

Solving sparse separable bilinear programs using lifted bilinear cover inequalities

Xiaoyi Gu*, Santanu S. Dey†, Jean-Philippe P. Richard‡

Abstract

Recently, in [12], we proposed a class of inequalities called *lifted bilinear cover inequalities*, which are second-order cone representable convex inequalities, and are valid for a set described by a separable bilinear constraint together with bounds on variables. In this paper, we study the computational potential of these inequalities for separable bilinear optimization problems. We first prove that the semi-definite programming relaxation provides no benefit over the McCormick relaxation for such problems. We then design a simple randomized separation heuristic for lifted bilinear cover inequalities. In our computational experiments, we separate many rounds of these inequalities starting from McCormick’s relaxation of instances where each constraint is a separable bilinear constraint set. We demonstrate that there is a significant improvement in the performance of a state-of-the-art global solver in terms of gap closed, when these inequalities are added at the root node compared to when they are not.

1 Introduction.

Lifting is a technique used to generate cutting planes for a set from a seed inequality valid for a restriction of this set. Lifting was first studied in the context of mixed integer linear programming (MILP); see [20] for a review. Cover inequalities ([24, 19, 1, 2, 15]) are valid for 0-1 knapsack sets with coefficients satisfying the minimal cover property. The inequalities obtained after lifting cover inequalities are called *lifted cover inequalities*; see [15, 28, 25, 26, 13, 14]. They form a family of strong valid inequalities for general 0-1 knapsack sets that is very important for state-of-the-art MILP solvers; see [3].

Inspired by the success of these inequalities we introduced in [12] a class of inequalities we call *lifted bilinear cover inequalities* for separable bilinear constraints. These second-order cone representable (SOCR) convex inequalities are derived using lifting and are valid for a set described by a separable bilinear constraint together with bounds on variables. In Section 1.1 below, we give a detailed description of lifted bilinear cover inequalities and the separable bilinear sets they are valid for.

In this paper, we study the computational potential of these inequalities. We are inspired by the classical paper of [7]. This paper is one of the first to highlight the computational importance of lifted cover inequalities in MILP. Specifically, this paper considered sparse instances – since most practical instances are sparse. Here, sparsity means that the support of each constraint is significantly smaller than the number of variables in the entire problem; see [8, 9] for discussions and

*xiaoyigu@gatech.edu, H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332.

†santanu.dey@isye.gatech.edu, H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332.

‡jrichar@umn.edu, Department of Industrial and Systems Engineering, University of Minnesota.

results regarding sparsity of problems and cutting-planes. Similarly, we generate sparse separable bilinear instances to test lifted bilinear cover inequalities. We numerically show that a significant improvement in the performance (in terms of gap closed) of a global solver can be observed when these inequalities are added at the root node compared to when they are not.

1.1 Lifted bilinear cover inequalities for separable bilinear programs.

For a positive integer n , we use the notation $[n]$ to describe the set $\{1, \dots, n\}$.

In [12], we derive inequalities that can be applied to improve convex relaxations of the feasible region of separable bilinear programs, which we call inequalities *lifted bilinear cover inequalities*.

Definition 1. A separable bilinear program is an optimization problem of the form

$$\begin{aligned} \min \quad & \sum_{i \in [n]} (c_i^x x_i + c_i^y y_i) \\ \text{s.t.} \quad & \sum_{i \in [n]} a_i^j x_i y_i \geq d_j, \quad \forall j \in [m], \\ & x_i, y_i \in [0, 1], \quad \forall i \in [n]. \end{aligned} \tag{1}$$

Each bilinear constraint in (1) is in the form of a *separable bilinear set*, defined as follows.

Definition 2. A set S is called a separable bilinear set if it is of the form

$$S := \left\{ (x, y) \in [0, 1]^n \times [0, 1]^n \mid \sum_{i \in [n]} a_i x_i y_i \geq d \right\},$$

where $d \in \mathbb{R}$ and $a_i \in \mathbb{R}$ for all $i \in [n]$.

For each $i \in [n]$, variables x_i and y_i in S appear in only one term in the left-hand-side. The convex hull of the set S has been studied in [11, 21] and the convex hull of the relaxation of S obtained by dropping the upper bounds on the variables is presented in [23] for the case where all coefficients a_i are non-negative. These convex hull results (especially the exact ones presented in [11] and [21]) have limited computational usefulness, since the description of the convex hull is exponential in the number of variables. This motivates the derivation of families of cutting planes for S .

In [12] we proposed using *minimal covers* to generate seed inequalities we call *bilinear cover inequalities* and then perform lifting to obtain valid inequalities for separable bilinear sets. We next briefly introduce the main results of [12]. We start by introducing the notions of *minimal cover* and *minimal cover yielding partition*.

Definition 3. A set $\{a_i \in \mathbb{R} \mid i \in [k]\}$, with $k \in \mathbb{Z}_+$ being a positive integer, is said to form a *minimal cover* of $d \in \mathbb{R}$ if

- (i) $a_i > 0$ for all $i \in [k]$, $d > 0$,
- (ii) $\sum_{i \in [k]} a_i > d$,
- (iii) $\sum_{i \in K} a_i \leq d$, $\forall K \subsetneq [k]$.

For a separable bilinear set S , we say that a partition $\Lambda = \{I, J_0, J_1\}$ of $[n]$, where $I \neq \emptyset$, is a *minimal cover yielding partition* if: $\{a_i \mid i \in I\}$ forms a minimal cover of $d^\Lambda := d - \sum_{i \in J_1} a_i$. For a minimal cover yielding partition, we let $J_0^+ := \{i \in J_0 \mid a_i > 0\}$ and $J_0^- := \{i \in J_0 \mid a_i < 0\}$. We define J_1^+ and J_1^- similarly.

Remark 1. When $k \geq 2$, conditions (ii) and (iii) in the definition of minimal cover imply condition (i). For example, if $a_i \leq 0$ for some $i \in [k]$, then (ii) implies $\sum_{j \in [k] \setminus \{i\}} a_j > d$, contradicting (iii). Now (iii) together with $a_i > 0$ for $i \in [k]$ implies that $d > 0$.

Notation 1. Assuming that $\{a_i \mid i \in [k]\}$ forms a minimal cover of d^Λ , we use

1. $\Delta := \sum_{i \in [k]} a_i - d^\Lambda$,
2. $d_i := d^\Lambda - \sum_{j \in [k] \setminus \{i\}} a_j = a_i - \Delta$,
3. $I^> := \{i \in [k] \mid a_i > \Delta\}$,
4. when $I^> \neq \emptyset$, let i_0 denote any index in $I^>$ such that $a_{i_0} = \min\{a_i \mid i \in I^>\}$. We say that i_0 does not exist when $I^> = \emptyset$.

The process to generate lifted bilinear cover inequalities for S from the minimal cover yielding partition is to (i) first fix $x_i = y_i = 0$ for $i \in J_0$ and $x_i = y_i = 1$ for $i \in J_1$, (ii) generate valid seed inequality for the restricted region, and (iii) finally lift the seed inequality to obtain a valid cut for S .

The following result from [12] establishes the existence of minimal cover yielding partitions in many instances of S .

Theorem 1. For a nonempty separable bilinear set S , either there exists at least one minimal cover yielding partition or $\text{conv}(S)$ is polyhedral.

The seed inequality used for lifting is presented in Theorem 2, which also discusses its strength.

Theorem 2. For a separable bilinear set S as in Definition 2 where $\{a_i \mid i \in [n]\}$ forms a minimal cover of d , the following bilinear cover inequality is valid:

$$\sum_{i \in [n]} \frac{\sqrt{a_i}}{\sqrt{a_i} - \sqrt{d_i}} (\sqrt{x_i y_i} - 1) \geq -1. \quad (2)$$

Further, the set $R := \{(x, y) \in \mathbb{R}_+^{2n} \mid (2)\}$ is such that $(4 \cdot R) \cap [0, 1]^{2n} \subseteq \text{conv}(S) \subseteq R \cap [0, 1]^{2n}$.

With the seed bilinear cover inequality, the following lifted bilinear cover inequality was proposed, which is valid for the whole separable bilinear set.

Theorem 3. Consider a separable bilinear set S as in Definition 2. Let $\Lambda = \{I, J_0, J_1\}$ be a minimal cover yielding partition, let $\Delta, a_{i_0}, d_i, l_+, l_-$ be defined as in Notation 1, and let $J_0^+, J_0^-, J_1^+, J_1^-$ be as in Definition 3. Then inequality

$$\sum_{i \in I} \frac{\sqrt{a_i}}{\sqrt{a_i} - \sqrt{d_i}} (\sqrt{x_i y_i} - 1) + \sum_{i \notin I} \gamma_i(x_i, y_i) \geq -1, \quad (3)$$

is valid for S where $\gamma_i : \mathbb{R}^2 \rightarrow \mathbb{R}$ for $i \in [n] \setminus I$ are the concave functions:

- (i) $\gamma_i(x, y) = l_+ a_i \min\{x, y\}$ for $i \in J_0^+$;
- (ii) $\gamma_i(x, y) = -l_+ a_i \min\{2 - x - y, 1\}$ for $i \in J_1^-$;
- (iii) $\gamma_i(x, y) = \min\{l_- a_i (x + y - 1), l_+ a_i (x + y - 1) + l_+ \Delta - 1, 0\}$ for $i \in J_0^-$;

(iv) $\gamma_i(x, y) = \min\{\tilde{g}_i(x, y), \tilde{h}_i(x, y), g_i(x, y), h_i(x, y)\}$, for $i \in J_1^+$ with $a_i \geq a_{i_0}$ when $I^> \neq \emptyset$, and $\gamma_i(x, y) = \min\{\tilde{g}_i(x, y), \tilde{h}_i(x, y)\}$ in all other cases where $i \in J_1^+$, with

$$\begin{aligned}\tilde{g}_i(x, y) &= l_+ a_i (\min\{x, y\} - 1) + l_+ \Delta - 1 \\ \tilde{h}_i(x, y) &= l_- a_i (\min\{x, y\} - 1) \\ g_i(x, y) &= \sqrt{a_i - \Delta} \sqrt{a_i} l_+ \sqrt{xy} - l_+ (a_i - \Delta) - 1 \\ h_i(x, y) &= \frac{\sqrt{a_i}}{\sqrt{a_i} - \sqrt{d_i}} (\sqrt{xy} - 1)\end{aligned}$$

with $l_- = \frac{1}{\Delta}$ and $l_+ = \frac{\sqrt{a_{i_0}} + \sqrt{d_{i_0}}}{\Delta \sqrt{d_{i_0}}}$ if i_0 exists and $l_+ = \frac{1}{\Delta}$ otherwise.

The lifted bilinear cover inequality (3) is convex and second-order cone representable (SOCR), making it easy to incorporate in relaxations, given the enormous progress in SOC solvers.

1.2 Main contributions.

While the lifted bilinear cover inequality was derived in [12], its computational usefulness has not been evaluated. In this paper, we utilize lifted bilinear cover inequalities in a computational study and illustrate their benefit.

We expand on our main contributions next:

- Instead of using lifted bilinear cover inequalities, we could consider using the semidefinite programming (SDP) relaxation for strengthening the McCormick relaxation. In general, it is known that combining the McCormick relaxation with the SDP relaxation produces good bounds; see [5] for instance. We show that, surprisingly, for the class of separable instances, the SDP relaxation gives a bound equal to that of the McCormick relaxation. We shall see in the later sections that the lifted bilinear cover inequalities are able to close significant root gap over the McCormick relaxation – thus showing that lifted bilinear cover inequalities are important in solving separable bilinear instances where SDP relaxations are of limited use.
- We design a simple randomized separation heuristic for lifted bilinear cover inequalities. We use this separation heuristic to add many rounds of cuts starting from a natural relaxation of the problem that uses McCormick inequalities. We then solve the instances using a commercial global solver with and without these inequalities added to the root node and compare the results obtained.
- We discover that the inequalities separated using the simple heuristic provide a major performance boost to a commercial global solver in terms of overall gap closed on sparse test instances.

The rest of the paper is organized as follows. In Section 2, we compare the SDP and McCormick relaxations for separable programs. In Section 3, we present our separation heuristic. In Section 4, we describe the procedure we use to generate synthetic test instances and discuss various other parameters regarding the use of the separation heuristic on these instances. In Section 5, we present detailed results of our experiments. We give concluding remarks in Section 6.

2 Comparing SDP and McCormick relaxations of separable programs.

We show that, for separable bilinear programs, the relaxation obtained by adding the traditional SDP constraint to the McCormick relaxation gives the same bound as that given by just the McCormick relaxation. In fact, we prove this result in the following slightly more general setting. Consider the following quadratically constrained quadratic program (QCQP):

$$\min \sum_{i \in [n]} a_i^0 x_i y_i + \sum_{i \in [n]} c_i^{x,0} x_i + \sum_{i \in [n]} c_i^{y,0} y_i \quad (4a)$$

$$\text{s.t.} \quad \sum_{i \in [n]} a_i^j x_i y_i + \sum_{i \in [n]} c_i^{x,j} x_i + \sum_{i \in [n]} c_i^{y,j} y_i \geq b_j, \quad \forall j \in [m], \quad (4b)$$

$$x, y \in [0, 1]^n. \quad (4c)$$

We consider next two relaxations of (4). In both of these relaxations, we introduce variables $w_{i,j}$ to represent the products $x_i x_j$, variables $w_{i,n+j}$ and $w_{n+i,j}$ to represent the products $x_i y_j$ and $y_i x_j$, respectively, and variables $w_{n+i,n+j}$ to represent the products $y_i y_j$. Clearly, only variables $w_{i,n+i}$ are needed to relax (4b). The first relaxation, which we call *McCormick relaxation*, is obtained by using McCormick inequalities to approximate the relationships between variables w , x , and y :

$$z_{Mc} := \min \sum_{i \in [n]} a_i^0 w_{i,n+i} + \sum_{i \in [n]} c_i^{x,0} x_i + \sum_{i \in [n]} c_i^{y,0} y_i \quad (5a)$$

$$\text{s.t.} \quad \sum_{i \in [n]} a_i^j w_{i,n+i} + \sum_{i \in [n]} c_i^{x,j} x_i + \sum_{i \in [n]} c_i^{y,j} y_i \geq b_j, \quad \forall j \in [m], \quad (5b)$$

$$x_i = u_i, \quad \forall i \in [n], \quad (5c)$$

$$y_i = u_{n+i}, \quad \forall i \in [n], \quad (5d)$$

$$u \in [0, 1]^{2n}, \quad (5e)$$

$$\max\{0, u_i + u_k - 1\} \leq w_{i,k} \leq \min\{u_i, u_k\}, \quad \forall (i, k) \in [2n] \times [2n]. \quad (5f)$$

The second relaxation, which we call *McCormick+SDP relaxation*, is obtained from the McCormick relaxation by including the traditional SDP relaxation of the property that W is a rank-1 matrix:

$$z_{MS} := \min \sum_{i \in [n]} a_i^0 w_{i,n+i} + \sum_{i \in [n]} c_i^{x,0} x_i + \sum_{i \in [n]} c_i^{y,0} y_i \quad (6a)$$

$$\text{s.t.} \quad (5b), (5c), (5d), (5e), (5f) \quad (6b)$$

$$W := \begin{bmatrix} 1 & u^\top \\ u & w \end{bmatrix} \succeq 0, \quad (6c)$$

where $W \succeq 0$ denotes that W is positive semi-definite.

Proposition 1. *For the optimization problem (4), $z_{Mc} = z_{MS}$, or both (5) and (6) are infeasible.*

Proof. Note that both (5) and (6) are bounded, as all variables are bounded by $[0, 1]$. Therefore, if either problem is feasible, then the feasible region is compact and thus a corresponding optimal solution must exist (objective function is linear). Since the feasible region of (6) is contained in that of (5), we only need to prove that if (5) is feasible, then (6) is also feasible and that $z_{Mc} \leq z_{MS}$. Let (x^*, y^*, u^*, w^*) be an optimal solution corresponding to (5).

Then, construct the matrix $W^* \in \mathbb{R}^{(2n+1) \times (2n+1)}$ as follows:

$$W_{ik}^* = \begin{cases} 1, & \text{if } i = k = 1, \\ u_{i-1}^*, & \text{if } k = 1, i \geq 2, \\ u_{k-1}^*, & \text{if } i = 1, k \geq 2, \\ u_{i-1}^*, & \text{if } i = k > 1, \\ u_{i-1}^* u_{k-1}^*, & \text{if } i \neq k, i > 1, k > 1, |i - k| \neq n, \\ w_{\min\{i-1, k-1\}, \max\{i-1, k-1\}}^*, & \text{if } i \neq k, i > 1, k > 1, |i - k| = n. \end{cases}$$

It is easy to verify that (x^*, y^*, u^*, W^*) satisfy all the linear constraints corresponding to z_{MS} . It remains to show that $W^* \succeq 0$. This would show that there exists a feasible solution to (6) with objective function value equal to that of z_{Mc} , completing the proof.

Construct the matrix $G \in \mathbb{R}^{(2n+1) \times (2n+1)}$ as follows:

$$G_{ik} = \begin{cases} 1, & \text{if } i = k = 1, \\ u_{i-1}^*, & \text{if } k = 1, \\ u_{k-1}^*, & \text{if } i = 1, \\ u_{i-1}^* u_{k-1}^*, & \text{if } j \neq k, i > 1, k > 1. \end{cases}$$

Clearly, $G \succeq 0$ since $G = \begin{bmatrix} 1 \\ u^* \end{bmatrix} [1 \quad (u^*)^\top]$. So it is sufficient to show that $W^* - G \succeq 0$. We compute that

$$(W^* - G)_{ik} = \begin{cases} 0, & \text{if } i = 1 \text{ or } k = 1, \\ u_{i-1}^* - (u_{i-1}^*)^2, & \text{if } i = k > 1, \\ 0, & \text{if } i \neq k, i > 1, k > 1, |i - k| \neq n, \\ w_{\min\{i-1, k-1\}, \max\{i-1, k-1\}}^* - u_{i-1}^* u_{k-1}^*, & \text{if } i \neq k, i > 1, k > 1, |i - k| = n. \end{cases}$$

In words, the first row and column are all zero. For each of the remaining rows, there is exactly one non-zero off-diagonal term. Therefore (without loss of generality) by the Gershgorin circle theorem (see [16]) it is sufficient to prove that

$$u_{i-1}^* - (u_{i-1}^*)^2 \geq |w_{i-1, n+i-1}^* - u_{i-1}^* u_{i+n-1}^*|, \quad \forall i \in \{2, \dots, n+1\},$$

that is, it is sufficient to prove that

$$x_i^* - (x_i^*)^2 \geq |w_{i, n+i}^* - x_i^* y_i^*|, \quad \forall i \in [n]$$

where $\min\{x_i^*, y_i^*\} \geq w_{i, n+i}^* \geq \max\{x_i^* + y_i^* - 1, 0\}$ for $i \in [n]$. This can be done by considering the four possible extreme values of w^* :

- **Case 1:** $w_{i, n+i}^* = 0$. In this case, $x_i^* + y_i^* \leq 1$. Then

$$1 - x_i^* \geq y_i^* \quad \Rightarrow \quad x_i^*(1 - x_i^*) \geq x_i^* y_i^* \quad \Rightarrow \quad x_i^* - (x_i^*)^2 \geq x_i^* y_i^* = |w_{i, n+i}^* - x_i^* y_i^*|.$$

- **Case 2:** $w_{i, n+i}^* = x_i^* + y_i^* - 1 \geq 0$. In this case, $x_i^* + y_i^* \geq 1$. First, observe that $w_{i, n+i}^* - x_i^* y_i^* = x_i^* + y_i^* - 1 - x_i^* y_i^* = (-1 + y_i^*)(1 - x_i^*) \leq 0$. Thus, $|w_{i, n+i}^* - x_i^* y_i^*| = -(-1 + y_i^*)(1 - x_i^*)$. Then,

$$\begin{aligned} x_i^* \geq 1 - y_i^* & \Rightarrow x_i^*(1 - x_i^*) \geq -(-1 + y_i^*)(1 - x_i^*) \\ & \Rightarrow x_i^* - (x_i^*)^2 \geq -(-1 + y_i^*)(1 - x_i^*) = |w_{i, n+i}^* - x_i^* y_i^*|. \end{aligned}$$

- **Case 3:** $w_{i,n+i}^* = x_i^*$. In this case, $x_i^* \leq y_i^*$. Then,

$$x_i^* \leq y_i^* \quad \Rightarrow \quad (x_i^*)^2 \leq x_i^* y_i^* \quad \Rightarrow \quad x_i^* - (x_i^*)^2 \geq x_i^* - x_i^* y_i^* = |w_{i,n+i}^* - x_i^* y_i^*|.$$

- **Case 4:** $w_{i,n+i}^* = y_i^*$. In this case, $x_i^* \geq y_i^*$. Then,

$$x_i^* \geq y_i^* \quad \Rightarrow \quad x_i^*(1 - x_i^*) \geq y_i^*(1 - x_i^*) \quad \Rightarrow \quad x_i^* - (x_i^*)^2 \geq y_i^* - x_i^* y_i^* = |w_{i,n+i}^* - x_i^* y_i^*|.$$

□

3 Improved relaxation through a heuristic separation algorithm.

In this section, we describe a procedure that, starting from the McCormick relaxation of the problem, produces an improved relaxation by adding a selection of lifted bilinear cover inequalities. This procedure, whose details are presented in Algorithm 1, solves a convex relaxation of the problem at each step and separates the obtained optimal relaxation solution using a heuristic separation procedure. The steps are repeated until either a preset time limit is reached, a preset iteration limit is reached, or the improvement in relaxation bounds between iterations becomes too small.

Algorithm 1 Improving Root Relaxation with Lifted Bilinear Cover Cuts

Require: bilinear programming problem; non-convex nonlinear solver `globalsolver`; convex second-order cone programming (SOCP) solver `convxsolver`;

parameters:

minimum improvement threshold ϵ_z , iteration number limit T , heuristic time limit T_{heu} , parameters for heuristic and solver

- 1: generate McCormick relaxation (5)
 - 2: solve (5) using `convxsolver`;
obtain optimal solution (x^0, y^0) and lower bound z^0 ; set $t = 1$
 - 3: **while** (heuristic time limit T_{heu} not reached) **do**
 - 4: generate bilinear cover cuts using Algorithm 2 and add them to the problem
 - 5: **if** (violated bilinear cover cuts were not generated) **break**
 - 6: solve the improved relaxation with `convxsolver`;
obtain optimal solution (x^t, y^t) and lower bound z^t
 - 7: **if** $|z^t - z^{t-1}|/|z^{t-1}| < \epsilon_z$ (improvement is too small), **break**
 - 8: $t \leftarrow t + 1$
 - 9: **if** $t > T$ (iteration limit is reached) **break**
 - 10: **end while**
 - 11: solve the problem using `globalsolver`
-

The workhorse of the procedure described above is a heuristic separation algorithm that, given a relaxation solution, seeks to generate a violated lifted bilinear cover cut from each of the bilinear constraints of the problem. The heuristic, which we explain next, is formally presented in Algorithm 2.

The simple randomized heuristic that we implement can be intuitively described as “guess a reasonable partition (I, J_0, J_1) ” and “randomly adjust it if it does not yield a violated cut”. Given a solution to separate, it considers each constraint j of the bilinear problem in sequence. For separating from a given row j , the heuristic must decide for each variable, whether it is in the set $I, J_0,$

Algorithm 2 Heuristic Separation Algorithm for Lifted Bilinear Cover Cuts

Require: bilinear programming problem with relaxation solution (x^t, y^t) ;
 parameters:
 approximation threshold ϵ , search iteration limit S_j for $j \in [m]$

- 1: **for** $j = 1 : m$ **do**
- 2: **if** heuristic time limit T_{heu} reached, **break**
- 3: **if** $\sum_{i=1}^n a_i^j x_i y_i - d_j \geq 0$ (bilinear constraint j not violated), **continue**
- 4: **for** $i = 1 : n$ **do**
- 5: **if** $a_i^j = 0$, set label l_i^j as inactive ((x_i, y_i) is not in the bilinear cover cut for row j)
- 6: **else if** $x_i y_i < \epsilon$, set l_i^j as J_0
- 7: **else if** $x_i y_i > 1 - \epsilon$, set l_i^j as J_1
- 8: **else if** $a_i^j > 0$, set l_i^j as I
- 9: **else** (*i.e.*, $a_i^j < 0$), set l_i^j as J_1 or J_0 with probability $x_i y_i$ and $1 - x_i y_i$ respectively
- 10: **end for**
- 11: **if** partition Λ with labels l is not a minimal cover yielding partition for row j **then**
- 12: **for** $s = 1 : S_j$ **do**
- 13: (check reason for failure)
- 14: **if** $d_j^\Lambda = d_j - \sum_{i \in J_1} a_i^j \leq 0$ (right-hand-side is not positive):
 select randomly $i^* \in \{i \mid a_i^j > 0, l_i^j = J_1\} \cup \{i \mid a_i^j < 0, l_i^j = J_0\}$ and flip $l_{i^*}^j$ from J_1 to I
 (in the first case) or from J_0 to J_1 (in the second case); **break** if not found
- 15: **else if** $\Delta_j = \sum_I a_i^j - d_j^\Lambda \leq 0$ (Δ_j is not positive, *i.e.*, not a cover):
 select randomly $i^* \in \{i \mid a_i^j > 0, l_i^j = J_0\} \cup \{i \mid a_i^j < 0, l_i^j = J_1\}$ and flip $l_{i^*}^j$ from J_0 to J_1
 or from J_1 to J_0 ; **break** if not found
- 16: **else** ($\exists i \in I, a_i^j < \Delta_j$, *i.e.*, not a minimal cover):
 find $i^* \in \operatorname{argmin}_I a_i^j$ and set $l_{i^*}^j$ as J_1
- 17: **if** partition Λ with adjusted labels l is a minimal cover yielding partition, **break**
- 18: **end for**
- 19: **end if**
- 20: **if** lifted bilinear cover cut is generated and violated by (x^t, y^t) :
 add the lifted bilinear cover cut to the problem
- 21: **end for**

or J_1 . The resulting cut is entirely determined based on this choice. Formally, the following steps are performed:

- The heuristic starts with an intuitive guess (steps 4-10), where the label l_i^j for variable i is assigned to J_0 if $x_i y_i$ is close to 0 and to J_1 if $x_i y_i$ close to 1. For all other indices i , l_i^j is set to I if $a_i^j > 0$, and randomly to J_0 or J_1 otherwise.
- After generating the initial guess, if (I, J_0, J_1) fails to be a minimal cover yielding partition, the heuristic performs several attempts at randomly adjusting the partition (steps 12-19), where labels l_i^j are flipped based on different cases of failure (steps 14, 15, 16 respectively). We allow the number S_j of attempts taken to depend on the row j , as we expect separation to be harder to perform in denser rows.
- If a partition (I, J_0, J_1) is found that corresponds to a violated lifted cover, the associated inequality is added, and we proceed to attempting to generate a cut from the following problem row.

Algorithm 2 is a basic guess-and-adjust heuristic, where the adjust step shares similarity to the WALK SAT approach of [22]. Nevertheless, our experiments presented in the later sections will illustrate the power of this basic heuristic, thereby also demonstrating the potential of the lifted bilinear cover cuts.

4 Experimental setup.

In this section, we present the way we setup our experiments, including instance generation, environment, and various other parameters.

4.1 Randomly generated instances

We consider instances of the form

$$\begin{aligned} \min \quad & \sum_{i \in [n]} (c_i^x x_i + c_i^y y_i) \\ \text{s.t.} \quad & \sum_{i \in [n]} a_i^j x_i y_i \geq d_j, \quad \forall j \in [m], \\ & x_i, y_i \in [0, 1], \quad \forall i \in [n]. \end{aligned}$$

Inspired by the classical paper [7] and by the fact most instances in practice are sparse, we generate sparse instances of the above model.

We generate test instances based on three parameters: the number of rows (*i.e.*, bilinear constraints) m , the number of variables n , and the density parameter p . Moreover, we generate two groups of instances, with the first group having non-negative coefficients a_i^j (which we call *non-negative instances*), and the second group having both non-negative and non-positive coefficients a_i^j (which we call *mixed-signs instances*).

After fixing the three parameters m , n , and p , we randomly generate the objective coefficients c , the bilinear coefficients a and the right-hand-sides d . The coefficients c_i^x and c_i^y are generated independently from a uniform distribution on $[0, 1]$; a_i^j is set to be non-zero with probability p , and when non-zero, it is generated from a uniform distribution on $[0, 1]$ for non-negative instances, or from a uniform distribution on $[-1, 1]$ for mixed-signs instances. For d_j , we let $s_j := \sum_{i \in [n]} a_i^j$ and

we set $d_j = r_j s_j$, where r_j is drawn from a uniform distribution on $[0, 1]$ when $s_j > 0$ and r_j is drawn from a uniform distribution on $[1, 2]$ otherwise.

For our experimental design, we chose $m, n \in [100, 250, 500]$ and $p \in [0.01, 0.02, 0.05]$. To make the problem “reasonably” sparse (*i.e.*, neither too sparse nor too dense), we only select those combinations where $np \in [5, 20]$ where np can be understood as the expected number of non-zero terms for each constraint. There is a total of 15 different settings of m , n , and p that satisfy this requirement. For each setting, we randomly generate ten different instances with non-negative coefficients and ten instances with mixed signs.

4.2 Environment.

The experiments were implemented in Python 3.9.7 with Gurobi 9.5.0 for both convex SOCP solver `convexsolver` and non-convex nonlinear solver `globalsolver`, with parameter `NonConvex=2` for the latter. All experiments were run in parallel on the high-performance computational cluster of ISyE, Georgia Tech, which contains roughly 2,340 cores of x86-64 processing with over 28.9TB of memory spread across the systems. Each task was processed on a dedicated core with 8GB of memory.

4.3 Testing.

We compare two settings for evaluating the efficacy of the lifted bilinear cover inequalities:

- We separate the lifted bilinear cover inequalities as described in the previous section. Then we give as input to Gurobi the separable bilinear optimization problem augmented with these separated cuts.
- We directly input the separable bilinear optimization problem to Gurobi.

In both the above settings, we use the default parameters of Gurobi.

For non-negative instances, we also add a comparison to the valid inequality described in [23]:

$$\sum_{i \in [n]} \sqrt{a_i^j x_i y_i} \geq \sqrt{d_j}, \quad (7)$$

for the j -th row, where the a_i^j s are non-negative. Together with nonnegativity constraints, inequality (7) gives the convex hull of the j -th row if there is no upper bound on the variables. For testing these cuts, we add them to the problem (effectively bolstering the McCormick relaxation at the root node) and then solve the resulting problem with default Gurobi.

We note that there is no obvious way to use the cuts (7) for mixed-signs instances, as there is no easy way to complement the variables in mixed-signs instances so as to produce equivalent bilinear instances with non-negative coefficients.

In Algorithm 1, we set both the heuristic time limit and the time limit for Gurobi non-linear non-convex solver to 1,800 seconds (half an hour). While it might seem that our algorithm has an unfair advantage with extra time for heuristic, we will see later that the difference is negligible as the time spent on the separation heuristic is significantly smaller compared to the time for solving the non-convex problem.

Moreover, for our heuristic, we set $\epsilon_z = 5 \times 10^{-3}$, $\epsilon = 10^{-2}$, $T = 10 \cdot np$ and $S_j = 10 \cdot \#(a_i^j \mid a_i^j \neq 0)$, where $\#(a_i^j \mid a_i^j \neq 0)$ represents the number of non-zero coefficients in the j -th constraint.

The main measure we use for comparison is the quality of the lower bound, which is measured by relative gap closed ρ . Using z_{Mc} as the value of McCormick relaxation, and z_{opt} as the best primal value (best solution found among both settings), we define the relative gap closed in percentage as

$$\rho := \frac{z^* - z_{Mc}}{z_{opt} - z_{Mc}} \times 100\%,$$

where z^* is the final lower bound achieved by the corresponding method.

Similarly, we consider the relative gap improvement $\Delta\rho$ to measure the gain by using lifted bilinear cover cuts, which is defined as

$$\Delta\rho := \frac{z_{BC}^* - z_{GRB}^*}{z_{opt} - z_{Mc}} \times 100\% = \rho_{BC} - \rho_{GRB},$$

where z_{BC}^* is the final lower bound achieved by Gurobi for the model where the lifted bilinear cover cuts are added, whereas z_{GRB}^* is the lower bound returned by the default Gurobi non-convex solver without any lifted bilinear cover cuts.

5 Evaluation.

In this section, we evaluate the performance of lifted bilinear cover cuts separated with the heuristic, compared to the default Gurobi solver. Our main focus is the relative gap closed ρ and the relative gap improvement $\Delta\rho$ defined in the previous section. We will however also discuss the number and the quality of bilinear cover cuts produced as well as associated computation times.

5.1 Non-negative instances.

5.1.1 Relative gap closed at the end of the time limit.

We begin by presenting in Figure 1 results regarding the relative gap closed (ρ) for non-negative instances.

As can be seen in the figure, with the bilinear cover cuts applied at the root node, Gurobi closes more gap than the default method, resulting in a final lower bound much closer to the primal value. The improvement reduces when the density p increases, which is anticipated as it becomes more difficult to generate violated lifted bilinear cover cuts using our separation heuristic.

In Figure 2, we take a closer look at the relative gap improvement $\Delta\rho$. Not only is the average $\Delta\rho$ (as shown in the figure) positive, it is in fact positive for most instances.

From analyzing these two figures, we draw the conclusion that lifted bilinear cover cuts provide a stable and robust performance boost for sparse non-negative instances.

5.1.2 Lifted bilinear cover cuts - number and root gap closed.

Next, we investigate the number and the quality of generated lifted bilinear cover cuts. In Figure 3, we present the number of cuts generated. In Figure 4, we display the relative gap closed solely at the root node by lifted bilinear cover cuts separated using the heuristics, which we compute as

$$\rho_{Heu} := \frac{z_{rootBC} - z_{Mc}}{z_{opt} - z_{Mc}}.$$

We observe that the number of separated lifted bilinear cover cuts reduces when p increases. Further, as can be anticipated, the number of cuts generated increases with the number of constraints.

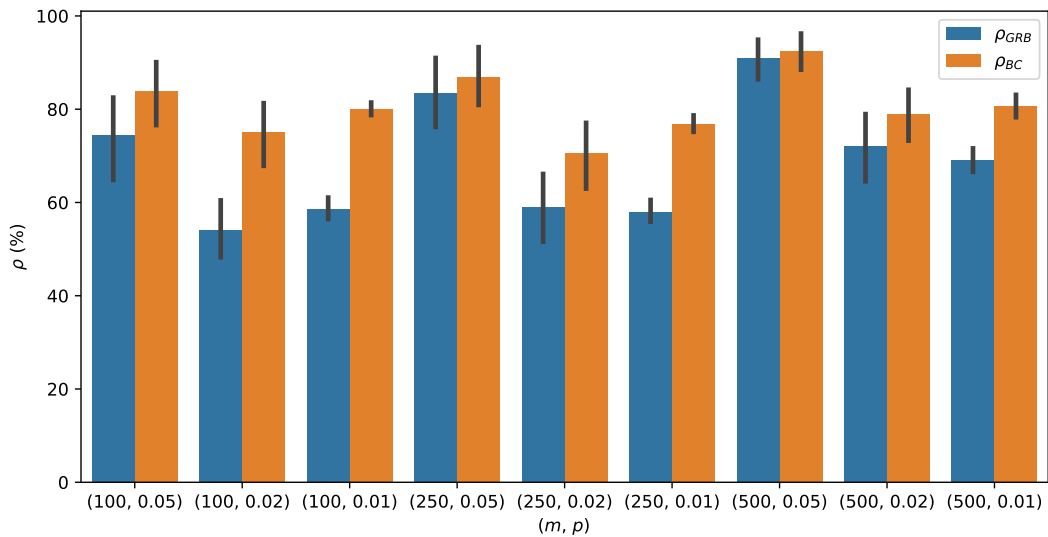


Figure 1: Relative gap closed ρ (%) for non-negative instances: ρ_{GRB} is the gap closed by Gurobi on the separable bilinear instances, and ρ_{BC} is the gap closed by Gurobi on the separable bilinear instances when lifted bilinear cover cuts are added at the root node.

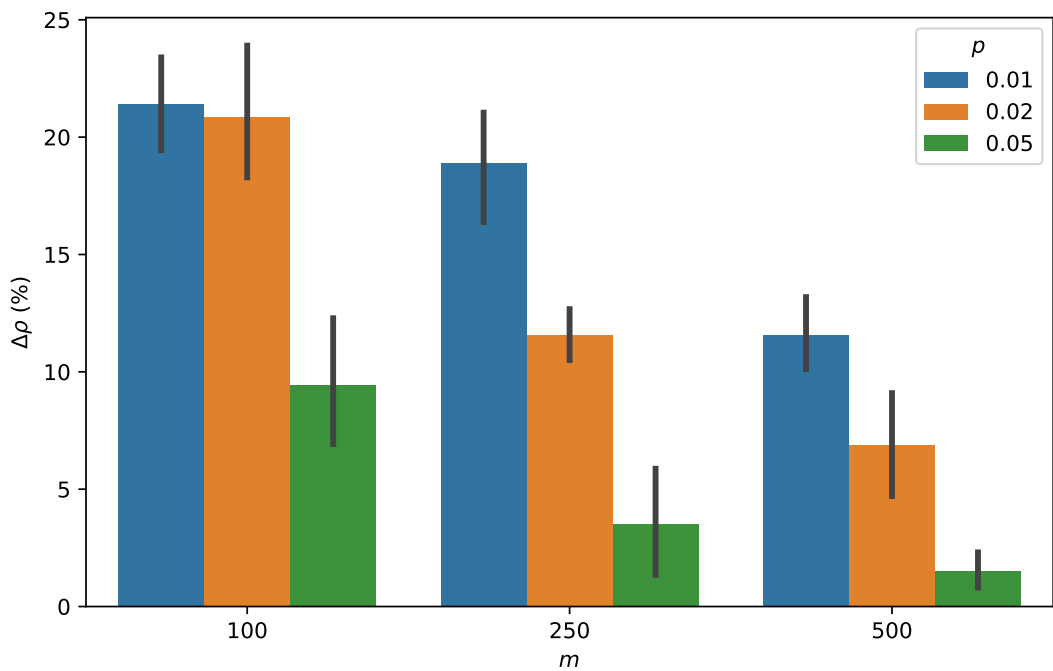


Figure 2: Relative gap improvement $\Delta\rho$ (%) for non-negative instances.

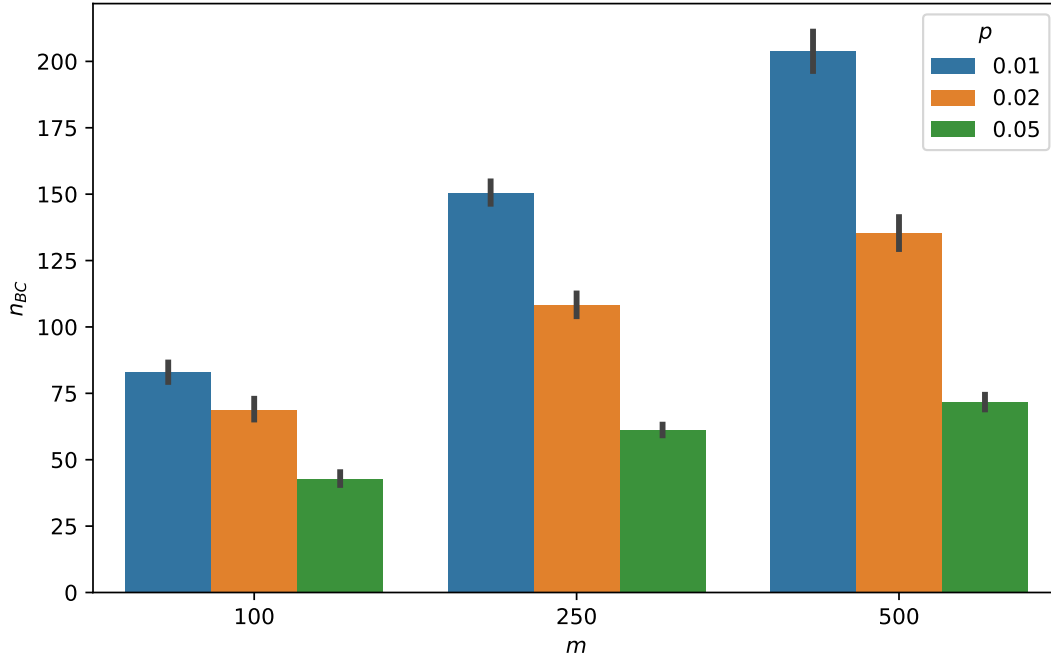


Figure 3: Number of lifted bilinear cover cuts n_{BC} for non-negative instances.

The results in Figure 4 illustrate the strength of lifted bilinear cover cuts over non-negative instances, where much of the gap can be closed by solving the convex relaxation obtained after applying the separation heuristic, which is significantly cheaper than solving the original problem. Interestingly, as far as the root gap closed, the performance of these cuts appears to be independent of the sparsity level p : it is approximately 60% for most classes of instances.

5.1.3 Time.

We consider now the time taken by the heuristic to separate cuts at the root node. As mentioned earlier, we impose a time restriction of 30 minutes for this step. As it turns out, the heuristic (including cut separation and solution of the SOCR convex relaxations) takes only a few seconds; see results in Figure 5.

The time taken by the heuristic increases as the size of the instances becomes larger, though it is still relatively negligible compared to the time spent by the non-convex solver.

Next we present the total times in Figure 6. The difference in time is nearly non-existent, as the root node heuristic takes very little time and the instances remained unsolved after 30 minutes of Gurobi’s run time in both settings. The similarity in overall time supports the fairness of the comparison of the two settings and confirms the effectiveness of the lifted bilinear cover cuts.

5.1.4 Comparison with the cuts (7).

We next compare the lifted bilinear cover cuts with the cuts (7) described in [23]. Because a single one of these cuts can be generated from each row, no involved separation is required. The results are presented in Figure 7 and Figure 8.

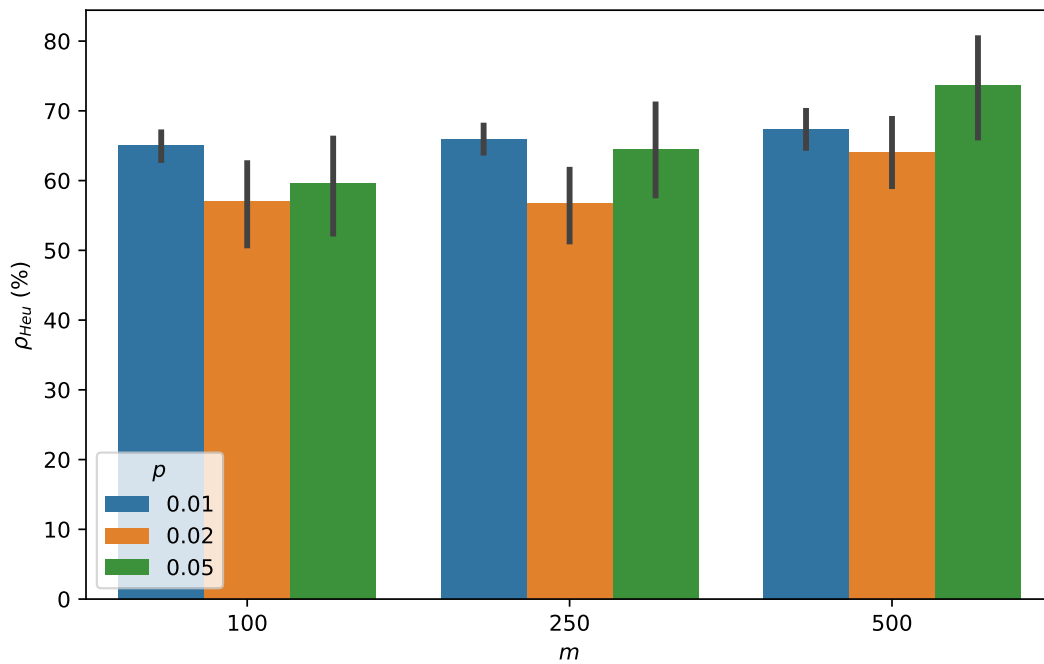


Figure 4: Relative gap closed at root node by lifted bilinear cover cuts (ρ_{Heu}) for non-negative instances.

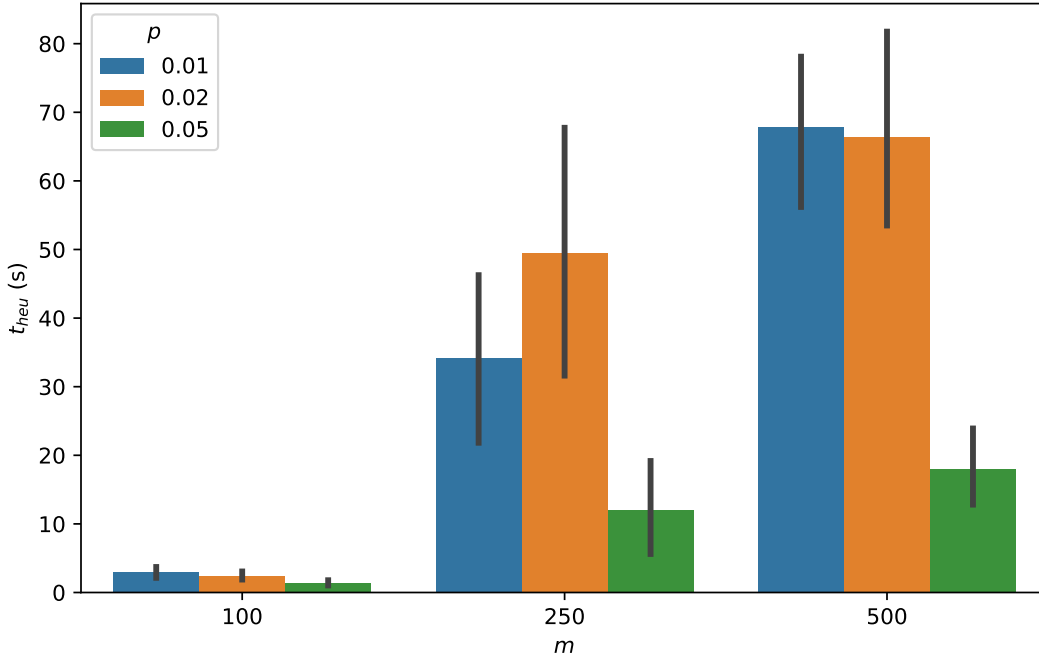


Figure 5: Heuristic CPU time (s) for non-negative instances.

In Figure 7, we present the relative gap closed results at the root node. While it has already been shown from previous sections that the lifted bilinear cover cuts could close a significant portion of gap at the root node, we note that (7) closes a significantly smaller portion of the gap at root node. In Figure 8, the final relative gap closed results are presented. we observe a similar performance between default Gurobi (ρ_{GRB}) and the method applying (7) (ρ_{MT}), with lifted bilinear cover cuts closing more gap (ρ_{BC}).

Combining the above observations, it is clear that the lifted bilinear cover cuts are able to close more gap at the root node and provide a more robust and stable performance boost overall, compared against both default Gurobi as well as cuts (7).

5.2 Mixed-signs instances.

We now focus on mixed-signs instances. We expect that such instances are inherently harder compared to non-negative instances, since it is more difficult to generate lifted bilinear cover cuts for them. Nevertheless, we next show that our heuristic still works well in separating lifted bilinear cover cuts in this case.

5.2.1 Relative gap at the end of the time limit.

We first present in Figure 9 the results regarding the relative gap closed ρ .

For the more difficult mixed-signs instances, our algorithm with root node heuristic again showed improvement, which was especially marked for the sparser instances.

A closer look at Figure 10 confirms the benefits, as we can observe a positive $\Delta\rho$ across every instance, which is especially noticeable for sparse settings (*i.e.*, small values of p).

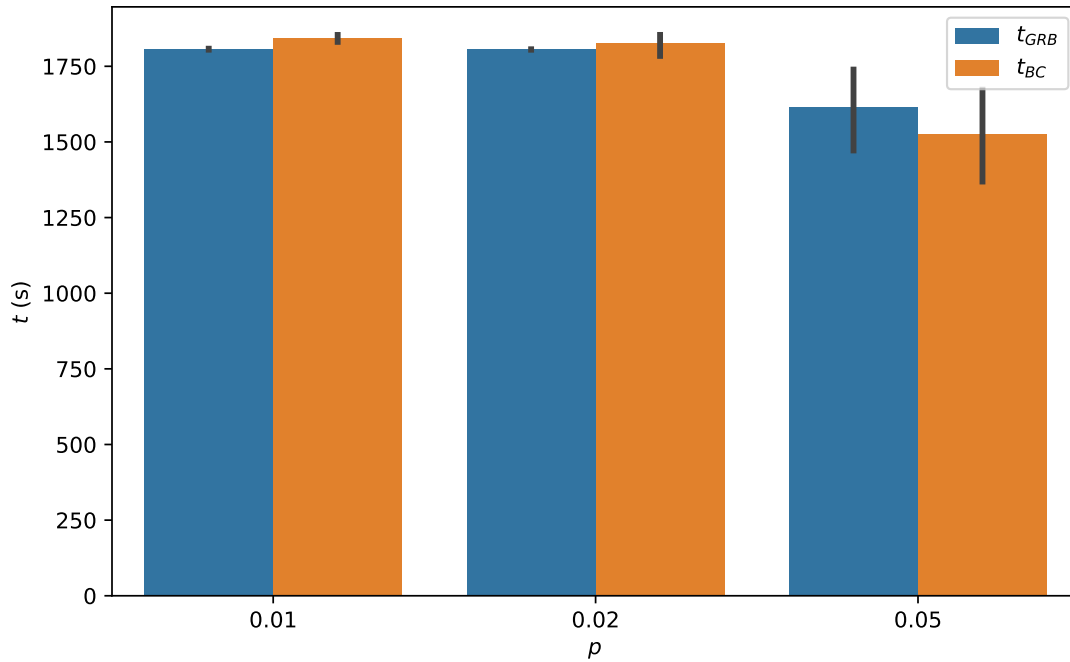


Figure 6: Total CPU time (s) for non-negative instances: t_{BC} is the time for the case where lifted bilinear cover cuts are added at the root node, and t_{GRB} is the time for the other setting.

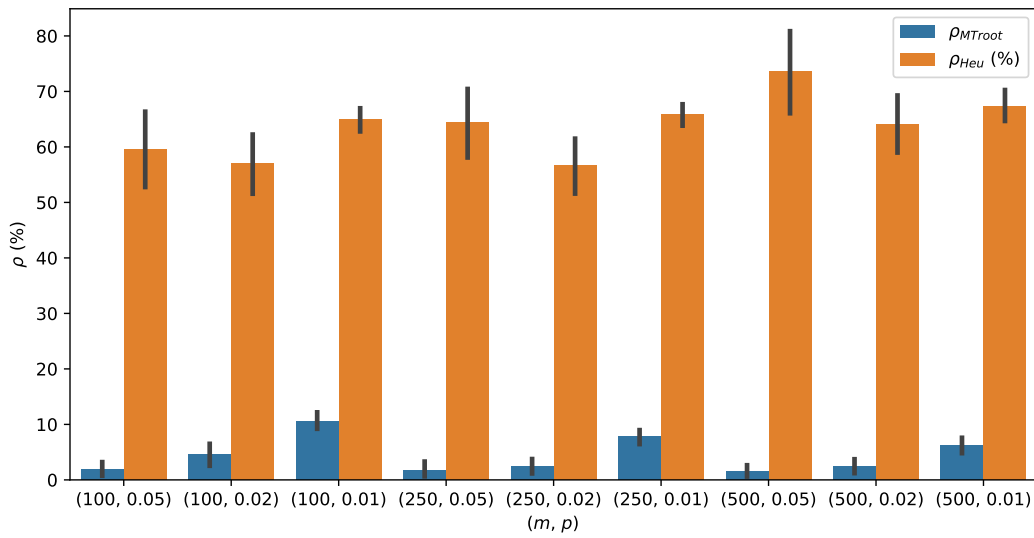


Figure 7: Relative gap closed at root node by lifted bilinear cover cuts (ρ_{Heu}) or by (7) (ρ_{MTroot}) for non-negative instances.

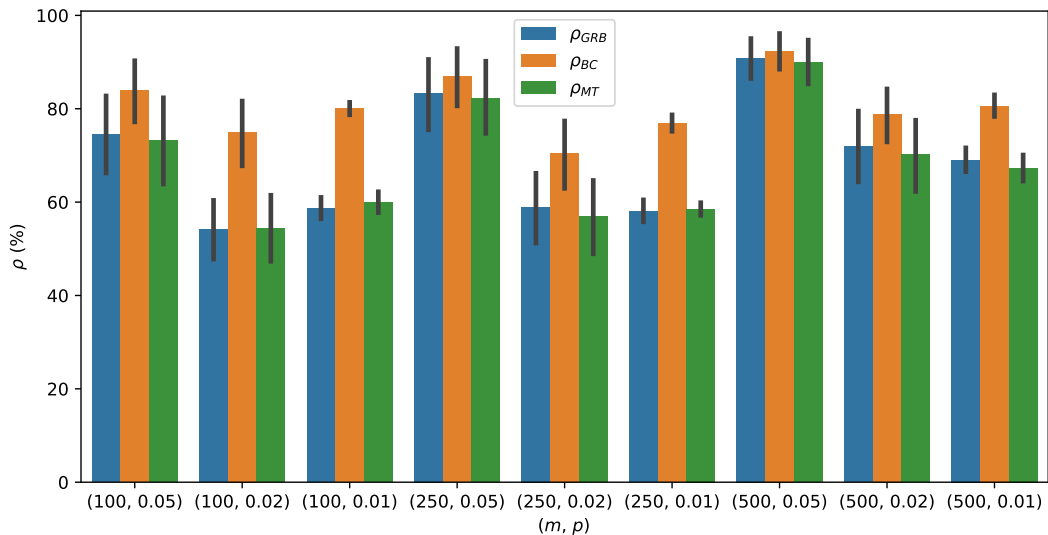


Figure 8: Relative gap closed ρ (%): ρ_{GRB} is the gap closed by Gurobi, ρ_{BC} is the gap closed by Gurobi when lifted bilinear cover cuts are added at the root node, and ρ_{MT} is the gap closed by Gurobi when cuts (7) are added at the root node

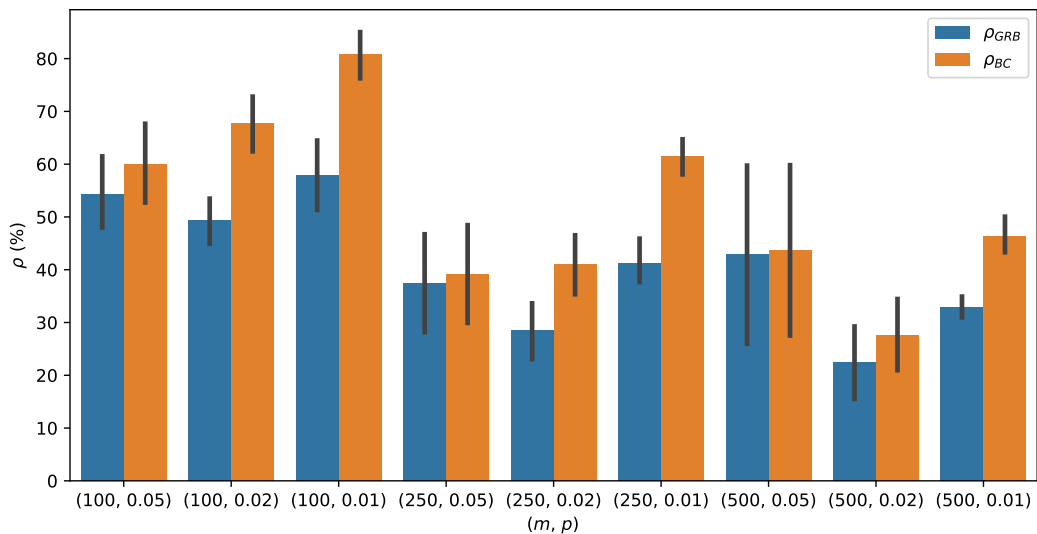


Figure 9: Relative gap closed ρ (%) for mixed-signs instances: ρ_{GRB} is the gap closed by Gurobi on the separable bilinear instances, and ρ_{BC} is the gap closed by Gurobi on the separable bilinear instances when lifted bilinear cover cuts are added at the root node.

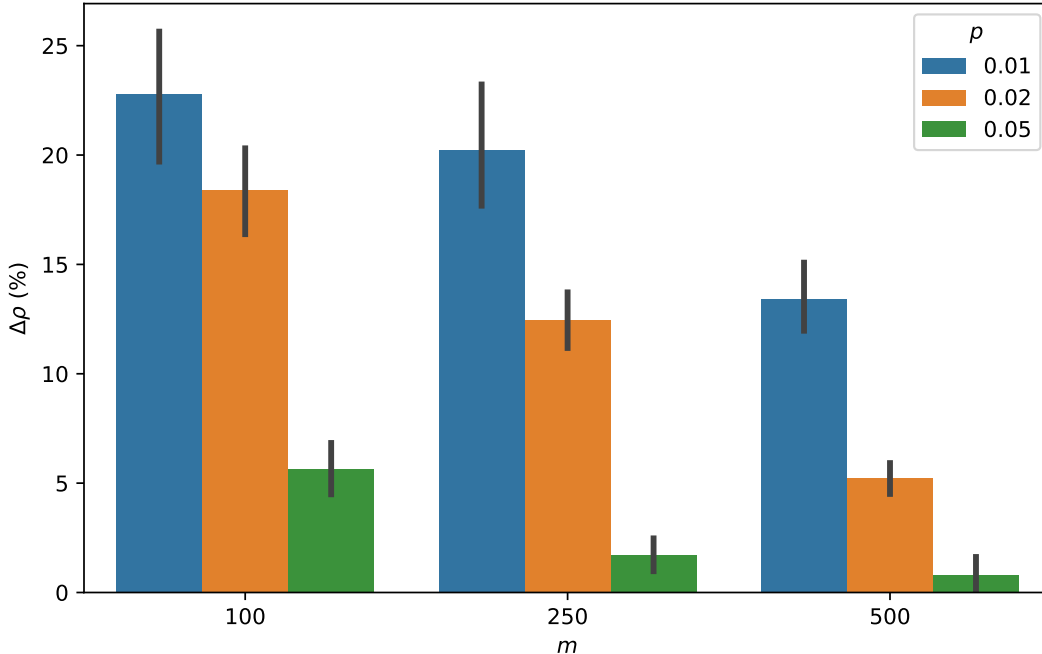


Figure 10: Relative gap improvement $\Delta\rho$ (%) for mixed-signs instances.

Therefore, for the more difficult mixed-signs instances, Figure 9 and Figure 10 indicate that our root node separation heuristic retains its strength by providing robust improvements across instances.

5.2.2 Lifted bilinear cover cuts - number and root gap closed.

We now turn to the number of lifted bilinear cover cuts added; see Figure 11. Figure 12 presents the gap closed by these inequalities at the root node.

As shown in the figures, our algorithm excels at separating lifted bilinear cover cuts for sparser problems, as more cuts were generated for smaller p while closing more gap (*i.e.*, larger ρ_{Heu}). Comparing Figure 3 and Figure 11, it appears (as we expected) that negative coefficients in bilinear rows make separation harder, as evidenced by fewer cuts being generated for mixed-signs instances and the root gap closed also reducing for mixed-signs instances.

5.2.3 Time.

We now focus on total computational time expanded. Figure 13 presents the time for the root node separation heuristic whereas Figure 14 presents the overall time results for mixed-signs instances. These numbers are very similar to the case of non-negative instances: the heuristic itself takes very little time and most of the time is consumed by the non-convex solver.

5.3 Overall Evaluation.

The results for both non-negative instances and mixed-signs instances indicate that with minimal consumption of cpu time, the heuristic we propose is capable of separating very good lifted bilinear

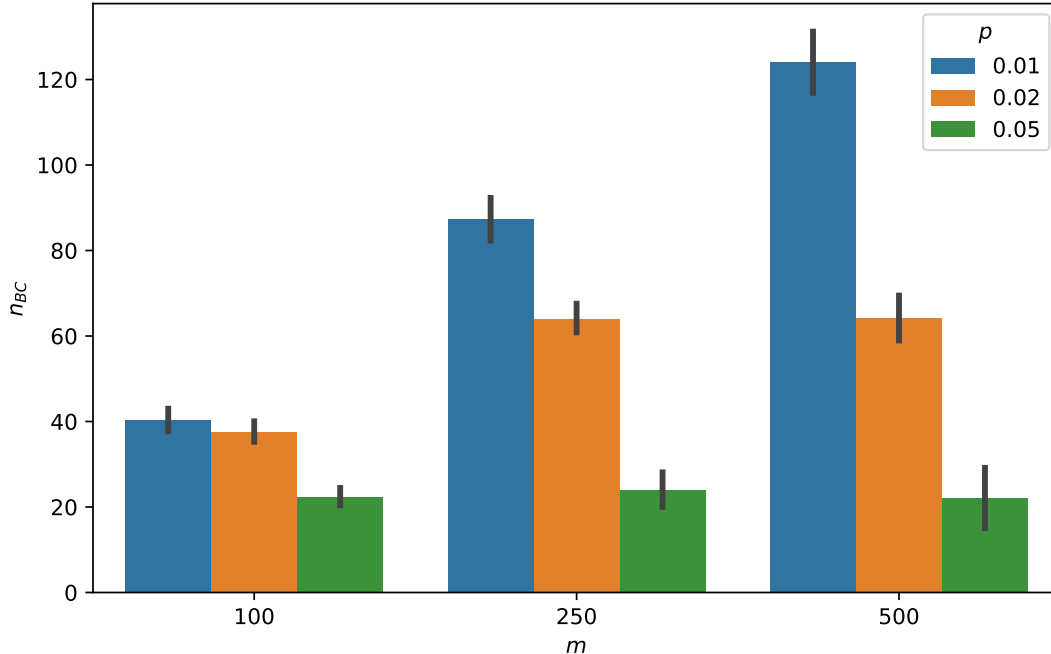


Figure 11: Number of lifted bilinear cover cuts n_{BC} for mixed-signs instances.

cover cuts, thus leading to a stable and robust performance improvement across instances, as measured in the relative gap closed ρ as well as relative gap improvement $\Delta\rho$. Such improvement in performance is especially noticeable when the instances are sparse.

Overall, we believe the strength and potential of the lifted bilinear cover cuts were demonstrated by our experiments, which were most conclusive for sparser problems.

6 Conclusions and Future Directions.

In this paper, we designed a simple heuristic to separate the lifted bilinear cover cuts introduced in [12]. In the computational experiments, the cuts demonstrate major and robust improvement for sparse problems, illustrating the potential of these cuts in situations where SDP-based formulations might not be as helpful.

We envision multiple future directions that will better uncover and utilize the potential of the lifted bilinear cover inequalities, as well as achieving better results:

- **Complexity of separation:** We developed a heuristic for separation since we could not design a polynomial-time exact separation routine. We conjecture that the separation problem is NP-hard although proving this result appears nontrivial.
- **Improving the heuristic for cutting planes:** While the heuristic we developed is quite simple, it uses randomness. It would be interesting to devise a deterministic heuristic and explore other more sophisticated variants of the heuristic. Another direction is to consider separating cuts from aggregations of constraints, an approach that has proven to be useful for MILPs (see [17, 4]) and for QCQPs (see [27, 6, 18, 10]).

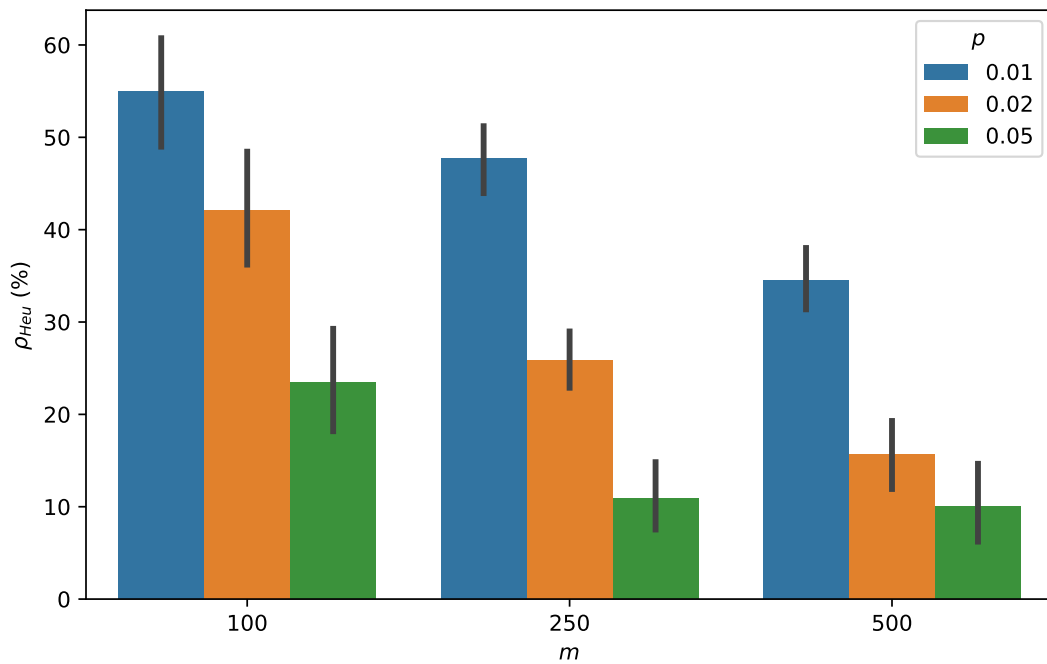


Figure 12: Relative gap closed at root node by lifted bilinear cover cuts (ρ_{Heu}) for mixed-signs instances.

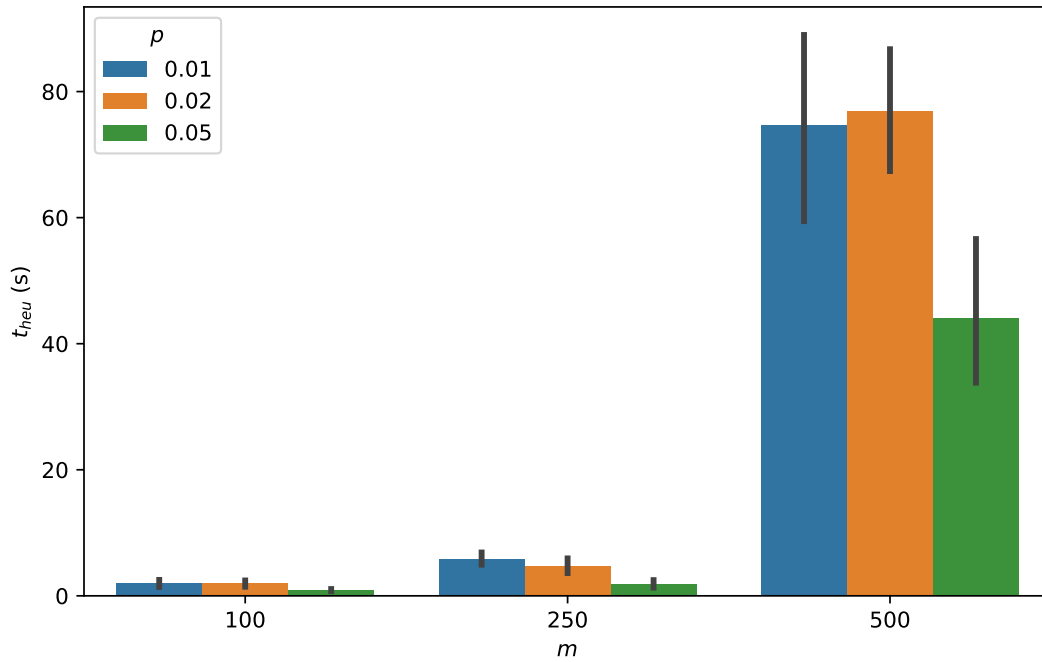


Figure 13: Heuristic CPU time (s) for mixed-signs instances.

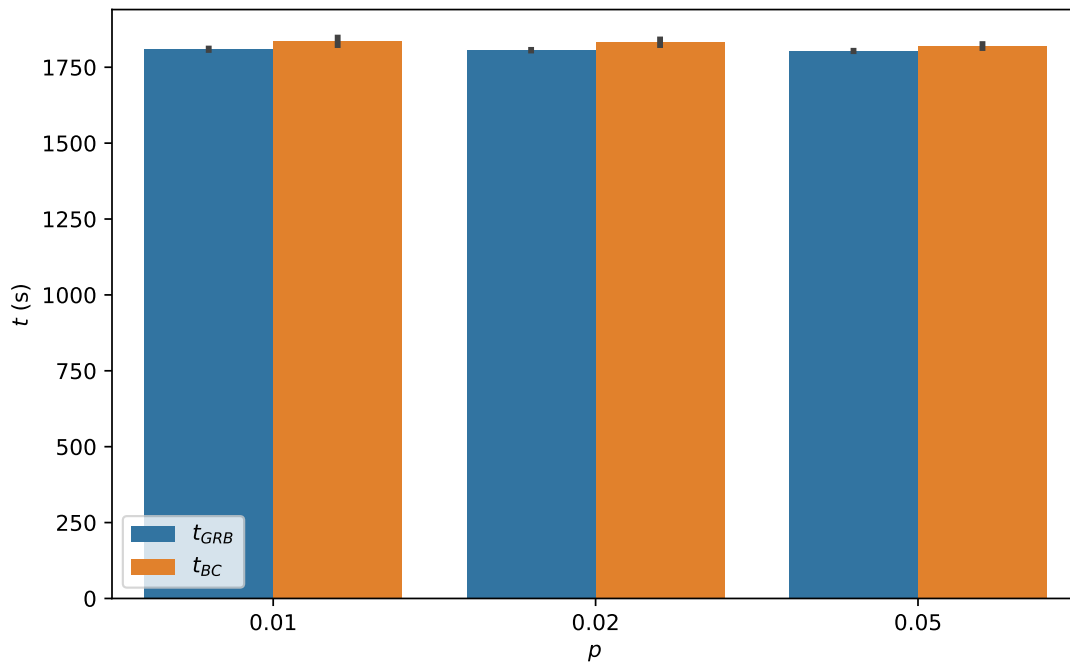


Figure 14: Total CPU time(s) for mixed-signs instances.

- **Integrating the bilinear cover cuts within the solver:** The procedure described in the paper only applies lifted bilinear cover cuts at the root node, and its potential cannot be fully realized without integration within a global solver. It is clear that the method could be applied at children nodes during the branch-and-bound process; evaluating the benefits this provides is an open question.
- **Other problem types:** It is easy to see that any quadratic constraint can be relaxed (at the expense of an increase in the number of variables) into a separable bilinear constraint. An interesting direction of research is to determine if adding these lifted bilinear cover cuts to such relaxation is useful in the solution of nonconvex QCQPs.

7 Acknowledgments

Santanu S. Dey gratefully acknowledges the support by ONR under grant N000141912323.

References

- [1] Egon Balas. Facets of the knapsack polytope. *Mathematical Programming*, 8(1):146–164, 1975.
- [2] Egon Balas and Eitan Zemel. Facets of the knapsack polytope from minimal covers. *SIAM Journal on Applied Mathematics*, 34(1):119–148, 1978.
- [3] Robert Bixby and Edward Rothberg. Progress in computational mixed integer programming—a look back from the other side of the tipping point. *Annals of Operations Research*, 149(1):37, 2007.
- [4] Merve Bodur, Alberto Del Pia, Santanu S Dey, Marco Molinaro, and Sebastian Pokutta. Aggregation-based cutting-planes for packing and covering integer programs. *Mathematical Programming*, 171(1):331–359, 2018.
- [5] Samuel Burer. A gentle, geometric introduction to copositive optimization. *Mathematical Programming*, 151(1):89–116, 2015.
- [6] Samuel Burer and Fatma Kiliç-Karzan. How to convexify the intersection of a second order cone and a nonconvex quadratic. *Mathematical Programming*, 162(1):393–429, 2017.
- [7] Harlan Crowder, Ellis L Johnson, and Manfred Padberg. Solving large-scale zero-one linear programming problems. *Operations Research*, 31(5):803–834, 1983.
- [8] Santanu S Dey, Marco Molinaro, and Qianyi Wang. Approximating polyhedra with sparse inequalities. *Mathematical Programming*, 154(1):329–352, 2015.
- [9] Santanu S Dey, Marco Molinaro, and Qianyi Wang. Analysis of sparse cutting planes for sparse milps with applications to stochastic milps. *Mathematics of Operations Research*, 43(1):304–332, 2018.
- [10] Santanu S Dey, Gonzalo Munoz, and Felipe Serrano. On obtaining the convex hull of quadratic inequalities via aggregations. *SIAM Journal on Optimization*, 32(2):659–686, 2022.
- [11] Santanu S. Dey, Asteroide Santana, and Yang Wang. New SOCP relaxation and branching rule for bipartite bilinear programs. *Optimization and Engineering*, 20(2):307–336, 2019.

- [12] Xiaoyi Gu, Santanu S Dey, and Jean-Philippe P Richard. Lifting convex inequalities for bipartite bilinear programs. *Mathematical Programming*, pages 1–33, 2022.
- [13] Zonghao Gu, George L. Nemhauser, and Martin W. P. Savelsbergh. Lifted flow cover inequalities for mixed 0-1 integer programs. *Mathematical Programming*, 85(3):439–467, 1999.
- [14] Zonghao Gu, George L. Nemhauser, and Martin W. P. Savelsbergh. Sequence independent lifting in mixed integer programming. *Journal of Combinatorial Optimization*, 4(1):109–129, 2000.
- [15] Peter L. Hammer, Ellis L. Johnson, and Uri N. Peled. Facet of regular 0–1 polytopes. *Mathematical Programming*, 8(1):179–206, 1975.
- [16] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- [17] Hugues Marchand and Laurence A Wolsey. Aggregation and mixed integer rounding to solve MIPs. *Operations Research*, 49(3):363–371, 2001.
- [18] Sina Modaresi and Juan Pablo Vielma. Convex hull of two quadratic or a conic quadratic and a quadratic inequality. *Mathematical Programming*, 164(1):383–409, 2017.
- [19] Manfred W. Padberg. A note on zero-one programming. *Operations Research*, 23(4):833–837, 1975.
- [20] Jean-Philippe P. Richard. Lifting techniques for mixed integer programming. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [21] Asteroide Santana and Santanu S. Dey. The convex hull of a quadratic constraint over a polytope. *SIAM Journal on Optimization*, 30(4):2983–2997, 2020.
- [22] T Schoning. A probabilistic algorithm for k-SAT and constraint satisfaction problems. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 410–414. IEEE, 1999.
- [23] Mohit Tawarmalani, Jean-Philippe P. Richard, and Kwanghun Chung. Strong valid inequalities for orthogonal disjunctions and bilinear covering sets. *Mathematical Programming*, 124(1-2):481–512, 2010.
- [24] Laurence A. Wolsey. Faces for a linear inequality in 0–1 variables. *Mathematical Programming*, 8(1):165–178, 1975.
- [25] Laurence A. Wolsey. Facets and strong valid inequalities for integer programs. *Operations Research*, 24(2):367–372, 1976.
- [26] Laurence A. Wolsey. Valid inequalities and superadditivity for 0–1 integer programs. *Mathematics of Operations Research*, 2(1):66–77, 1977.
- [27] Uğur Yildiran. Convex hull of two quadratic constraints is an LMI set. *IMA Journal of Mathematical Control and Information*, 26(4):417–450, 2009.
- [28] Eitan Zemel. Lifting the facets of zero–one polytopes. *Mathematical Programming*, 15(1):268–277, 1978.