Inverse Mixed-Integer Programming: Polyhedral Approaches and Algorithms

Samuel Garvin[1], Fatemeh Nosrat[2], Andrew J. Schaefer[1]
[1]Rice University, [2]University of Florida

Problem Setup

Let $X \subseteq \mathbb{R}^n, x_0 \in X.$ The inverse-feasible region is defined as the set of all objective vectors for which this chosen point is optimal over the whole set:

$$\mathrm{IFR}(X, x_0) := \{d \in \mathbb{R}^n \mid d^\top x \leq d^\top x_0,\ \forall x \in X\}$$

We first focus on polyhedral sets, defining

$$P := \{x \in \mathbb{R}^n \mid Ax \leq b\}.$$

Geometry of the IFR

Let F_0 be the face of P such that $x_0 \in inner(F_0)$.

The feasible directions within this face define the subspace

$$V_0 := \mathrm{span}\{x - x_0 \mid x \in F_0\}$$

We show that the inverse-feasible region is restricted to its **orthogonal complement**, which also defines its dimension:

$$\mathrm{IFR}(P, x_0) \subseteq V_0^\perp \text{ and } \dim(\mathrm{IFR}(P, x_0)) = \dim(V_0^\perp).$$

Moreover, equality holds exactly when x_0 is an interior point:

$$\mathrm{IFR}(P, x_0) = V_0^\perp \iff x_0 \in \mathrm{inner}(P)$$

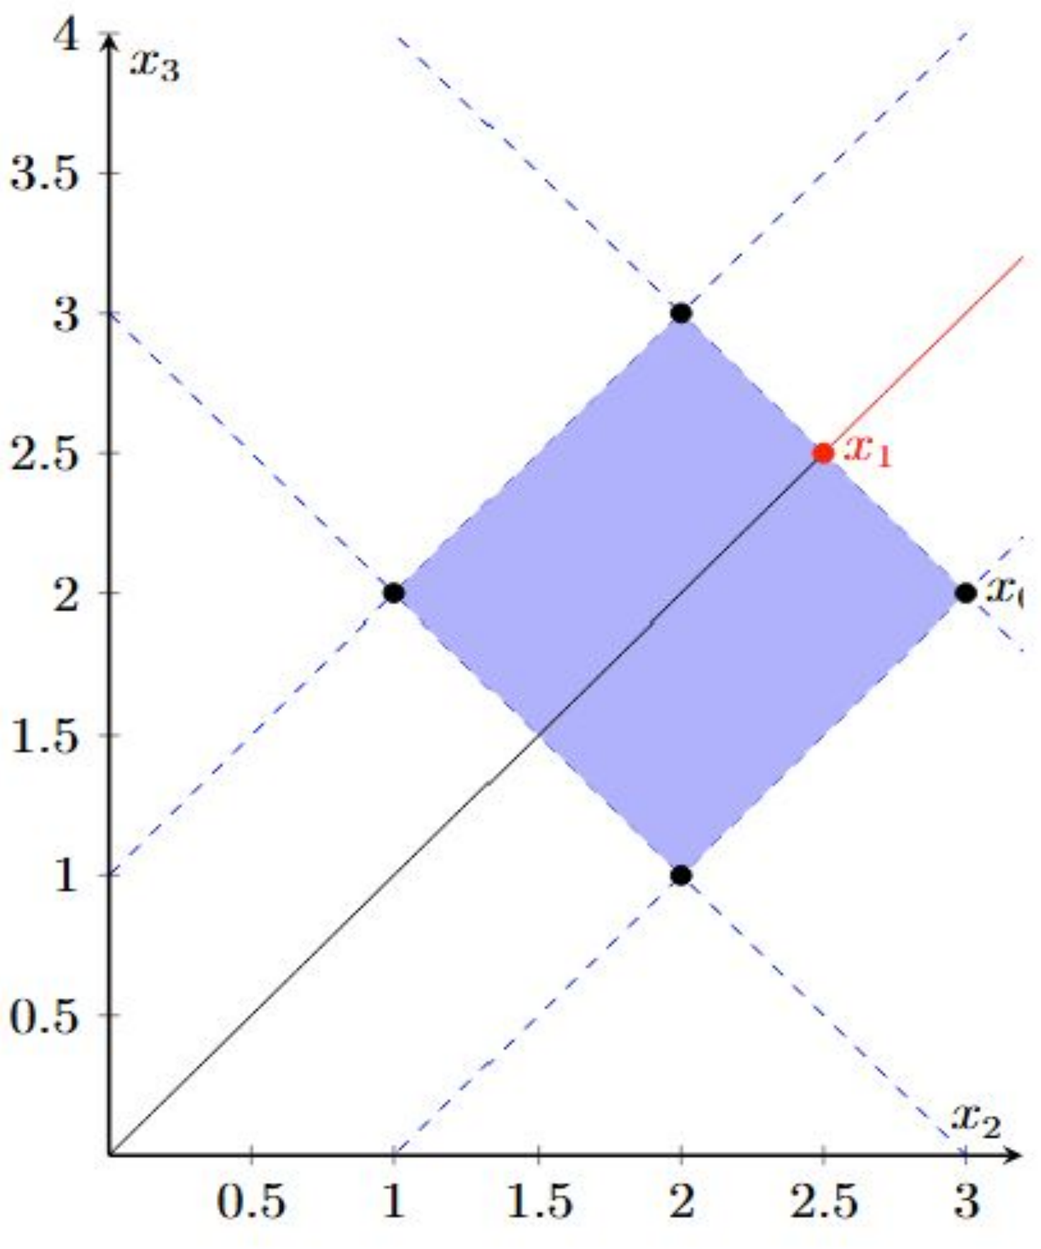

$P \subseteq \mathbb{R}^2$, with $\mathrm{IFR}(P, x_0)$ shown as a subset of $V_1^\perp$.

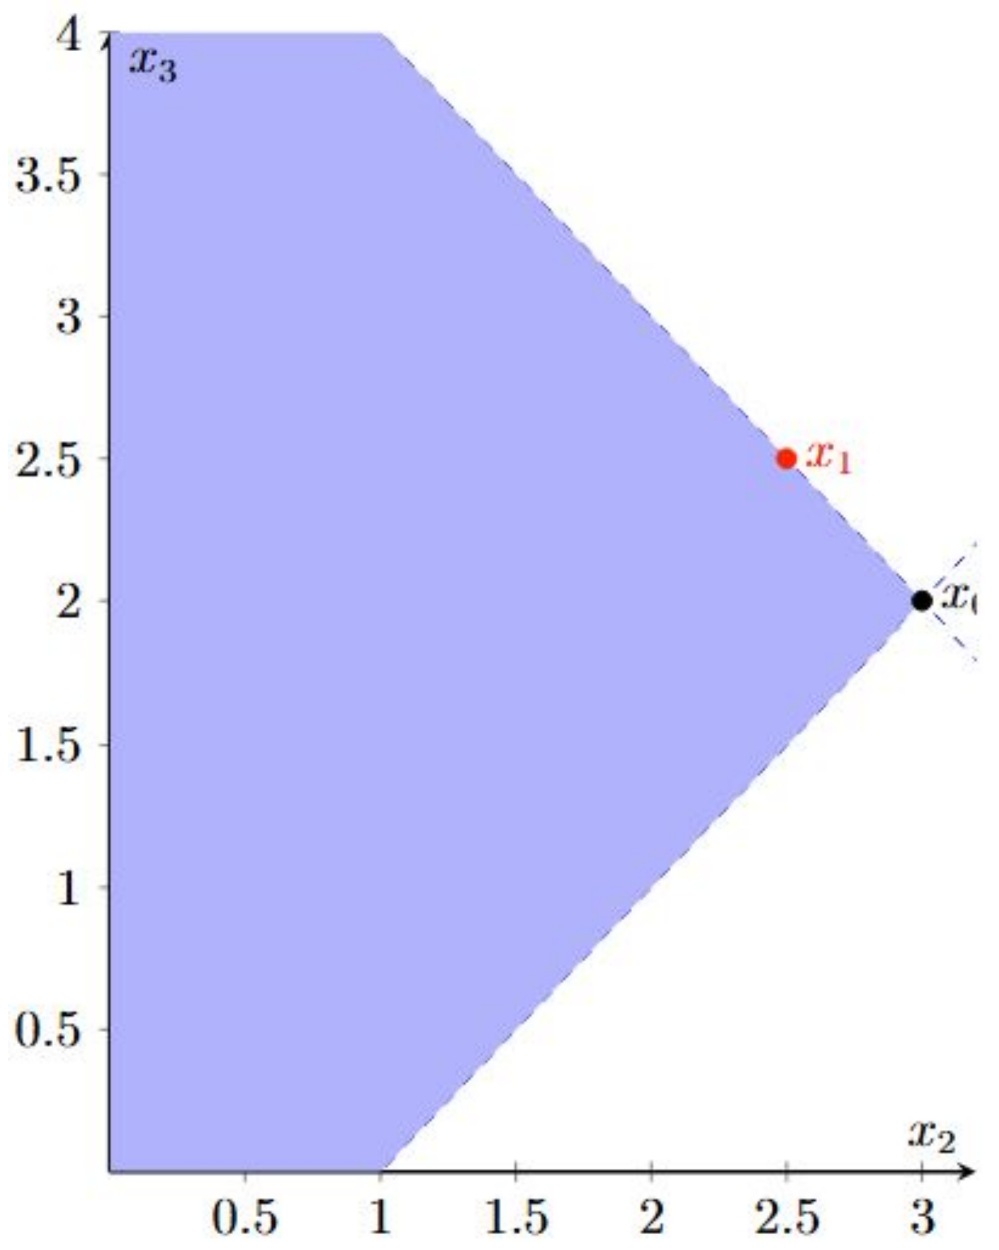

$P_{x_0}^=$, with only 2 active constraints.

Tavaslioglu et al. define $I^= = \{i | a_i x = b_i,\ \forall x \in P\},\ I_{x_0} = \{i | a_i x_0 = b_i\}$.

We construct a reduced system

$$P_{x_0}^= := \{x \in \mathbb{R}^n | \, a_i x = b_i, \forall i \in I^=,\ a_i x \leq b_i, \forall i \in I_{x_0} \backslash I^=\}$$

that keeps only structurally relevant constraints for inverse feasibility, showing that $\mathrm{IFR}(P, x_0) = \mathrm{IFR}(P_{x_0}^=, x_0)$

We also show that the **extreme rays** of the inverse-feasible region are exactly $\{a_i \mid i \in I^{x_0}\}$ when P is full-dimensional.

As an example, the inverse of the maximum weight matching problem on K_{20} may be formulated as an LP with 524,498 constraints, versus only 10, 440 in the reduced formulation.

Mixed-Integer Sets

We extend our results to the case where some variables are restricted to integer values. Let $S := P \cap (\mathbb{Z}^{n-l} \times \mathbb{R}^l)$. We consider the forward mixed-integer program

$$\mathcal{MIP} : \max\{c^\top x \mid x \in S\}$$

The inverse problem seeks to find the closest inverse-feasible vector to the "target":

$$\mathrm{Inv}(\mathcal{MIP}, x_0) := \min\{f(c - d) \mid d \in \mathrm{IFR}(S, x_0)\},$$

where f is **any** differentiable norm. Existing approaches rely on LP reformulations or iterative LP-based methods. While theoretically valid, LP reformulations often become super-exponentially large, while iterative approaches typically only guarantee feasibility at termination without giving any progress guarantees during runtime. These limitations motivate a method that provides structured, **incremental improvement.**

A Cutting-Plane Algorithm

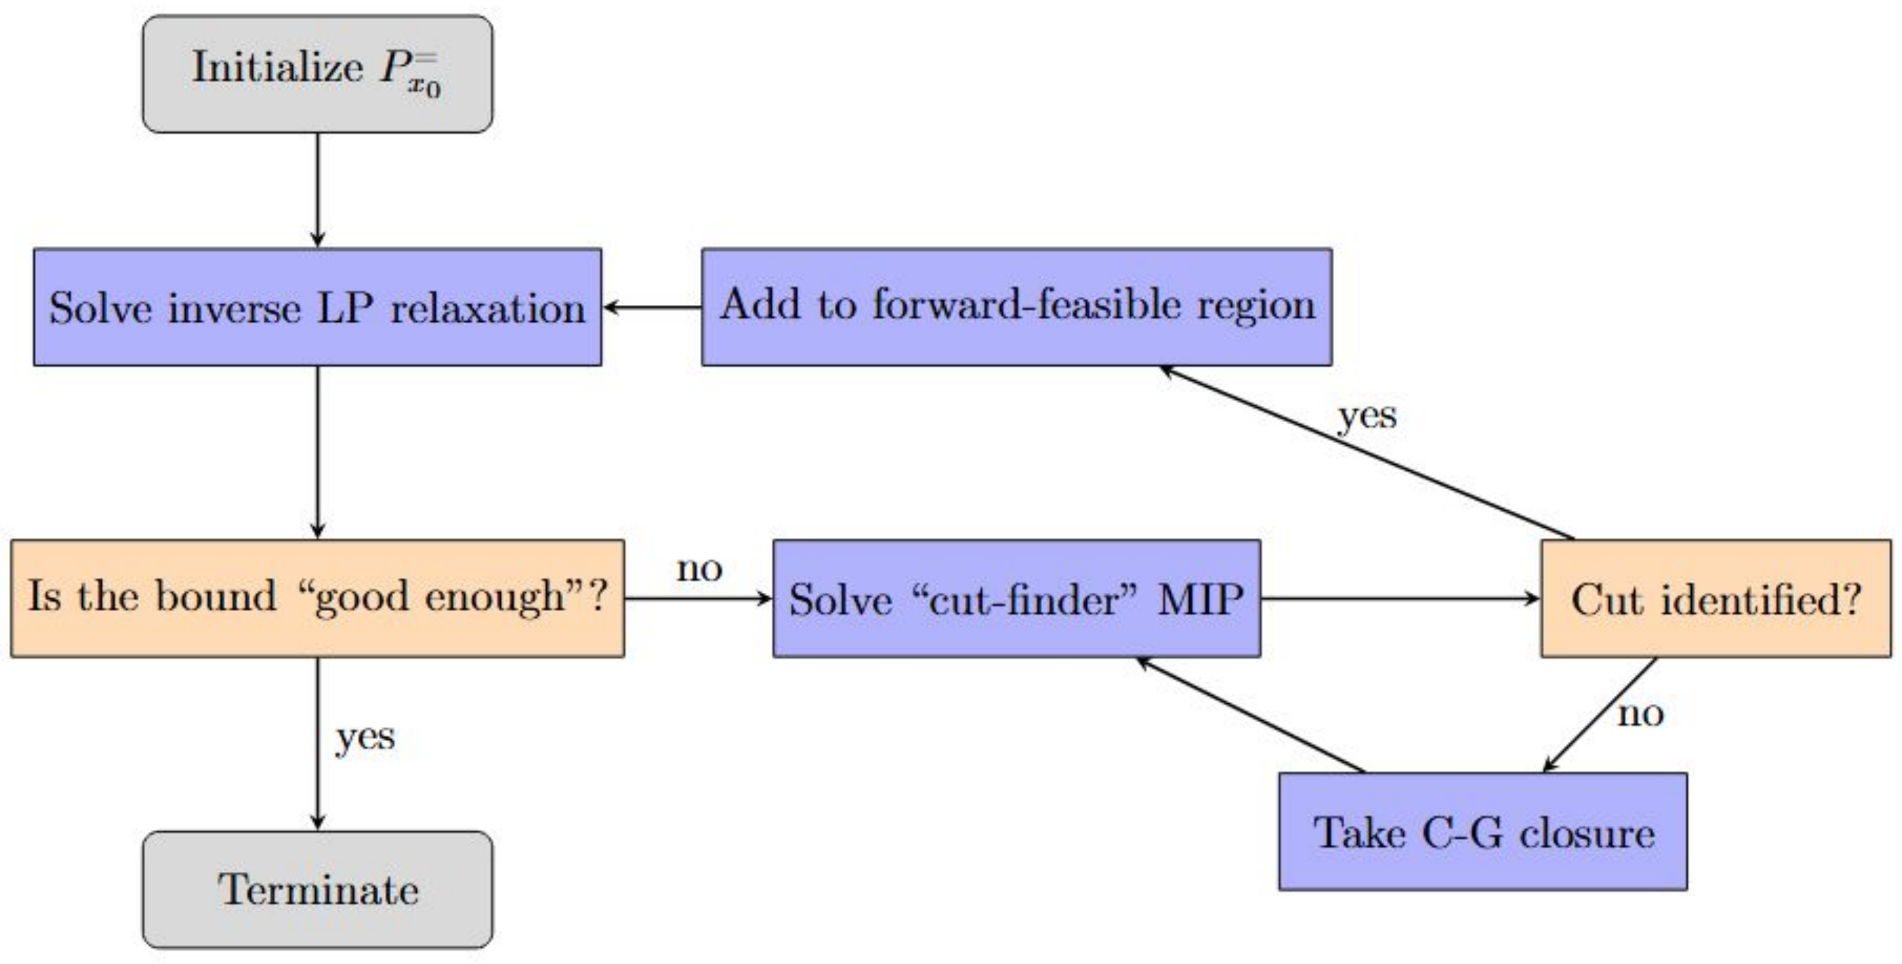
Initialize $P_{x_0}^=$

Solve inverse LP relaxation

Add to forward-feasible region

Is the bound "good enough"?

Solve "cut-finder" MIP

Cut identified?

Take C-G closure

Terminate

A flowchart for our cutting-plane algorithm.

We iteratively add cuts $\alpha^\top x \leq \alpha^\top x_0$ valid for $conv(S)$ to the forward problem. In general, we show that if

$P' := P_{x_0}^= \cap \{x \mid \alpha^\top x \leq \alpha^\top x_0\}$ has $P' \subsetneq P$, then

$\alpha \in \mathrm{IFR}(P', x_0) \setminus \mathrm{IFR}(P_{x_0}^=, x_0)$.

$$\begin{aligned}
\min \quad & \theta \\
\text{s.t.} \quad & (b_{(k)} - A_{(k)} x_0) y - \theta \leq 1, \\
& \nabla f(d_l)(A_{(k)}^T y) - \theta \leq 0, \\
& A^{(k)T} y + z = \mathbf{0}, \\
& 0 \leq y_i \leq 1,\ \forall i \in [m], \\
& z_j \in \mathbb{Z},\ \forall j \in [n], \\
& \theta \in \mathbb{R}
\end{aligned}$$

The MIP subproblem for cut generation.

Each iterative, improving cut $\alpha^\top x \leq \alpha^\top x_0$ in our algorithm comes from the optimal solution of the above MIP that "searches" over the current Chvatal-Gomory closure of the forward feasible region (encoded by $A_{(k)}, b_{(k)}$ and a current best solution d_l). Among our chief contributions is a proof that the MIP has a negative objective value **iff such a cut exists**. The resulting bound sequence is **strictly decreasing**, and finite termination is guaranteed.

Conclusions/Contributions

To our knowledge, this is the first algorithm for the inverse of MIP that is guaranteed to give feasible solutions during runtime **and** to give strictly better bounds. It is also the first work to characterize the IFR in terms of orthogonality and to suggest solution via a reduction of the forward LP. Taking the C-G closure is computationally expensive; future work should leverage repeated structure to improve on this.