed-integer programming library. *Optimization online preprint: http://www.optimization-online.org/DB_HTML/2019/07/7285.html*. Submitted to *Mathematical Programming Computation* .
- Güler, Çiğdem, Horst W Hamacher. 2010. Capacity inverse minimum cost flow problem. *Journal of Combinatorial Optimization* **19**(1) 43–59.
- Heuberger, Clemens. 2004. Inverse combinatorial optimization: A survey on problems, methods, and results. *Journal of combinatorial optimization* **8**(3) 329–361.
- Houghtalen, Lori, Özlem Ergun, Joel Sokol. 2011. Designing mechanisms for the management of carrier alliances. *Transportation Science* **45**(4) 465–482.
- Iyengar, Garud, Wanmo Kang. 2005. Inverse conic programming with applications. *Operations Research Letters* **33**(3) 319–330.
- Keshavarz, Arezou, Yang Wang, Stephen Boyd. 2011. Imputing a convex objective function. *2011 IEEE International Symposium on Intelligent Control*. IEEE, 613–619.
- Kuiteing, Aimé Kamgaing, Patrice Marcotte, Gilles Savard. 2017. Network pricing of congestion-free networks: The elastic and linear demand case. *Transportation Science* **51**(3) 791–806.

- Kuiteing, Aimé Kamgaing, Patrice Marcotte, Gilles Savard. 2018. Pricing and revenue maximization over a multicommodity transportation network: the nonlinear demand case. *Computational Optimization and Applications* **71**(3) 641–671.
- Lamperski, Jourdain B, Andrew J Schaefer. 2015. A polyhedral characterization of the inverse-feasible region of a mixed-integer program. *Operations Research Letters* **43**(6) 575–578.
- Linderoth, Jeff, Stephen Wright. 2003. Decomposition algorithms for stochastic programming on a computational grid. *Computational Optimization and Applications* **24**(2-3) 207–250.
- Marcotte, Patrice, Anne Mercier, Gilles Savard, Vedat Verter. 2009. Toll policies for mitigating hazardous materials transport risk. *Transportation science* **43**(2) 228–243.
- Naghavi, Mostafa, Ali Asghar Foroughi, Masoud Zarepisheh. 2019. Inverse optimization for multi-objective linear programming. *Optimization Letters* **13**(2) 281–294.
- Rahmaniani, Ragheb, Teodor Gabriel Crainic, Michel Gendreau, Walter Rei. 2017. The benders decomposition algorithm: A literature review. *European Journal of Operational Research* **259**(3) 801–817.
- Rathore, Himanshu, Suresh Kumar Jakhar. 2021. Differential carbon tax policy in aviation: One stone that kills two birds? *Journal of Cleaner Production* **296** 126479.
- Ruiz, Carlos, Antonio J Conejo, Dimitris J Bertsimas. 2013. Revealing rival marginal offer prices via inverse optimization. *IEEE Transactions on Power Systems* **28**(3) 3056–3064.
- Saez-Gallego, Javier, Juan M Morales, Marco Zugno, Henrik Madsen. 2016. A data-driven bidding model for a cluster of price-responsive consumers of electricity. *IEEE Transactions on Power Systems* **31**(6) 5001–5011.
- Schaefer, Andrew J. 2009. Inverse integer programming. *Optimization Letters* **3**(4) 483–489.
- Troutt, Marvin D, Wan-Kai Pang, Shui-Hung Hou. 2006. Behavioral estimation of mathematical programming objective function coefficients. *Management science* **52**(3) 422–434.
- Turner, Sarina DO, Timothy CY Chan. 2013. Examining the leed rating system using inverse optimization. *Journal of Solar Energy Engineering* **135**(4) 040901.
- Utz, Sebastian, Maximilian Wimmer, Markus Hirschberger, Ralph E Steuer. 2014. Tri-criterion inverse portfolio optimization with application to socially responsible mutual funds. *European Journal of Operational Research* **234**(2) 491–498.
- Wang, Lizhi. 2009. Cutting plane algorithms for the inverse mixed integer linear programming problem. *Operations Research Letters* **37**(2) 114–116.
- Xu, Susan Jia, Mehdi Nourinejad, Xuebo Lai, Joseph YJ Chow. 2018. Network learning via multiagent inverse transportation problems. *Transportation Science* **52**(6) 1347–1364.
- Yu, Shi, Haoran Wang, Chaosheng Dong. 2020. Learning risk preferences from investment portfolios using inverse optimization. *arXiv preprint arXiv:2010.01687* .

-
- Zhang, Jianzhong, Zhenhong Liu. 1996. Calculating some inverse linear programming problems. *Journal of Computational and Applied Mathematics* **72**(2) 261–273.
- Zhang, Jianzhong, Chengxian Xu. 2010. Inverse optimization for linearly constrained convex separable programming problems. *European Journal of Operational Research* **200**(3) 671–679.
- Zhang, Jianzhong, Liwei Zhang, Xiantao Xiao. 2010. A perturbation approach for an inverse quadratic programming problem. *Mathematical Methods of Operations Research* **72**(3) 379–404.
- Zhang, Jing, Sepideh Pourazarm, Christos G Cassandras, Ioannis Ch Paschalidis. 2018. The price of anarchy in transportation networks: Data-driven evaluation and reduction strategies. *Proceedings of the IEEE* **106**(4) 538–553.
- Zhou, Ying, Lizhi Wang, James D McCalley. 2011. Designing effective and efficient incentive policies for renewable energy in generation expansion planning. *Applied Energy* **88**(6) 2201–2209.
- Zverovich, Victor, Csaba I Fábrián, Eldon FD Ellison, Gautam Mitra. 2012. A computational study of a solver system for processing two-stage stochastic lps with enhanced benders decomposition. *Mathematical Programming Computation* **4**(3) 211–238.

Online Supplement

This is the Online Supplement to the paper “*Inverse Mixed Integer Optimization: Polyhedral Insights and Trust Region Methods*”. Appendix A.1 contains all proofs for theoretical results presented in the paper. Appendix A.2 provides an illustration of why inverse linear optimization generates poor approximations to inverse mixed integer optimization. Appendix A.3 illustrates how the algorithms perform over a particular instance, whereas Appendix A.4 shows the computational results over all the instances.

A.1. Proofs

Several of the proofs in this section use the result that the inverse-feasible region of the sets \mathcal{X} , $\text{conv}(\mathcal{X})$, and $\mathcal{E}(\mathcal{X})$ are the same, i.e., $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) = \mathcal{C}(\hat{\mathbf{x}}, \text{conv}(\mathcal{X})) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{E}(\mathcal{X}))$. This result was previously established in Wang (2009).

Proof of Lemma 1 The main result that must be proved is that for any given $\mathbf{x} \in \mathcal{X}$, $\mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{x}\}) = \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}) \forall \lambda > 0$. Given that this statement is true, the result of $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{E}(\mathcal{X})) = \mathcal{C}(\hat{\mathbf{x}}, \bar{\mathcal{E}}(\mathcal{X}))$ comes trivially. The proof comes in two parts. we first prove that (i) for any given $\mathbf{x} \in \mathcal{X}$, $\mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{x}\}) \subseteq \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}) \forall \lambda > 0$. We then prove that (ii) for any given $\mathbf{x} \in \mathcal{X}$ and $\lambda > 0$, $\mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}) \subseteq \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{x}\})$.

Proof: (i) Let $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{x}\})$. By definition, $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) \leq 0$. Similarly, $\lambda \mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) \leq 0 \forall \lambda > 0$. Since $\lambda \mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) = \mathbf{c}^\top(\hat{\mathbf{x}} - \hat{\mathbf{x}} - \lambda(\mathbf{x} - \hat{\mathbf{x}})) = \mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x}))$, the statement $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}) \forall \lambda > 0$ must be true. (ii) For any $\hat{\mathbf{x}} \in \mathcal{X}$ and $\lambda > 0$, let \mathbf{c} denote any vector such that $\lambda \mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) \leq 0$, i.e., $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\})$. Because λ is positive, $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) \leq 0$. Thus, $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \{\mathbf{x}\})$. \square

Proof of Theorem 1 We prove by contradiction. (\Leftarrow) We assume $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) = \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$, but \mathcal{G} is not a generator set. This implies that there exists some \mathbf{c} such that either (i) $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$ but $\mathbf{c} \notin \mathcal{C}(\hat{\mathbf{x}}, \mathcal{G})$ or (ii) $\mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{G})$ but $\mathbf{c} \notin \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$. If (i) is true, then there exists a cost vector \mathbf{c} such that $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{x}) \leq 0 \forall \mathbf{x} \in \mathcal{X}$ but $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{g}) > 0$ for some $\mathbf{g} \in \mathcal{G}$. By definition, this \mathbf{g} cannot be in \mathcal{X} , i.e., $\mathbf{g} \notin \mathcal{X}$, and thus $\mathbf{g} \notin \text{conv}(\mathcal{X})$. Furthermore, there cannot exist any $\mathbf{x} \in \mathcal{X}, \lambda > 0$ such that $\hat{\mathbf{x}} + \lambda(\mathbf{x} - \hat{\mathbf{x}}) = \mathbf{g}$. This implies that $\mathbf{g} \notin \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$, which is a contradiction. The same argument can be applied to show that (ii) leads to a contradiction.

(\Rightarrow) We assume that \mathcal{G} is in fact a generator set, but $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) \neq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. This implies that there exists a point \mathbf{y}^* such that either (i*) $\mathbf{y}^* \in \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ but $\mathbf{y}^* \notin \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ or (ii*) $\mathbf{y}^* \in \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ but $\mathbf{y}^* \notin \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$. If (i*) is true, then there exists a separating hyperplane between $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ and \mathbf{y}^* . The normal vector of this hyperplane that is pointing in the direction of $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ defines a cost vector \mathbf{c} such that all points $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{y}) \leq 0 \forall \mathbf{y} \in \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$, which implies that $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{g}) \leq 0 \forall \mathbf{g} \in \mathcal{G}$. On

the other hand, \mathbf{y}^* lies on the opposite side of the hyperplane, which implies that $\mathbf{c}(\hat{\mathbf{x}} - \mathbf{y}^*) > 0$. Finally, because $\mathbf{y}^* \in \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ and thus must be constructed by a convex combination of the rays $\{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}_{\lambda \geq 0, \mathbf{x} \in \mathcal{X}}$, there must exist at least one $\mathbf{x}^* \in \mathcal{X}$ such that $\mathbf{c}(\hat{\mathbf{x}} - \mathbf{x}^*) > 0$. This contradicts the initial assumption that $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{G})$. The same argument can be applied to show that (ii*) leads to a contradiction. \square

Proof of Corollary 1 (\Rightarrow) Note that by the definition of $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$, $\mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$. If $\mathcal{G} \subseteq \mathcal{X}$ is a generator set, then $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G}) = \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ by Theorem 1. Thus, $\mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. These statements together yield $\mathcal{G} \subseteq \mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. (\Leftarrow) Suppose $\mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ for some $\mathcal{G} \subseteq \mathcal{X}$. This implies that for any $\mathbf{x} \in \mathcal{X}$, $\mathbf{x} \in \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ and $\{\mathbf{y}(\lambda, \hat{\mathbf{x}}, \mathbf{x})\}_{\lambda \geq 0} \in \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. Therefore, $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. Since $\mathcal{G} \subseteq \mathcal{X}$, the reverse must also be true, i.e., $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G}) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$. Thus, $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G}) = \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$. Applying Theorem 1, $\mathcal{G} \subseteq \mathcal{X}$ must be a generator set. \square

Proof of Theorem 2 (\Rightarrow) We first prove the forward direction, i.e., that if $\mathcal{G} \subseteq \mathcal{X}$ is a generator set, then $\mathcal{B}_\epsilon(\hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X})$ must be a subset of $\text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\})$. We prove by contradiction. Suppose that there exists a point \mathbf{y} in $\mathcal{B}_\epsilon(\hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X})$ such that $\mathbf{y} \notin \text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\})$. This implies that there exists a separating hyperplane between \mathbf{y} and $\text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\})$. The point $\hat{\mathbf{x}}$ must lie in only of side of this hyperplane. Taking the normal vector of this hyperplane to be \mathbf{c} , we must have that either $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{g}) \leq 0 \forall \mathbf{g} \in \mathcal{G}$ and $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{y}) > 0$, or that $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{g}) > 0 \forall \mathbf{g} \in \mathcal{G}$ and $\mathbf{c}^\top(\hat{\mathbf{x}} - \mathbf{y}) \leq 0$. Given that $\mathbf{y} \in \text{conv}(\mathcal{X})$, this contracts the initial assumption that $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) = \mathcal{C}(\hat{\mathbf{x}}, \text{conv}(\mathcal{X})) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{G})$.

(\Leftarrow) We now prove the reverse direction. In particular, we prove that if the convex hull of $\mathcal{G} \subseteq \mathcal{X}$ with $\hat{\mathbf{x}}$ contains the set $\mathcal{B}_\epsilon(\hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X})$ for some $\epsilon > 0$, then \mathcal{G} must be a forward-feasible generator set, i.e., $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}) = \mathcal{C}(\hat{\mathbf{x}}, \mathcal{G})$. We first observe that if $\mathcal{B}_\epsilon(\hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X}) \subseteq \text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\})$, then

$$\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{B}(\epsilon, \hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X})) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\})).$$

We will now prove that both sides of this equation can be simplified to obtain $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$, at which point we can apply Corollary 1 to show that $\mathcal{G} \subseteq \mathcal{X}$ is a generator set. The major steps in this proof rely on the definition of the set $\mathcal{Y}(\hat{\mathbf{x}}, \cdot)$, and we refer to Theorem 1 for its definition.

We first simplify the left side of the equation. Note that by the definition of $\mathcal{Y}(\hat{\mathbf{x}}, \cdot)$, which is a set polyhedral cone pointed at $\hat{\mathbf{x}}$, $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{B}(\epsilon_1, \hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X})) = \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{B}(\epsilon_2, \hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X}))$ for any $\epsilon_1, \epsilon_2 > 0$. This implies that $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{B}(\epsilon, \hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X})) = \mathcal{Y}(\hat{\mathbf{x}}, \text{conv}(\mathcal{X}))$, since $\lim_{\epsilon \rightarrow \infty} \mathcal{B}(\epsilon, \hat{\mathbf{x}}) \cap \text{conv}(\mathcal{X}) = \text{conv}(\mathcal{X})$. The set $\mathcal{Y}(\hat{\mathbf{x}}, \text{conv}(\mathcal{X}))$ can be further simplified into $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$ by the definition of $\mathcal{Y}(\hat{\mathbf{x}}, \cdot)$. We now simplify the right side of the equation. Given that $\hat{\mathbf{x}} \in \mathcal{Y}(\hat{\mathbf{x}}, \text{conv}(\mathcal{G} \cup \{\hat{\mathbf{x}}\}))$ by definition, the right side of the equation can be simplified to $\mathcal{Y}(\hat{\mathbf{x}}, \text{conv}(\mathcal{G}))$. Furthermore, $\mathcal{Y}(\hat{\mathbf{x}}, \text{conv}(\mathcal{G})) = \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$.

These steps lead to the result that $\mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X}) \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$. Since $\mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{X})$, and we assumed $\mathcal{G} \subseteq \mathcal{X}$, $\mathcal{G} \subseteq \mathcal{X} \subseteq \mathcal{Y}(\hat{\mathbf{x}}, \mathcal{G})$ and by Corollary (1), \mathcal{G} must be a forward-feasible generator set. \square

A.2. Approximations using Inverse Linear Optimization

Let \mathcal{X}_{LP} denote a relaxed forward-feasible region where all integrality constraints present in $\mathcal{X} = \{\mathbf{x} \mid \mathbf{Ax} \geq \mathbf{b}, \mathbf{x} \in \mathbb{Z}^{n-q} \times \mathbb{R}^q\}$ are ignored. Similarly, let $\mathbf{c}_{LP} \in \arg \min_{\mathbf{c} \in \mathcal{P}} \{\|\mathbf{c} - \mathbf{c}^0\| : \mathbf{c} \in \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}_{LP})\}$ denote the solution to an arbitrary inverse linear optimization model where \mathcal{X}_{LP} is considered instead of \mathcal{X} . Since $\mathcal{X} \subseteq \mathcal{X}_{LP}$, the inverse-feasible region $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}_{LP})$ must be a subset of $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$, i.e., $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}_{LP}) \subseteq \mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$.

The main drawback of using inverse linear optimization over \mathcal{X}_{LP} is that the set $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}_{LP})$ can be very different from the true inverse-feasible region $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X})$. This means that \mathbf{c}_{LP} can be very far from the (or in fact any) optimal solution \mathbf{c}^* . Furthermore, there can be a large number of different regions \mathcal{X}_{LP} that define the same mixed integer set \mathcal{X} when the integrality constraints are considered. Specifically, for a fixed \mathbf{A} and \mathbf{b} generating $\mathcal{X} = \{\mathbf{x} \mid \mathbf{Ax} \geq \mathbf{b}, \mathbf{x} \in \mathbb{Z}^{n-q} \times \mathbb{R}^q\}$, there can exist many different $\hat{\mathbf{A}} \neq \mathbf{A}$ and $\hat{\mathbf{b}} \neq \mathbf{b}$ that can replace \mathbf{A} and \mathbf{b} and generate the same \mathcal{X} . However, each of these $\hat{\mathbf{A}}$'s and $\hat{\mathbf{b}}$'s will generate different \mathcal{X}_{LP} sets. This means that the estimates of \mathbf{c}_{LP} are entirely dependent on the particular definition of \mathbf{A} and \mathbf{b} , making \mathbf{c}_{LP} , and the quality of \mathbf{c}_{LP} relative to \mathbf{c}^* , unstable.




Figure A.1 Plot of four different linear constraint models that define the same integer forward-feasible region when integrality constraints are considered. The forward-region in all four figures consists of points $(0,0)$, $(0,1)$, $(1,0)$, and $(1,1)$. The true inverse-feasible region for the point $\hat{\mathbf{x}} = (1,1)$ is shown below each of the diagrams (as the cone between the two large perpendicular vectors), while the inverse-feasible regions of the relaxed forward-feasible region are highlighted in light blue.

For example, consider Figure A.1, where the forward-feasible region in each of the subfigures are the same, i.e., $\mathcal{X} = \{(0,0), (1,0), (0,1), (1,1)\}$. On the other hand, the linear constraints used to describe this region ($x_1, x_2 \geq 0$, and the two hyperplanes indicated by orange lines) are different, resulting in different \mathcal{X}_{LP} and very different inverse-feasible regions for the solution $\hat{\mathbf{x}} = (1,1)$. In particular, the inverse-feasible regions of $(\hat{\mathbf{x}}, \mathcal{X}_{LP})$ become increasingly smaller from Figures (1a)

to (1d), which would result in solutions \mathbf{c}_{LP} that become increasingly worse relative to \mathbf{c}^* for any choice of \mathbf{c}^0 . This example, with $\mathcal{C}(\hat{\mathbf{x}}, \mathcal{X}_{LP}) = \{\mathbf{0}\}$ in Figure (1d), also shows that there is no bound on the distance between \mathbf{c}_{LP} and \mathbf{c}^* (or similarly, between $\|\mathbf{c}_{LP} - \mathbf{c}^0\|_1$ and $\|\mathbf{c} - \mathbf{c}^0\|_1$). For these reasons, inverse linear optimization models may be poorly suited for solving inverse MILO models.

A.3. Example of Algorithm Performance over a Single Instance

Below, we provide an example of the performance of the five algorithms over the instance assign1-5-8_t1. In Figure A.2, the total solution time is plotted against the iteration count, and the slope of the line at any iteration reflects the cut generation speed at that particular iteration. The instance cannot be solved by CP, which has particularly high cut generation times, computing only 3 cuts by the one hour time limit. Adding the early-stop heuristic on CP (forming CP-ES) lowers cut generation times and results in a solution time of 2061 seconds with 250 iterations. On the other hand, adding trust regions on the CP (forming CPTR) lowers cut generation times more significant, *and* improves the strength of cuts, resulting in the instance being solved in 590 seconds while requiring only 81 iterations. We observe through Figure A.2 that cut generation times are initially very low, and gradually increase with iteration count (as a result of larger trust regions). CPTR-ES further reduces cut generation time while requiring the same number of iterations. Finally, CPTR-ES-DR requires 9 more iterations than CPTR-ES, but the decrease in average cut generation time dominates the increase in iteration count.




Figure A.2 Total solution time over iteration count for the inverse MILO instance assign1-5-8_t1. The instance is not solved by the CP algorithm, which manages only to compute 3 cuts within the time limit. CP-ES solves the instance in 2061 seconds using 250 iterations (not shown).

A.4. Table of Solved Instances

Table A.1: Performance profile of all instances solved by at least one of the algorithms.

Instance	Problem Size		CP		CP-ES		CPTR		CPTR-ES		CPTR-ES-DR	
	vars.	cons.	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)
50v-10.t1	2013	233	0	>3600	4964	>3600	1567	1919	1265	206	2460	370
50v-10.t2	2013	233	0	>3600	4567	>3600	1002	992	990	121	1731	240
50v-10.t3	2013	233	0	>3600	4732	>3600	1456	237	1351	164	3046	402
assign1-5-8.t1	156	161	2	>3600	250	2061	81	590	81	372	90	266
assign1-5-8.t2	156	161	4	>3600	71	455	26	38	26	11	26	14
assign1-5-8.t3	156	161	3	>3600	139	682	50	46	50	47	69	290
bppc4-08.t1	1456	111	0	>3600	1	108	0	>3600	1	108	1	108
bppc4-08.t2	1456	111	0	>3600	1	13	0	>3600	1	11	1	11
bppc4-08.t3	1456	111	0	>3600	1	5	0	>3600	1	5	1	6
cod105.t1	1024	1024	5	>3600	6	>3600	10	>3600	17	1894	27	2557
cod105.t2	1024	1024	0	127	0	121	0	122	0	100	0	117
cod105.t3	1024	1024	6	>3600	9	>3600	10	>3600	18	1474	24	1916
csched007.t2	1758	351	3	>3600	256	>3600	20	>3600	147	>3600	144	3475
csched008.t1	1536	351	17	1533	23	260	4	666	5	64	6	91
csched008.t2	1536	351	17	2137	115	2122	10	1489	40	2748	41	1164
csched008.t3	1536	351	18	>3600	193	2768	1	>3600	26	130	53	2711
cvs16r128-89.t2	3472	4633	0	>3600	489	>3600	6	>3600	200	1260	270	963
cvs16r128-89.t3	3472	4633	0	>3600	483	>3600	5	>3600	170	1368	270	1680
drayage-100-23.t1	11090	4630	660	>3600	788	>3600	174	245	174	274	195	242
drayage-100-23.t2	11090	4630	935	>3600	755	2370	124	218	124	247	150	235
drayage-100-23.t3	11090	4630	492	>3600	796	>3600	160	>3600	176	>3600	178	3520
drayage-25-23.t1	11090	4630	765	>3600	861	>3600	162	277	159	270	173	217
drayage-25-23.t2	11090	4630	912	3303	875	2657	129	283	129	321	116	217
drayage-25-23.t3	11090	4630	461	>3600	821	>3600	137	1069	139	659	172	684
eil33-2.t1	4516	32	89	75	89	71	93	102	104	107	102	102
eil33-2.t2	4516	32	86	101	86	88	87	127	97	142	116	161
eil33-2.t3	4516	32	51	105	52	105	67	405	80	386	93	661
enlight_hard.t1	200	100	0	0	0	0	0	0	0	0	0	0
enlight_hard.t2	200	100	0	0	0	0	0	0	0	0	0	0
enlight_hard.t3	200	100	0	0	0	0	0	0	0	0	0	0
gen-ip002.t1	41	24	63	704	60	229	44	600	44	90	53	105
gen-ip002.t2	41	24	28	615	28	51	34	559	34	44	44	52
gen-ip002.t3	41	24	65	893	61	309	52	688	54	144	70	189
germanrr.t1	10813	10779	0	>3600	1	24	0	>3600	1	21	1	23
germanrr.t2	10813	10779	0	>3600	1	16	0	>3600	1	22	1	22

Continued on next page

Instance	Problem Size		CP		CP-ES		CPTR		CPTR-ES		CPTR-ES-DR	
	vars.	cons.	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)
germanrr.t3	10813	10779	0	>3600	1	17	0	>3600	1	22	1	22
glass-sc.t1	214	6119	0	>3600	292	>3600	161	961	161	140	188	151
glass-sc.t2	214	6119	0	>3600	512	1826	125	1314	125	126	125	125
glass-sc.t3	214	6119	0	>3600	640	2074	131	1704	136	146	136	145
gmu-35-40.t1	1205	424	593	934	411	176	93	1323	96	137	98	143
gmu-35-50.t2	1919	435	442	>3600	621	292	86	1270	86	45	87	43
gmu-35-50.t3	1919	435	3	>3600	699	284	10	>3600	87	74	102	89
leo1.t1	6731	593	1	359	1	5	1	351	1	5	1	6
leo1.t2	6731	593	1	395	1	5	1	351	1	5	1	7
leo1.t3	6731	593	0	198	0	203	0	182	0	209	0	207
leo2.t1	11100	593	1	919	1	6	1	887	1	6	1	7
leo2.t2	11100	593	1	1034	1	5	1	885	1	6	1	7
leo2.t3	11100	593	0	496	1	5	0	432	1	11	1	11
markshare.4_0.t1	34	4	4	0	4	0	4	0	4	0	4	0
markshare.4_0.t2	34	4	4	0	4	0	4	0	4	0	4	0
markshare.4_0.t3	34	4	8	2	8	2	9	1	9	2	9	2
markshare2.t1	74	7	0	>3600	8	7	0	>3600	8	5	8	5
markshare2.t2	74	7	0	>3600	7	7	0	>3600	7	5	7	5
markshare2.t3	74	7	0	>3600	8	8	0	>3600	8	5	8	5
mas74.t1	151	13	21	445	1	5	1	0	1	0	1	0
mas74.t2	151	13	21	467	1	5	1	0	1	0	1	0
mas74.t3	151	13	21	458	1	5	1	0	1	0	1	0
mas76.t1	151	12	16	20	1	21	1	0	1	0	1	0
mas76.t2	151	12	16	21	1	21	1	0	1	0	1	0
mas76.t3	151	12	16	21	1	22	1	0	1	0	1	0
mcsched.t1	1747	2107	1	90	1	7	1	88	1	7	1	7
mcsched.t2	1747	2107	1	93	1	7	1	88	1	7	1	7
mcsched.t3	1747	2107	1	135	1	11	1	131	1	11	1	11
mik-250-20-75-4.t1	270	195	2626	580	2626	624	121	4	131	7	124	7
mik-250-20-75-4.t2	270	195	2336	>3600	2312	>3600	154	145	161	161	170	171
mik-250-20-75-4.t3	270	195	2480	451	2480	484	191	13	185	10	191	12
mzzv11.t1	10240	9499	616	>3600	640	>3600	109	653	118	604	137	568
mzzv11.t2	10240	9499	765	>3600	732	3181	104	423	99	404	112	453
mzzv11.t3	10240	9499	723	>3600	689	3318	132	698	156	815	152	527
mzzv42z.t1	11717	10460	764	2624	777	3027	86	258	93	292	112	259
mzzv42z.t2	11717	10460	740	2208	788	2326	100	186	103	257	111	186
mzzv42z.t3	11717	10460	953	>3600	945	>3600	80	230	90	285	102	314

Continued on next page

Instance	Problem Size		CP		CP-ES		CPTR		CPTR-ES		CPTR-ES-DR	
	vars.	cons.	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)
n5-3.t1	2550	1062	8994	2167	7731	1889	1432	245	737	121	816	213
n5-3.t2	2550	1062	5645	1301	4945	982	396	108	842	138	595	173
n5-3.t3	2550	1062	5866	1159	5544	988	397	99	2355	461	495	163
neos-2657525-crna.t2	524	342	0	>3600	19	116	0	>3600	18	55	18	55
neos-3083819-nubu.t1	8644	4725	403	1265	319	1009	180	537	162	533	156	304
neos-3083819-nubu.t2	8644	4725	116	1768	273	1977	116	1094	218	1967	145	777
neos-3083819-nubu.t3	8644	4725	364	2854	250	1837	165	244	188	343	193	243
neos-3381206-awhea.t2	2375	479	104	>3600	1490	>3600	119	>3600	379	3597	330	3177
neos-3627168-kasai.t1	1462	1655	0	>3600	3	36	0	>3600	3	36	4	11
neos-3627168-kasai.t2	1462	1655	0	>3600	3	33	0	>3600	3	34	4	11
neos-3627168-kasai.t3	1462	1655	0	>3600	3	69	0	>3600	3	64	3	34
neos-4338804-snowy.t1	1344	1701	0	>3600	614	105	36	2	36	2	39	6
neos-4338804-snowy.t2	1344	1701	0	>3600	1079	652	37	8	37	8	62	9
neos-4338804-snowy.t3	1344	1701	0	>3600	1453	1391	95	50	95	51	143	60
neos-4954672-berkel.t1	1533	1848	0	>3600	1909	>3600	10	>3600	455	258	476	322
neos-4954672-berkel.t2	1533	1848	0	>3600	3511	>3600	10	>3600	424	205	462	215
neos-4954672-berkel.t3	1533	1848	0	>3600	1423	>3600	10	>3600	549	1601	567	1547
neos-860300.t1	1385	850	20	139	16	58	20	101	24	86	33	192
neos-860300.t2	1385	850	19	122	20	66	31	165	35	144	40	162
neos-860300.t3	1385	850	17	185	28	180	22	172	25	146	24	193
neos17.t2	535	486	2457	>3600	675	1770	65	252	139	1612	144	238
neos17.t3	535	486	441	>3600	536	3221	150	974	155	1958	152	556
neos5.t1	63	63	6	>3600	69	688	38	0	38	0	38	0
neos5.t2	63	63	71	1318	67	470	36	0	36	0	36	0
neos5.t3	63	63	84	3343	81	2760	38	0	38	0	38	0
qap10.t1	4150	1820	19	514	35	568	24	641	34	579	34	585
qap10.t2	4150	1820	23	461	32	428	30	571	39	543	39	542
qap10.t3	4150	1820	27	451	27	264	37	854	51	1022	51	1063
ran14x18-disj-8.t1	504	447	0	>3600	994	726	625	180	662	96	731	96
ran14x18-disj-8.t2	504	447	0	>3600	1777	1289	569	88	569	95	685	103
ran14x18-disj-8.t3	504	447	0	>3600	1250	788	563	110	522	60	667	71
rocI-4-11.t1	6839	10883	1	268	1	296	1	321	1	369	1	368
rocI-4-11.t2	6839	10883	2	368	2	318	2	681	2	668	2	645
rocI-4-11.t3	6839	10883	3	226	5	83	3	146	4	56	4	63
rococoB10-011000.t1	4456	1667	0	>3600	1	9	0	>3600	1	9	1	9
rococoB10-011000.t2	4456	1667	0	>3600	1	9	0	>3600	1	9	1	9
rococoB10-011000.t3	4456	1667	0	>3600	1	11	0	>3600	1	10	1	10

Continued on next page

Instance	Problem Size		CP		CP-ES		CPTR		CPTR-ES		CPTR-ES-DR	
	vars.	cons.	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)	num.it.	time(s)
rococoC10-001000_t1	3117	1293	1	261	1	10	1	335	1	9	1	10
rococoC10-001000_t2	3117	1293	1	269	1	10	1	335	1	9	1	10
rococoC10-001000_t3	3117	1293	1	270	1	10	1	335	1	9	1	10
roi2alpha3n4_t1	6816	1251	20	>3600	154	>3600	149	>3600	146	1599	142	1455
roi2alpha3n4_t2	6816	1251	21	>3600	197	>3600	130	>3600	125	524	144	3294
roi2alpha3n4_t3	6816	1251	20	>3600	155	>3600	117	>3600	146	2830	148	3537
roll3000_t1	1166	2295	1	9	1	10	1	13	2	23	2	20
roll3000_t2	1166	2295	10	168	5	81	5	83	5	72	5	80
roll3000_t3	1166	2295	1	7	1	7	1	0	1	5	1	5
seymour_t1	1372	4944	0	>3600	443	>3600	10	>3600	1136	449	1426	431
seymour_t2	1372	4944	0	>3600	474	>3600	10	>3600	959	272	1036	253
seymour_t3	1372	4944	0	>3600	450	>3600	30	>3600	1194	614	1532	643
seymour1_t1	1372	4944	47	>3600	331	>3600	1679	800	1543	521	1663	454
seymour1_t2	1372	4944	48	>3600	396	>3600	956	495	970	263	1000	238
seymour1_t3	1372	4944	49	>3600	410	>3600	1669	920	1669	496	1875	498
sp150x300d_t1	600	450	4424	>3600	1414	133	248	35	248	36	205	24
sp150x300d_t2	600	450	2035	181	2035	181	222	26	222	29	228	19
sp150x300d_t3	600	450	1653	193	1606	154	199	18	199	19	212	27
splice1k1_t1	3253	6505	0	>3600	1	48	0	>3600	1	40	1	50
splice1k1_t2	3253	6505	0	>3600	1	47	0	>3600	1	49	1	41
splice1k1_t3	3253	6505	0	>3600	1	48	0	>3600	1	40	1	50
supportcase26_t1	436	870	38	729	36	37	16	3	16	3	16	3
supportcase26_t2	436	870	13	>3600	70	2018	29	1537	30	686	34	950
supportcase26_t3	436	870	9	487	14	8	3	0	3	0	3	0
wachplan_t1	3361	1553	0	1787	0	1845	0	1123	0	1087	0	1087
wachplan_t2	3361	1553	0	1804	0	1823	0	1122	0	1131	0	1135
wachplan_t3	3361	1553	0	1799	0	1823	0	1123	0	1089	0	1133