

# Learning to Accelerate the Global Optimization of Quadratically-Constrained Quadratic Programs

Rohit Kannan<sup>1,2</sup>, Harsha Nagarajan<sup>2</sup>, and Deepjyoti Deka<sup>2</sup>

<sup>1</sup>Center for Nonlinear Studies (T-CNLS), Los Alamos National Laboratory, Los Alamos, NM, USA.

<sup>2</sup>Applied Mathematics & Plasma Physics (T-5), Los Alamos National Laboratory, Los Alamos, NM, USA.  
E-mail: {rohitk@alum.mit.edu, harsha@lanl.gov, deepjyoti@lanl.gov}

Version 2 (this document): February 22, 2023

Version 1: December 31, 2022

## Abstract

We learn optimal instance-specific heuristics for the global minimization of nonconvex quadratically-constrained quadratic programs (QCQPs). Specifically, we consider partitioning-based mixed-integer programming relaxations for nonconvex QCQPs and propose the novel problem of *strong partitioning* to optimally partition variable domains *without* sacrificing global optimality. We design a local optimization method for solving this challenging max-min strong partitioning problem and replace this expensive benchmark strategy with a machine learning (ML) approximation for homogeneous families of QCQPs. We present a detailed computational study on randomly generated families of QCQPs, including instances of the pooling problem, using the open-source global solver Alpine. Our numerical experiments demonstrate that strong partitioning and its ML approximation significantly reduce Alpine’s solution time by factors of 3.5 – 16.5 and 2 – 4.5 *on average* and by maximum factors of 15 – 700 and 10 – 200, respectively, over the different QCQP families.

**Key words:** Quadratically-Constrained Quadratic Program, Piecewise McCormick Relaxations, Global Optimization, Machine Learning, Strong Partitioning, Sensitivity Analysis, Pooling Problem

## 1 Introduction

Many real-world applications involve the repeated solution of the same underlying quadratically-constrained quadratic program (QCQP) with slightly varying model parameters. Examples include the pooling problem with varying input qualities [53] and the cost-efficient operation of the power grid with varying loads and renewable sources [13]. These hard optimization problems are typically solved using off-the-shelf global optimization software [8, 12, 45, 54, 56, 61] that do not exploit the shared problem structure—heuristics within these implementations are engineered to work well *on average* over a diverse set of instances and may perform sub-optimally for instances from a specific application [46]. Recent work [9, 47] has shown that tailoring branching decisions can significantly accelerate branch-and-bound (B&B) algorithms for mixed-integer linear programs (MILPs). In contrast, only a few papers (see Section 2.2) attempt to use machine learning (ML) to accelerate the guaranteed global solution of nonconvex *nonlinear* programs (NLPs).

We use ML to accelerate the global minimization of nonconvex QCQPs. We focus on accelerating partitioning algorithms [11, 62, 68] since they are effective in solving this challenging class of problems [19, 56]. Partitioning algorithms determine lower bounds on the optimal value of a nonconvex QCQP using piecewise convex relaxations. They begin by selecting a subset of the continuous variables participating in nonconvex terms for partitioning. At each iteration, they refine the partitions of the domains of these variables and update the piecewise convex relaxations in their lower bounding formulation. They then solve a convex

mixed-integer program (MIP) [48, 65] to determine a lower bound. Partitioning algorithms typically use heuristics to specify the locations of partitioning points and continue to refine their variable partitions until the lower bounds converge to the optimal objective value. Since the complexity of their MIP relaxations can grow significantly at each iteration, the choice of partitioning points in the initial iterations can have a huge impact on the overall performance of these algorithms [56]. Despite their importance, the optimal choice of partitioning points is not well understood and current approaches resort to heuristics such as bisecting the active partition [19, 62, 68], or adding partitioning points at/around the lower bounding solution [11, 56] to refine variable partitions.

**Proposed approach.** We learn to optimally partition the domains of variables in a given QCQP *without* sacrificing global optimality. Similar to the concept of *strong branching* in B&B algorithms for MIPs [1, 8], we propose the novel concept of *strong partitioning* to choose partitioning points. The key idea of strong partitioning is to determine a specified number of partitioning points per variable such that the resulting piecewise convex relaxation-based lower bound is maximized. We formulate strong partitioning as a max-min problem, where the outer-maximization chooses the partitioning points and the inner-minimization solves the piecewise relaxation-based lower bounding problem for a given partition. We solve this max-min problem to local optimality by using generalized gradients of the value function of the inner-minimization problem within a bundle solver for nonsmooth nonconvex optimization. Because each iteration of the bundle method requires the solution of a MIP, solving this max-min problem may be computationally expensive. Therefore, we use ML to learn the strong partitioning strategy for homogeneous QCQP instances. We evaluate the performance of strong partitioning and an off-the-shelf ML approximation on randomly generated QCQPs, including instances of the pooling problem, using the open-source global solver Alpine [55, 56]. Our experiments demonstrate that using strong partitioning to initialize Alpine’s variable partitions can reduce its solution time by a factor of  $3.5 - 16.5$  on average and by a maximum factor of  $15 - 700$ . They also illustrate that our ML model is able to learn the strong partitioning strategy approximately, reducing Alpine’s solution time by a factor of  $2 - 4.5$  on average and by a maximum factor of  $10 - 200$  over the same set of instances.

This paper is organized as follows. Section 2 reviews approaches that use ML to accelerate the guaranteed global solution of MILPs, NLPs, and mixed-integer NLPs (MINLPs). Section 3 outlines partitioning-based lower bounding methods for nonconvex QCQPs. Section 4 introduces strong partitioning and designs an algorithm for its solution with theoretical guarantees. Section 5 presents detailed computational results demonstrating the effectiveness of using strong partitioning and an off-the-shelf ML approximation within Alpine for randomly generated QCQPs, including instances of the pooling problem. We conclude with directions for future research in Section 6.

**Notation.** Let  $[n] := \{1, 2, \dots, n\}$  and  $\mathbb{R}_+$  denote the set of non-negative reals. We write  $v_i$  to denote the  $i$ th component of vector  $v = (v_1, v_2, \dots, v_n) \in \mathbb{R}^n$ ,  $M_{ij}$  to denote the  $(i, j)$ th component of matrix  $M$ , and  $|S|$ ,  $\text{int}(S)$ , and  $\text{conv}(S)$  to denote the cardinality, interior, and convex hull of a set  $S$ , respectively.

## 2 Related work

Optimization solvers tune key algorithmic parameters by extensive testing on benchmark libraries. However, they typically only consider a narrow family of efficiently computable heuristics. Moreover, they do not tailor these heuristics to each problem instance but seek a universal setting for good *average* solver performance. Machine learning, on the other hand, can enable efficient approximations of better performing but expensive heuristics, potentially leading to significant computational gains for hard test instances. Several recent papers survey the burgeoning field of using ML to accelerate MILP and combinatorial optimization algorithms [9, 18, 42, 47]. In the next sections, we review related approaches on learning to branch for MILPs and learning to accelerate the guaranteed solution of (MI)NLPs.

## 2.1 Learning to branch for MILPs

Branch-and-bound and its variants form the backbone of modern MILP solvers. Besides heuristics for determining good feasible solutions, selecting the branching variable at each node of the B&B tree is probably the most impactful decision in terms of the run time of the algorithm [1]. Typically, a subset of the integer variables with fractional lower bounding solutions at a particular node are considered as the candidate branching variables for that node. The gold standard branching variable selection heuristic is *full strong branching* (FSB), which chooses as the branching variable the one that leads to the maximum product of improvements in the lower bounds of the two children nodes (assuming they are both feasible). FSB results in a 65% reduction in the number of nodes explored by the B&B tree (relative to the default branching strategy) on average over standard test instances; however, this comes with a 44% increase in the cost-per-node that is not tenable [1]. MILP solvers therefore tend to use computationally cheaper heuristic approximations of FSB such as reliability, pseudocost, or hybrid branching. Motivated by FSB’s promise, most approaches for learning to branch for MILPs aim to develop a computationally efficient ML approximation that retains its attractive performance.

Alvarez et al. [4] propose hand-crafted features to construct an ML approximation of FSB using Extremely Randomized Trees. Khalil et al. [38] propose an instance-specific on-the-fly ML approximation of FSB using Support Vector Machines (SVMs) and use the learned approximation for subsequent branching decisions in the same problem instance. Gasse et al. [28] and Nair et al. [57] design graph neural network (GNN) approximations of FSB that leverages the bipartite graph representation of MILPs and removes the need for feature engineering. Zarpellon et al. [70] seek to learn branching policies that generalize to heterogeneous MILPs by explicitly parametrizing the state of the B&B tree. A few works (see, e.g., [25, 32]) study online and reinforcement learning approaches for making branching decisions. Some others [5, 23] learn combinations of existing heuristics to come up with better branching decisions.

## 2.2 Learning for (MI)NLPs

There are relatively fewer approaches in the literature for accelerating the *guaranteed global solution* of nonconvex NLPs and MINLPs using ML. To the best of our knowledge, none of these approaches use ML to accelerate partitioning-based global optimization algorithms.

Baltean-Lugojan et al. [6] consider the global solution of nonconvex QCQPs using semi-definite programming (SDP) relaxations. They use ML to construct good outer-approximations of these SDP relaxations to mitigate the computational burden of SDP solvers. They train a neural network to select cuts based on their sparsity and predicted impact on the objective, and show that their approach results in computationally cheap relaxations that can be effectively integrated into global solvers.

Ghaddar et al. [29] consider branching variable selection for a B&B search tree that is embedded within the reformulation-linearization technique (RLT) for solving polynomial problems. They use ML to choose the “best branching strategy” from a portfolio of branching rules (cf. Di Liberto et al. [23]) based on violations of RLT-defining identities. They design several hand-crafted features and pick the branching strategy that optimizes a quantile regression forest-based approximation of their performance indicator. González-Rodríguez et al. [30] consider a portfolio of second-order cone and SDP constraints to strengthen the RLT formulation for polynomial problems and use ML to select constraints to add within a B&B framework.

Bonami et al. [15] train classifiers to predict whether linearizing products of binary variables or binary and bounded continuous variables may be computationally advantageous for solving MIQPs. Nannicini et al. [58] train an SVM classifier to decide if an expensive optimality-based bounds tightening (OBBT) routine should be used in lieu of a cheaper feasibility-based routine for nonconvex MINLPs. Cengil et al. [20] consider the AC optimal power flow problem and train a deep neural network to identify a small subset of lines and buses for which a reduced-cost OBBT routine is applied. Finally, Lee et al. [43] use classification and regression techniques to identify effective cuts for the generalized Benders decomposition master problem.

We review partitioning algorithms for the global minimization of QCQPs in the next section before outlining our proposal to accelerate them using ML.

### 3 Partitioning-based bounds for QCQPs

Consider the nonconvex QCQP

$$\begin{aligned} \min_{x \in [0,1]^n} \quad & x^\top Q^0 x + (r^0)^\top x \\ \text{s.t.} \quad & x^\top Q^i x + (r^i)^\top x \leq b_i, \quad \forall i \in [m_I], \end{aligned} \quad (1)$$

where  $b \in \mathbb{R}^{m_I}$  and for each  $k \in \{0\} \cup [m_I]$ ,  $r^k \in \mathbb{R}^n$  and  $Q^k \in \mathbb{R}^{n \times n}$  are symmetric but not necessarily positive semi-definite. QCQPs with equality constraints and different variable bounds can be handled using simple transformations. Polynomial optimization problems may also be reformulated as QCQPs through the addition of variables and constraints. It is well known that nonconvex QCQPs are NP-hard [66].

QCQPs arise in several applications (see [53] for a detailed list) such as facility location [41], refinery optimization [34, 69], electric grid optimization [13], and circle packing [21]. By introducing variables and constraints, we can reformulate (1) into the following equivalent form:

$$\begin{aligned} v^* := \min_{x \in [0,1]^n, w} \quad & c^\top x + d^\top w \\ \text{s.t.} \quad & Ax + Bw \leq b, \\ & w_{ij} = x_i x_j, \quad \forall (i, j) \in \mathcal{B}, \\ & w_{kk} = x_k^2, \quad \forall k \in \mathcal{Q}, \end{aligned} \quad (\text{QCQP})$$

for some vectors  $c$  and  $d$ , matrices  $A$  and  $B$ , and index sets  $\mathcal{B} \subset \{(i, j) \in [n]^2 : i \neq j\}$  and  $\mathcal{Q} \subset [n]$  of (pairs of) variables participating in bilinear and univariate quadratic terms. We assume (QCQP) is feasible for simplicity, and define the set  $F := \{(x, w) : x \in [0, 1]^n, Ax + Bw \leq b\}$  for convenience.

One of the earliest approaches for constructing lower bounds on the optimal value of (QCQP) uses termwise McCormick relaxations [2, 51] to yield the following lower bounding problem:

$$\begin{aligned} \min_{(x, w) \in F} \quad & c^\top x + d^\top w \\ \text{s.t.} \quad & (x_i, x_j, w_{ij}) \in \mathcal{M}_{ij}^B, \quad \forall (i, j) \in \mathcal{B}, \\ & (x_k, w_{kk}) \in \mathcal{M}_k^Q, \quad \forall k \in \mathcal{Q}, \end{aligned} \quad (2)$$

where the bilinear and quadratic equality constraints in (QCQP) have been replaced with the following valid convex relaxations on  $x \in [0, 1]^n$  for each  $(i, j) \in \mathcal{B}$  and  $k \in \mathcal{Q}$ :

$$\begin{aligned} \mathcal{M}_{ij}^B &:= \{(x_i, x_j, w_{ij}) : 0 \leq w_{ij} \leq x_i, x_i + x_j - 1 \leq w_{ij} \leq x_j\}, \\ \mathcal{M}_k^Q &:= \{(x_k, w_{kk}) : x_k^2 \leq w_{kk} \leq x_k\}. \end{aligned}$$

Problem (2) can be used within a spatial B&B framework for solving (QCQP) to global optimality. Several papers (see, e.g., [7, 11, 14, 17, 59, 62, 63, 68]) improve upon the termwise McCormick bound, usually at an increase in the computational cost but with the goal of reducing the overall time for the B&B algorithm to converge. In this work, we consider the so-called piecewise McCormick relaxation approach [10, 11, 19, 37, 40, 52, 56, 62, 68] for strengthening the termwise McCormick relaxations.

Piecewise McCormick relaxations begin by partitioning the domains of variables participating in nonconvex terms into sub-intervals. Assume for simplicity that we wish to partition the domain of each  $x_i$  into  $d + 1$  sub-intervals with  $d \geq 1$  (general partitioning schemes can be handled similarly). For each  $i \in [n]$ , let

$$\mathcal{P}_i := [p_0^i, p_1^i, p_2^i, \dots, p_d^i, p_{d+1}^i], \quad \text{with} \quad 0 =: p_0^i \leq p_1^i \leq p_2^i \leq \dots \leq p_d^i \leq p_{d+1}^i := 1,$$

denote the array of  $d + 2$  partitioning points for variable  $x_i$ , including the original variable bounds 0 and 1. Given partitions  $\mathcal{P}_i$ ,  $i \in [n]$ , the piecewise McCormick relaxation-based lower bounding problem to (QCQP)




Figure 1: The left and middle plots illustrate the lower and upper parts, respectively, of the piecewise McCormick relaxation for the bilinear term  $w_{12} = x_1x_2$  on the domain  $x_1, x_2 \in [-1, 1]$  (domain changed from  $[0, 1]$  for better illustration) with partitions  $\mathcal{P}_1 = \mathcal{P}_2 = [-1, 0, 1]$ . The right plot illustrates the piecewise McCormick relaxation for the quadratic term  $w_{11} = x_1^2$  on the domain  $x_1 \in [0, 1]$  (the lower part coincides with the red quadratic curve) with the partition  $\mathcal{P}_1 = [0, 0.5, 1]$ .

can be abstractly written as

$$\begin{aligned}
\min_{(x,w) \in F} \quad & c^T x + d^T w & (3) \\
\text{s.t.} \quad & (x_i, x_j, w_{ij}) \in \mathcal{PMR}_{ij}^B(\mathcal{P}_i, \mathcal{P}_j), \quad \forall (i, j) \in \mathcal{B}, \\
& (x_k, w_{kk}) \in \mathcal{PMR}_k^Q(\mathcal{P}_k), \quad \forall k \in \mathcal{Q},
\end{aligned}$$

where  $\mathcal{PMR}_{ij}^B(\mathcal{P}_i, \mathcal{P}_j)$  and  $\mathcal{PMR}_k^Q(\mathcal{P}_k)$  denote the feasible regions corresponding to the piecewise McCormick relaxations of the bilinear equation  $w_{ij} = x_i x_j$  and the quadratic equation  $w_k = x_k^2$ , respectively. While there are several ways of formulating these piecewise McCormick relaxations, we use the so-called “convex combination” or “lambda” formulation below (see [39] for enhancements in the multilinear setting). The piecewise McCormick relaxation for the bilinear constraint  $w_{ij} = x_i x_j$  can be represented as follows [65]:

$$\begin{aligned}
\mathcal{PMR}_{ij}^B(\mathcal{P}_i, \mathcal{P}_j) := \left\{ (x_i, x_j, w_{ij}) : \exists \lambda^{ij} \in \mathbb{R}_+^{(d+2)^2}, y^i \in \{0, 1\}^{(d+1)}, y^j \in \{0, 1\}^{(d+1)} \right. \\
\left. \text{s.t. } (x_i, x_j, w_{ij}, \lambda^{ij}, y^i, y^j) \text{ satisfies (4a) – (4d)} \right\},
\end{aligned}$$

where

$$x_i = \sum_{k=0}^{d+1} \sum_{l=0}^{d+1} \lambda_{k(d+2)+l+1}^{ij} p_l^i, \quad x_j = \sum_{k=0}^{d+1} \sum_{l=0}^{d+1} \lambda_{k(d+2)+l+1}^{ij} p_l^j, \quad w_{ij} = \sum_{k=0}^{d+1} \sum_{l=0}^{d+1} \lambda_{k(d+2)+l+1}^{ij} p_l^i p_l^j, \quad (4a)$$

$$\sum_{k=1}^{d+1} y_k^i = 1, \quad \sum_{k=1}^{d+1} y_k^j = 1, \quad \sum_{k=1}^{(d+2)^2} \lambda_k^{ij} = 1, \quad (4b)$$

$$\sum_{k=0}^{d+1} \lambda_{k(d+2)+1}^{ij} \leq y_1^i, \quad \sum_{k=1}^{d+2} \lambda_{k(d+2)}^{ij} \leq y_{d+1}^i, \quad \sum_{k=0}^{d+1} \lambda_{k(d+2)+l+1}^{ij} \leq y_l^i + y_{l+1}^i, \quad \forall l \in [d], \quad (4c)$$

$$\sum_{k=1}^{d+2} \lambda_k^{ij} \leq y_1^j, \quad \sum_{k=1}^{d+2} \lambda_{(d+2)^2-k}^{ij} \leq y_{d+1}^j, \quad \sum_{k=1}^{d+2} \lambda_{l(d+2)+k}^{ij} \leq y_l^j + y_{l+1}^j, \quad \forall l \in [d]. \quad (4d)$$

Note that only equations (4a) depend on the partitions  $\mathcal{P}_i$  and  $\mathcal{P}_j$  of  $x_i$  and  $x_j$ , respectively, which are parameters in these constraints. The binary vectors  $y^i$  and  $y^j$  denote the active partition of  $x_i$  and  $x_j$ —these variables may be reused in the piecewise McCormick relaxations of other nonconvex terms involving  $x_i$  or  $x_j$ .

The piecewise McCormick relaxation for the quadratic equality constraint  $w_{kk} = x_k^2$  can be represented as follows [48, 56]:

$$\mathcal{PMR}_k^Q(\mathcal{P}_k) := \left\{ (x_k, w_{kk}) : \exists \lambda^k \in \mathbb{R}_+^{(d+2)}, y^k \in \{0, 1\}^{(d+1)}, \text{ s.t. } (x_k, w_{kk}, \lambda^k, y^k) \text{ satisfies (5a) - (5c)} \right\},$$

where

$$x_k = \sum_{l=0}^{d+1} \lambda_{l+1}^k p_l^k, \quad w_{kk} \leq \sum_{l=0}^{d+1} \lambda_{l+1}^k (p_l^k)^2, \quad \sum_{l=1}^{d+1} y_l^k p_{l-1}^k \leq x_k \leq \sum_{l=2}^{d+2} y_{l-1}^k p_{l-1}^k, \quad (5a)$$

$$\sum_{l=1}^{d+1} y_l^k = 1, \quad \sum_{l=1}^{d+2} \lambda_l^k = 1, \quad w_{kk} \geq x_k^2, \quad (5b)$$

$$\lambda_1^k \leq y_1^k, \quad \lambda_{d+2}^k \leq y_{d+1}^k, \quad \lambda_{l+1}^k \leq y_l^k + y_{l+1}^k, \quad \forall l \in [d]. \quad (5c)$$

Note that only equations (5a) depend on the partition  $\mathcal{P}_k$  of  $x_k$ . Additionally, equations (5b) involve convex quadratic functions of  $x_k$ . Figure 1 illustrates the piecewise McCormick relaxations for a bilinear and a univariate quadratic term.

Using the above representations of  $\mathcal{PMR}_{ij}^B$  and  $\mathcal{PMR}_k^Q$ , we obtain the following extended convex MIP formulation for the piecewise McCormick relaxation-based lower bounding problem to (QCQP):

$$\begin{aligned} \min_{(x,w) \in F, \lambda \geq 0, y \in Y} \quad & c^T x + d^T w & (\text{PMR}) \\ \text{s.t.} \quad & (x_i, x_j, w_{ij}, \lambda^{ij}, y^i, y^j) \text{ satisfies (4a) - (4d), } \forall (i, j) \in \mathcal{B}, \\ & (x_k, w_{kk}, \lambda^k, y^k) \text{ satisfies (5a) - (5c), } \forall k \in \mathcal{Q}, \end{aligned}$$

where  $Y := \{y \in \{0, 1\}^{|\mathcal{N}\mathcal{C}| \times (d+1)} : \sum_{l=1}^{d+1} y_l^i = 1, \forall i \in \mathcal{N}\mathcal{C}\}$  is a special-ordered set of type 1, the set  $\mathcal{N}\mathcal{C} := \{i : \exists j \in [n] \text{ s.t. } (i, j) \in \mathcal{B}\} \cup \mathcal{Q}$  collects the indices of variables participating in nonconvex terms, variables  $\lambda$  comprise  $\lambda^{ij}$ ,  $(i, j) \in \mathcal{B}$ , and  $\lambda^k$ ,  $k \in \mathcal{Q}$ , and variables  $y$  comprise  $y^i$ ,  $i \in \mathcal{N}\mathcal{C}$ . Problem (PMR) is a convex mixed-integer QCQP (MIQCQP).

Algorithm 1 outlines a partitioning-based global optimization algorithm that solves problem (PMR) to determine a sequence of lower bounds on the optimal value of (QCQP). Note that we do not consider bound tightening steps in this work, but our approach can be readily adapted to the setting where bounds tightening is employed. As mentioned in the introduction, the choice of heuristics on line 7 of Algorithm 1 for refining the partitions  $\{\mathcal{P}_i\}_{i \in [n]}$  can greatly impact the number of iterations and time for convergence. We detail an adaptive partition refinement strategy proposed in the literature [56] in Section 3.1, and introduce the concept of strong partitioning in Section 4 to optimally specify variable partitions.

We further relax (PMR) by outer-approximating the convex quadratic terms in equation (5b) to obtain the following MILP relaxation for strong partitioning. Note that for purely bilinear problems (i.e.,  $|\mathcal{Q}| = 0$ ), problem (PMR) is already an MILP.

$$\begin{aligned} \underline{v}(P) := \min_{(x,w) \in F, \lambda \geq 0, y \in Y} \quad & c^T x + d^T w & (\text{PMR-OA}) \\ \text{s.t.} \quad & (x_i, x_j, w_{ij}, \lambda^{ij}, y^i, y^j) \text{ satisfies (4a) - (4d), } \forall (i, j) \in \mathcal{B}, \\ & (x_k, w_{kk}, \lambda^k, y^k) \text{ satisfies (5a) and (5c), } \forall k \in \mathcal{Q}, \\ & \sum_{l=1}^{d+2} \lambda_l^k = 1, \quad w_{kk} \geq 2\alpha_j^k x_k - (\alpha_j^k)^2, \quad \forall j \in \mathcal{J}_k, \quad k \in \mathcal{Q}. \end{aligned} \quad (6)$$

We explicitly indicate the dependence of the piecewise McCormick lower bound  $\underline{v}$  on the  $d \times n$  matrix of partitioning points  $P := (p^1, p^2, \dots, p^n)$ , excluding points  $p_0^i := 0$  and  $p_{d+1}^i := 1$ . Constraints (6) outer-approximate the inequalities  $w_{kk} \geq x_k^2$  in equation (5b) at the points  $\{\alpha_j^k\}_{j \in \mathcal{J}_k} \subset [0, 1]$  (we assume



---

**Algorithm 1** Partitioning algorithm for the global optimization of (QCQP)

---

- 1: **Input:** relative optimality tolerance  $\varepsilon_r > 0$
  - 2: **Initialization:** partitions  $\mathcal{P}_i^0 = [0, 1]$ ,  $i \in [n]$ , best found solution  $\{\hat{x}\} = \emptyset$  with objective  $UBD = +\infty$ , lower bound  $LBD = -\infty$ , iteration number  $l = 0$
  - 3: Solve (QCQP) locally. Update incumbent  $\hat{x}$  and upper bound  $UBD$  if relevant
  - 4: Solve the termwise McCormick relaxation (2). Update lower bound  $LBD$
  - 5: **while**  $\frac{UBD - LBD}{|UBD| + 10^{-6}} > \varepsilon_r$  **or**  $UBD = +\infty$  **do**
  - 6:   Update  $l \leftarrow l + 1$
  - 7:   Refine partitions  $\{\mathcal{P}_i^{l-1}\}_{i \in [n]}$  to  $\{\mathcal{P}_i^l\}_{i \in [n]}$  by adding partitioning points
  - 8:   Solve piecewise McCormick relaxation (PMR) with partitions  $\{\mathcal{P}_i^l\}_{i \in [n]}$ . Update lower bound  $LBD$
  - 9:   Solve (QCQP) locally. Update incumbent  $\hat{x}$  and upper bound  $UBD$  if needed
  - 10: **end while**
  - 11: Return  $\varepsilon_r$ -optimal solution  $\hat{x}$ , upper bound  $UBD$ , and lower bound  $LBD$
- 

$\{p_0^k, p_1^k, \dots, p_{d+1}^k\} \subset \{\alpha_j^k\}$ . We only use the outer-approximation (PMR-OA) for strong partitioning and revert to solving problem (PMR) while computing lower bounds (see Algorithm 1).

In preparation for Section 4, we recast (PMR-OA) in the following abstract form for suitably defined vectors  $\bar{b}$  and  $\bar{c}$ , matrix  $\bar{B}$ , matrix-valued function  $M$  with co-domain  $\mathbb{R}^{n_r \times n_c}$ , and variables  $z$  (that includes  $x$ ,  $w$ ,  $\lambda$ , and slack variables):

$$\begin{aligned} \underline{v}(P) := \min_{y \in Y} v(P, y), \quad (7) \quad & v(P, y) := \min_{z \geq 0} \bar{c}^\top z \quad (8) \\ & \text{s.t. } M(P, y)z = \bar{B}y + \bar{b}. \end{aligned}$$

We omit in problem (8) the third set of constraints in equation (5a) for simplicity because they only strengthen the LP relaxation of (PMR-OA) and are redundant for the piecewise McCormick relaxations  $\mathcal{PMR}_k^Q$ . Note that for any  $y \in Y$ , at most four of the  $\lambda^{ij}$  variables in the formulation of each  $\mathcal{PMR}_{ij}^B$  and at most two of the  $\lambda^k$  variables in the formulation of each  $\mathcal{PMR}_k^Q$  may be nonzero. Consequently, for each  $y \in Y$ , we eliminate the variables and equations corresponding to the  $\lambda$  variables that are fixed to zero and let  $M(P, y)$  denote the coefficient matrix of the remaining equations (the coefficients of the matrix  $M(P, y)$  themselves do not depend on  $y$ ).

### 3.1 Adaptive partitioning strategy

Algorithm 2 outlines the adaptive partitioning strategy proposed in Nagarajan et al. [56] that is implemented within the solver Alpine. This adaptive strategy empirically performs well on numerous test instances [56] and is motivated by the fact that uniformly partitioning variable domains (proposed, e.g., in [19, 62, 68]) creates many partitions that do not contribute significantly to improving the piecewise McCormick lower bound. Instead of selecting a subset of the variables participating in nonconvex terms for partitioning as in [56], Algorithm 2 partitions the domains of all variables in  $\mathcal{NC}$ . This is because partitioning only a subset of variables in  $\mathcal{NC}$  may result in an analogue of the so-called cluster problem in reduced-space global optimization [35, 36], potentially resulting in a significant increase in the number of iterations for convergence.

At each iteration  $l \in \mathbb{N}$ , Algorithm 2 adds (up to) two partitioning points per variable around a reference point  $\bar{x}^{l-1}$ . At the first iteration,  $\bar{x}^0$  is either set to the feasible local solution from presolve if one is found, or to a solution to the termwise McCormick relaxation (2) otherwise. At subsequent iterations  $l > 1$ ,  $\bar{x}^{l-1}$  is specified as the  $x$ -component of a solution to the piecewise McCormick relaxation lower bounding problem (PMR) at iteration  $l - 1$ . Relative to the solution  $\bar{x}^{l-1}$ , the  $j$ th partition  $[\hat{p}_{j-1}^i, \hat{p}_j^i]$  of variable  $x_i$  is said to be active at iteration  $l - 1$  (i.e.,  $\mathcal{A}(i, l - 1) = j$ ) if  $\hat{p}_{j-1}^i \leq \bar{x}_i^{l-1} \leq \hat{p}_j^i$ , or, equivalently, if there exists

---

**Algorithm 2** Adaptive partition refinement strategy at iteration  $l$ 


---

- 1: **Input:** parameter  $\Delta \geq 4$  and  $\forall i \in [n]$ : partition  $\mathcal{P}_i^{l-1} := [\hat{p}_0^i, \hat{p}_1^i, \dots, \hat{p}_{m_{l-1},i}^i]$ , index  $\mathcal{A}(i, l-1) \in [m_{l-1}, i]$  of the active partition at iteration  $l-1$ , and reference point  $\bar{x}_i^{l-1}$  in the active partition, i.e., the point  $\bar{x}_i^{l-1} \in [\hat{p}_{\mathcal{A}(i,l-1)-1}^i, \hat{p}_{\mathcal{A}(i,l-1)}^i]$
- 2: **Initialization:** partition  $\mathcal{P}_i^l = \mathcal{P}_i^{l-1}$ ,  $i \in [n]$
- 3: **for**  $i \in \mathcal{NC}$  **do**
- 4:     Refine  $\mathcal{P}_i^{l-1}$  to  $\mathcal{P}_i^l$  as follows:

$$\mathcal{P}_i^l = \left[ \hat{p}_0^i, \dots, \hat{p}_{\mathcal{A}(i,l-1)-1}^i, \max \left\{ \hat{p}_{\mathcal{A}(i,l-1)-1}^i, \bar{x}_i^{l-1} - \frac{\text{width}(\mathcal{A}(i, l-1))}{\Delta} \right\}, \right. \\ \left. \min \left\{ \hat{p}_{\mathcal{A}(i,l-1)}^i, \bar{x}_i^{l-1} + \frac{\text{width}(\mathcal{A}(i, l-1))}{\Delta} \right\}, \hat{p}_{\mathcal{A}(i,l-1)}^i, \dots, \hat{p}_{m_{l-1},i}^i \right],$$

where  $\text{width}(\mathcal{A}(i, l-1)) := \hat{p}_{\mathcal{A}(i,l-1)}^i - \hat{p}_{\mathcal{A}(i,l-1)-1}^i$

5: **end for**

6: **Output:** refined partitions  $\{\mathcal{P}_i^l\}_{i \in [n]}$

---

an optimal solution to (PMR) such that  $y_j^i = 1$ . The parameter  $\Delta$  (default value = 10) in Algorithm 2 is a dimensionless scaling factor for the size of the partition constructed around the reference point  $\bar{x}^{l-1}$ . Larger values of  $\Delta$  result in a narrower partition around  $\bar{x}^{l-1}$ .

In the next section, we propose strong partitioning to specify the partitions  $\{\mathcal{P}_i^1\}_{i \in [n]}$  in the first iteration instead of using the heuristic in Algorithm 2. If *UBD* and the lower bound *LBD* obtained using strong partitioning are not converged, we revert to using the adaptive partitioning strategy in Algorithm 2 to specify the partitions  $\{\mathcal{P}_i^l\}_{i \in [n]}$  in the subsequent iterations  $l \geq 2$  of our partitioning algorithm.

## 4 Strong partitioning for nonconvex QCQPs

The choice of partitioning points in the initial iterations can greatly impact the strength of lower bounds, number of iterations for convergence, and overall solution time. While there are some motivations for the adaptive partitioning strategy in Section 3.1, it is still ad hoc for a few reasons: it uses the same parameter  $\Delta$  to partition the domains of *all* variables, and it only considers symmetric partitions around the reference point  $\bar{x}^{l-1}$ . The quality of the partitions  $\{\mathcal{P}_i^1\}_{i \in [n]}$  in the first iteration also depend on the quality of the feasible solution determined during presolve, with sub-optimal presolve solutions potentially leading to sub-optimal initial partitions and slow convergence overall. We propose *strong partitioning* (SP) to address the above limitations of Algorithm 2.

The concept of strong partitioning is akin to strong branching in B&B algorithms for MI(N)LPs. Strong branching for MILPs only chooses the branching variable (a discrete choice) at a node to maximize some function of the lower bound improvements at its two children nodes. Strong partitioning, on the other hand, chooses partitioning points for each partitioned variable (continuous choices within the variable domains) such that the resulting piecewise McCormick relaxation lower bound is maximized. It can be formulated as the following max-min problem:

$$P^* \in \arg \max_{P \in \mathfrak{P}} \underline{v}(P), \tag{SP}$$

where  $\underline{v}(P)$  is the value function of (PMR-OA) and the set  $\mathfrak{P}$  is defined as

$$\mathfrak{P} := \{P := (p^1, p^2, \dots, p^n) \in \mathbb{R}_+^{d \times n} : 0 \leq p_1^i \leq p_2^i \leq \dots \leq p_d^i \leq 1, \forall i \in [n]\}.$$

The strong partitioning problem (SP) is challenging to solve even to local optimality because the inner-minimization problem (PMR-OA) includes binary decisions and its feasible region depends on  $P$  (variables



of the outer-maximization). While (SP) can be formulated as a generalized semi-infinite program [67], state-of-the-art global optimization algorithms for this problem class do not scale even for moderate problem dimensions. Therefore, we design a local optimization method for (SP) with the hope of determining partitioning points  $\bar{P} \in \mathfrak{P}$  that yield a tight lower bound  $\underline{v}(\bar{P})$ . We use the local solution of (SP) to specify the initial partitions  $\{\mathcal{P}_i^1\}_{i \in [n]}$  within Algorithm 1. If the resulting bounds  $LBD$  and  $UBD$  are not converged after one iteration of Algorithm 1, we use Algorithm 2 to specify the partitions  $\{\mathcal{P}_i^l\}_{i \in [n]}$  from the second iteration.

We use generalized gradients of the value function of the inner-minimization (PMR-OA) within a bundle solver for nonsmooth nonconvex optimization to solve problem (SP) to local optimality. Although the value function of an MILP might be discontinuous in general, (PMR-OA) possesses special structure because (outer-approximations of) piecewise McCormick relaxations are nonconvex piecewise-linear continuous functions (cf. Figure 1). This allows for the computation of sensitivity information in this setting. The bundle solver, MPBNGC [50], that we use, requires function and generalized gradient evaluations at points  $P \in \mathfrak{P}$  during the course of its algorithm. Each function evaluation  $\underline{v}(P)$  requires the solution of the MILP (PMR-OA). Under suitable assumptions, a generalized gradient  $\partial_P \underline{v}(P)$  can be obtained by fixing  $y$  to an optimal  $y$ -solution of (PMR-OA) and computing a generalized gradient of the resulting LP (8) using parametric sensitivity theory [22]. We formalize these details in the next section. Before we proceed, we include the convergence guarantees of MPBNGC [50] below for the sake of completeness.

**Definition 1.** Let  $Z \subset \mathbb{R}^N$  be open. A locally Lipschitz function  $f : Z \rightarrow \mathbb{R}$  is said to be weakly semismooth if the directional derivative  $f'(z, d) = \lim_{t \downarrow 0} \frac{f(z+td) - f(z)}{t}$  exists for all  $z \in Z$ ,  $d \in \mathbb{R}^N$  and  $f'(z, d) = \lim_{t \downarrow 0} \xi(z+td)^T d$  for  $\xi(z+td) \in \partial f(z+td)$ .

**Definition 2.** Let  $f : \mathbb{R}^N \rightarrow \mathbb{R}$ ,  $g : \mathbb{R}^N \rightarrow \mathbb{R}^M$  be locally Lipschitz continuous. Consider the problem  $\min_{z: g(z) \leq 0} f(z)$ . A feasible point  $z^*$  is said to be substationary if there exist multipliers  $\lambda \geq 0$  and  $\mu \in \mathbb{R}_+^M$ , with  $(\lambda, \mu) \neq (0, 0)$ , such that

$$0 \in \lambda \partial f(z^*) + \sum_{j=1}^M \mu_j \partial g_j(z^*), \quad \mu_j g_j(z^*) = 0, \quad \forall j \in [M].$$

**Theorem 1.** Suppose the function  $\underline{v}$  is weakly semismooth. Then MPBNGC either terminates finitely with a substationary point to (SP), or any accumulation point of a sequence of MPBNGC solutions is a substationary point to (SP).

*Proof.* See Theorem 9 of [50]. □

The example below shows the value function  $\underline{v}$  of (PMR-OA) may be nonsmooth.

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


$$\min_{x \in [0,1]} x \quad \text{s.t.} \quad x^2 \geq (0.4)^2.$$

The optimal solution is  $x^* = 0.4$  with optimal value  $v^* = 0.4$ . Suppose we wish to partition the domain of  $x$  into two sub-intervals (i.e.,  $d = 1$ ). Let  $\mathcal{P} = [0, p, 1]$  denote the partition of  $x$  with  $0 \leq p \leq 1$ . After some algebraic manipulation, the outer-approximation problem (PMR-OA) can be reformulated as

$$\underline{v}(p) = \min_{x \in [0,1]} x \quad \text{s.t.} \quad w \geq (0.4)^2, \quad w \leq \max\{px, (1+p)x - p\}, \quad w \geq 2\alpha_j x - \alpha_j^2, \quad \forall j \in \mathcal{J},$$

where  $\{\alpha_j^k\}_{j \in \mathcal{J}} \subset [0, 1]$  and we write  $\underline{v}(p)$  to indicate the dependence on the partitioning point  $p$ . We can derive the piecewise McCormick lower bound to be

$$\underline{v}(p) = \begin{cases} \frac{0.16+p}{1+p}, & \text{if } 0 \leq p \leq 0.4 \\ \frac{0.16}{p}, & \text{if } 0.4 < p \leq 1 \end{cases},$$



which shows that  $\underline{v}$  is continuous and piecewise differentiable at  $p = 0.4$ .

#### 4.1 Computing generalized gradients of $\underline{v}$

We identify conditions under which a generalized gradient of the value function  $\underline{v}$  may be computed in practice (we consider generalized gradients of  $\underline{v}$  on  $\mathfrak{P}$  instead of the open set  $\text{int}(\mathfrak{P})$  for simplicity). We begin with the following useful result; its assumption that the  $y$ -solution of (7) is unique can be verified by adding a “no-good cut” and re-solving (7) to check if the second-best solution for  $y$  has a strictly greater objective than  $\underline{v}(P)$ .

**Lemma 2.** Suppose problem (7) has a unique  $y$  solution  $y^* \in Y$  at  $P \in \mathfrak{P}$  and  $v(\cdot, y^*)$  is continuous at  $P$ . Then  $\underline{v}(\tilde{P}) = v(\tilde{P}, y^*)$ ,  $\forall \tilde{P} \in \mathfrak{P}$  in a neighborhood of  $P$ .

*Proof.* Because  $y^*$  is the unique  $y$  solution to (7) at  $P \in \mathfrak{P}$ ,  $v(P, y^*) < v(P, y)$ ,  $\forall y \in Y \setminus \{y^*\}$ . To see  $\underline{v}(\cdot) \equiv v(\cdot, y^*)$  in a neighborhood of  $P$ , we show that the value function  $v(\cdot, y)$  is lower semicontinuous on  $\mathfrak{P}$  for each  $y \in Y$ . The stated result then holds since  $v(\cdot, y^*)$  is assumed to be continuous at  $P$ .

The set-valued mapping  $P \in \mathfrak{P} \mapsto \{z \geq 0 : M(P, y) = \bar{B}y + \bar{b}\}$  is locally compact for each  $y \in Y$  by virtue of the continuity of the mapping  $M(\cdot, y)$  and the finite bounds that can be deduced on all of the variables in problem (PMR-OA). Lemma 5.3 of Still [64] then implies that  $v(\cdot, y)$  is lower semicontinuous on  $\mathfrak{P}$  for each  $y \in Y$ .  $\square$

The next result characterizes the gradient of  $\underline{v}$  in the non-degenerate case.

**Theorem 3.** Suppose  $P \in \mathfrak{P}$  and problem (7) has a unique  $y$  solution  $y^* \in Y$ . Consider the LP (8) with  $y$  fixed to  $y^*$ . If this LP has a unique primal solution  $z^*$  and a unique dual solution  $\pi^*$ , then

$$\frac{\partial \underline{v}}{\partial p_j^i}(P) = \frac{\partial v}{\partial p_j^i}(P, y^*) = \sum_{k=1}^{n_r} \sum_{l=1}^{n_c} \pi_k^* z_l^* \frac{\partial M_{kl}}{\partial p_j^i}(P, y^*), \quad \forall i \in [n], j \in [d].$$

*Proof.* Lemma 2 implies  $\underline{v}(\cdot) \equiv v(\cdot, y^*)$  in a neighborhood of  $P$  provided  $v(\cdot, y^*)$  is continuous at  $P$ . Theorem 1 of Freund [26] (cf. Proposition 4.1 of [22]) and the fact that the function  $M(\cdot, y^*)$  is continuously differentiable on  $\mathfrak{P}$  together imply  $v(\cdot, y^*)$  is continuously differentiable at  $P$  and the stated equalities hold.  $\square$

Next, we derive a formula for the generalized gradient  $\partial_P \underline{v}(P)$  when the assumption that the LP (8) is non-degenerate fails to hold.

**Theorem 4.** Suppose  $P \in \mathfrak{P}$  and problem (7) has a unique  $y$  solution  $y^* \in Y$ . Consider the LP (8) with  $y$  fixed to  $y^*$ . Suppose  $v(\cdot, y^*)$  is finite and locally Lipschitz in a neighborhood of  $P$ . Then

$$\partial_P \underline{v}(P) = \partial_P v(P, y^*) \subset \text{conv} \left( \left\{ \sum_{k=1}^{n_r} \sum_{l=1}^{n_c} \pi_k^* x_l^* \partial_P M_{kl}(P) : (x^*, \pi^*) \text{ is a primal-dual optimal pair for (8)} \right\} \right).$$

*Proof.* Lemma 2 implies  $\underline{v}(\cdot) \equiv v(\cdot, y^*)$  in a neighborhood of  $P$ . The stated equalities hold by mirroring the proof of Theorem 5.1 of De Wolf and Smeers [22] and noting that the function  $M(\cdot, y^*)$  is continuously differentiable on  $\mathfrak{P}$ .  $\square$

De Wolf and Smeers [22] (see Assumption 5.1) and Im [33] argue the following result with  $\bar{y} = y^*$  ensures  $v(\cdot, y^*)$  is locally Lipschitz in a neighborhood of  $P \in \mathfrak{P}$ .

**Lemma 5.** Suppose  $P \in \mathfrak{P}$  and  $\bar{y} \in Y$ . Consider the LP (8) with  $y$  fixed to  $\bar{y}$ . If the matrix  $M(P, \bar{y})$  has full row rank and  $\bar{B}\bar{y} + \bar{b} \in \text{int}(\{M(P, \bar{y})z : z \geq 0\})$ , then  $v(\cdot, \bar{y})$  is finite and locally Lipschitz in a neighborhood of  $P$ .

*Proof.* See Proposition 5.3 of [22] and pages 73 to 76 of [33].  $\square$

We now verify that the full rank assumption in Lemma 5 holds in general.

**Lemma 6.** The matrix  $M(P, y)$  has full row rank,  $\forall P \in \text{int}(\mathfrak{P})$  and  $y \in Y$ .

*Proof.* Fix  $y \in Y$ . Since  $P \in \text{int}(\mathfrak{P})$ , we have  $0 < p_1^i < p_2^i < \dots < p_d^i < 1$  for each  $i \in [n]$ . We show that for each  $(i, j) \in \mathcal{B}$  and  $k \in \mathcal{Q}$ , the equality constraints in equations (4a)-(4d) and equations (5a)-(5c) have full row rank, which readily imply that  $M(P, y)$  has full row rank. We ignore inequality constraints because they are transformed into equality constraints by the addition of unique slack variables.

We begin by focusing on the equality constraints in (4a)-(4d) involving the  $x$ ,  $w$ , and  $\lambda$  variables. Consider a fixed  $(i, j) \in \mathcal{B}$ . Since at most four of the  $\lambda^{ij}$  variables may be nonzero, we can rewrite these equality constraints as follows after a change of variables (here,  $\mathcal{A}(i)$  denotes the active partition of  $x_i$ ):

$$\begin{pmatrix} -1 & 0 & 0 & p_{\mathcal{A}(i)-1}^i & p_{\mathcal{A}(i)-1}^i & p_{\mathcal{A}(i)}^i & p_{\mathcal{A}(i)}^i \\ 0 & -1 & 0 & p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(j)}^j & p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(j)}^j \\ 0 & 0 & -1 & p_{\mathcal{A}(i)-1}^i p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(i)-1}^i p_{\mathcal{A}(j)}^j & p_{\mathcal{A}(i)}^i p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(i)}^i p_{\mathcal{A}(j)}^j \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_i \\ x_j \\ w_{ij} \\ \lambda_1^{ij} \\ \lambda_2^{ij} \\ \lambda_3^{ij} \\ \lambda_4^{ij} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}.$$

We argue that the following sub-matrix is of full rank whenever  $P \in \text{int}(\mathfrak{P})$ :

$$\begin{pmatrix} p_{\mathcal{A}(i)-1}^i & p_{\mathcal{A}(i)-1}^i & p_{\mathcal{A}(i)}^i & p_{\mathcal{A}(i)}^i \\ p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(j)}^j & p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(j)}^j \\ p_{\mathcal{A}(i)-1}^i p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(i)-1}^i p_{\mathcal{A}(j)}^j & p_{\mathcal{A}(i)}^i p_{\mathcal{A}(j)-1}^j & p_{\mathcal{A}(i)}^i p_{\mathcal{A}(j)}^j \\ 1 & 1 & 1 & 1 \end{pmatrix}.$$

Subtracting the first column from the second column, the third from the fourth column, and finally the first from the third column yields the column vectors

$$\begin{aligned} & (p_{\mathcal{A}(i)-1}^i, p_{\mathcal{A}(j)-1}^j, p_{\mathcal{A}(i)-1}^i p_{\mathcal{A}(j)-1}^j, 1), \\ & (0, (p_{\mathcal{A}(j)}^j - p_{\mathcal{A}(j)-1}^j), p_{\mathcal{A}(i)-1}^i (p_{\mathcal{A}(j)}^j - p_{\mathcal{A}(j)-1}^j), 0), \\ & ((p_{\mathcal{A}(i)}^i - p_{\mathcal{A}(i)-1}^i), 0, p_{\mathcal{A}(j)-1}^j (p_{\mathcal{A}(i)}^i - p_{\mathcal{A}(i)-1}^i), 0), \\ & (0, (p_{\mathcal{A}(j)}^j - p_{\mathcal{A}(j)-1}^j), p_{\mathcal{A}(i)}^i (p_{\mathcal{A}(j)}^j - p_{\mathcal{A}(j)-1}^j), 0). \end{aligned}$$

It is easy to see these vectors are linearly independent if  $0 < p_{\mathcal{A}(i)-1}^i < p_{\mathcal{A}(i)}^i < 1, \forall i$ .

Next, we focus on the equality constraints in (5a)-(5c) involving the  $x$ ,  $w$ , and  $\lambda$  variables for a fixed  $k \in \mathcal{Q}$ . Since at most two of the  $\lambda^k$  variables may be nonzero, we can rewrite these equality constraints as follows after a change of variables:

$$\begin{pmatrix} -1 & p_{\mathcal{A}(k)-1}^k & p_{\mathcal{A}(k)}^k \\ 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_k \\ \lambda_1^k \\ \lambda_2^k \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The last two matrix columns are linearly independent if  $0 < p_{\mathcal{A}(k)-1}^k < p_{\mathcal{A}(k)}^k < 1$ .  $\square$

Finally, we show that for almost every (a.e.)  $P \in \mathfrak{P}$  with respect to the uniform measure, problem (7) either has a unique  $y$  solution (which ensures that we can use Theorem 3 or 4 to compute a generalized of  $v$  under mild conditions), or  $v(P) = v^*$  (i.e., the optimal values of (QCQP) and problem (7) are equal, which implies that partitioning points  $P$  are sufficient for convergence of the lower bound), or both.

**Theorem 7.** At least one of the following holds for a.e.  $P \in \mathfrak{P}$ :

1.  $v(P) = v^*$ ,
2. Problem (7) has a unique  $y$  solution.

*Proof.* Consider  $P \in \mathfrak{P}$ , and let  $\hat{y} \in Y$  and  $\hat{x} \in [0, 1]^n$  denote the  $y$  and  $x$  components of an optimal solution to problem (7). We consider the following cases:

- (a) For each  $k \in \mathcal{Q}$ , we have  $\hat{x}_k \in \{p_0^k, p_1^k, \dots, p_{d+1}^k\}$ . Additionally, for each  $(i, j) \in \mathcal{B}$ , we either have  $\hat{x}_i \in \{p_0^i, p_1^i, \dots, p_{d+1}^i\}$ , or  $\hat{x}_j \in \{p_0^j, p_1^j, \dots, p_{d+1}^j\}$ , or both.
- (b) Case (a) does not hold, i.e., there either exists at least one index  $k \in \mathcal{Q}$  such that  $\hat{x}_k \notin \{p_0^k, p_1^k, \dots, p_{d+1}^k\}$ , or there exists at least one pair of indices  $(i, j) \in \mathcal{B}$  such that both  $\hat{x}_i \notin \{p_0^i, p_1^i, \dots, p_{d+1}^i\}$  and  $\hat{x}_j \notin \{p_0^j, p_1^j, \dots, p_{d+1}^j\}$ .

Suppose case (a) holds. Since we assume that  $\{p_0^k, p_1^k, \dots, p_{d+1}^k\} \subset \{\alpha_j^k\}$  for each  $k \in \mathcal{Q}$ , our outer-approximation of the piecewise McCormick relaxation (5a)-(5c) for the quadratic constraint  $w_{kk} = x_k^2$  is exact (i.e., there is no relaxation gap) at the partitioning points  $x_k \in \{p_0^k, p_1^k, \dots, p_{d+1}^k\}$ . Additionally, the piecewise McCormick relaxation (4a)-(4d) for the bilinear constraint  $w_{ij} = x_i x_j$  is exact either when  $x_i \in \{p_0^i, p_1^i, \dots, p_{d+1}^i\}$ , or when  $x_j \in \{p_0^j, p_1^j, \dots, p_{d+1}^j\}$ , or both. Therefore, the point  $\hat{x}$  is feasible to the original QCQP (1), which implies  $v(P) = v^*$ .

Suppose instead that case (b) holds. Additionally, suppose there exist multiple  $y$  solutions to problem (7). Let  $\tilde{y} \in Y$  and  $\tilde{x} \in [0, 1]^n$  denote the  $y$  and  $x$  components of *another* optimal solution to problem (7) with  $\tilde{y} \neq \hat{y}$ . Since case (a) does not hold and  $\tilde{y} \neq \hat{y}$ , we have  $\tilde{x} \neq \hat{x}$ . Moreover, there exist non-singular basis matrices  $\hat{M}(P, \hat{y})$  and  $\tilde{M}(P, \tilde{y})$  for the LPs (8) corresponding to  $\hat{y}$  and  $\tilde{y}$ , respectively, such that

$$v(P) = v(P, \hat{y}) = \hat{c}^T [\hat{M}(P, \hat{y})]^{-1} (\bar{B}\hat{y} + \bar{b}) = \tilde{c}^T [\tilde{M}(P, \tilde{y})]^{-1} (\bar{B}\tilde{y} + \bar{b}) = v(P, \tilde{y}) \quad (9)$$

for suitable vectors  $\tilde{c}$  and  $\hat{c}$ , which only include the components of  $\bar{c}$  corresponding to the basic variables of these LPs. Since not all components of  $\hat{x}$  equal 0 or 1, at least some of the entries of  $\hat{M}(P, \hat{y})$  are functions of the partitioning points  $P$ . Moreover,  $v(P, \hat{y})$  and  $v(P, \tilde{y})$  are not identical functions of  $P$  since  $\tilde{x} \neq \hat{x}$ . Equation (9) thus yields a polynomial equation in  $P$ . Therefore, the set of all  $P \in \mathfrak{P}$  such that (9) holds has measure zero. Noting that  $|Y| < +\infty$  and the number of possible bases is finite for each  $y \in Y$  concludes the proof.  $\square$

## 4.2 Algorithmic enhancements

We design preprocessing and postprocessing steps that can be used to mitigate the computational burden of solving (SP) locally and enable our ML model to more effectively learn its solution.

The outer-maximization in problem (SP) involves  $n \times d$  partitioning variables. Since larger problem dimensions may increase both the per-iteration cost and number of iterations taken by the bundle solver to converge, we propose preprocessing heuristics to fix a subset of the partitioning points  $P$  and to compute an initial guess  $P^0$  for the bundle method. After solving the max-min problem (SP) (line 18), we propose postprocessing steps to eliminate partitioning points in its solution  $\bar{P}$  that do not significantly affect the lower bound  $v(\bar{P})$ . Algorithm 3 includes detailed pseudocode of our preprocessing (lines 1–17) and postprocessing steps (lines 19–28).

---

**Algorithm 3** Preprocessing and postprocessing steps for strong partitioning

---

**Preprocessing steps**

- 1: **Initialization:** partitions  $\hat{\mathcal{P}}_i^0 := [0, 1], \forall i \in [n]$
- 2: Solve the McCormick relaxation (2) to compute a lower bounding solution  $\bar{x}^0$
- 3: **for**  $k = 1, 2, \dots, d$  **do**
- 4:     **for**  $i = 1, 2, \dots, n$  **do**
- 5:         **if**  $i \notin \mathcal{NC}$  **or**  $\bar{x}_i^{k-1} \approx \tilde{x}_i$  for some  $\tilde{x}_i \in \hat{\mathcal{P}}_i^{k-1}$  **then**
- 6:             Set  $\hat{\mathcal{P}}_i^k = \hat{\mathcal{P}}_i^{k-1}$
- 7:         **else**
- 8:             Insert  $\bar{x}_i^{k-1}$  in  $\hat{\mathcal{P}}_i^{k-1}$  to obtain  $\hat{\mathcal{P}}_i^k$
- 9:         **end if**
- 10:     **end for**
- 11:     Solve (PMR-OA) with partitions  $\{\hat{\mathcal{P}}_i^k\}_{i \in [n]}$  to determine solution  $\bar{x}^k$
- 12: **end for**
- 13: Let  $n_i := |\hat{\mathcal{P}}_i^d| - 2, \forall i \in [n]$
- 14: Let  $[0, p_{d-n_i+1}^{i0}, p_{d-n_i+2}^{i0}, \dots, p_d^{i0}, 1]$  denote  $\hat{\mathcal{P}}_i^d$  and set  $p_j^{j0} := 0, \forall j \in [d - n_i]$
- 15: Set the initial guess for (SP) to  $P^0$ , where  $P_{ji}^0 := p_j^{j0}, \forall i \in [n], j \in [d]$
- 16: Fix variables  $p_j^i, j \in [d - n_i]$ , to 0 while solving (SP)
- 17: **Output:** initial guess (and variable fixings)  $P^0$  for problem (SP)
- 18: Solve max-min problem (SP) to obtain a solution  $\bar{P} \in \mathfrak{P}$  with objective  $\bar{v} := \underline{v}(\bar{P})$

**Postprocessing steps**

- 19: **for**  $j = 1, 2, \dots, d$  **do**
  - 20:     **for**  $i \in \mathcal{NC}$  **do**
  - 21:         Set  $\hat{P} = \bar{P}$  and replace the element  $\hat{P}_{ji}$  with zero
  - 22:         Solve (PMR-OA) with partitioning points  $\hat{P}$  to obtain bound  $\hat{v} := \underline{v}(\hat{P})$
  - 23:         **if**  $\hat{v} \leq \bar{v} + 10^{-6}|\bar{v}|$  **then**
  - 24:             Update  $\bar{P} = \hat{P}$  and sort it such that  $\bar{P}_{ki} \leq \bar{P}_{(k+1)i}, \forall k \in [d - 1]$
  - 25:         **end if**
  - 26:     **end for**
  - 27: **end for**
  - 28: **Output:** partitions  $\{\mathcal{P}_i^1\}_{i \in [n]}$  for Algorithm 1 based on postprocessed  $\bar{P}$
- 

## 5 Numerical experiments

We study the impact of using strong partitioning to specify Alpine’s partitions at the first iteration and investigate an off-the-shelf ML model for learning these partitions for homogeneous QCQPs. We begin by describing the setup for our computational experiments. In Section 5.1, we outline the procedure for generating families of random QCQP instances, including instances of the pooling problem. We detail our ML approximation of strong partitioning in Section 5.2, and compare the performance of strong partitioning and its ML approximation against Alpine’s default partitioning strategy (see Section 3.1 for details) in Section 5.3.

Our strong partitioning code is written in Julia 1.6.3 and implemented within Alpine.jl v0.4.1<sup>1</sup>. We use JuMP.jl v1.1.1 and use Gurobi 9.1.2 via Gurobi.jl v0.11.3 for solving LPs, MILPs, and convex MIQCQPs (with MIPGap =  $10^{-6}$ ). We use either Ipopt 3.14.4 via Ipopt.jl v1.0.3 (with max\_iter =  $10^4$ ), or Artelys

<sup>1</sup><https://github.com/lanl-ansi/Alpine.jl>

Knitro 12.4.0 via KNITRO.jl v0.13.0 (with `algorithm = 3`) to solve NLPs locally within Alpine<sup>2</sup>. We use the bundle solver MPBNGC 2.0 [50] via MPBNGCInterface.jl<sup>3</sup> (with options `OPT_LMAX = 20`, `OPT_EPS = 10-9`, and `OPT_NITER = OPT_NFASG = 500`) to solve the max-min problem (SP) to local optimality. We consider strong partitioning with either two or four partitioning points per partitioned variable ( $d = 2$  or  $d = 4$ ) in addition to the variable bounds and use scikit-learn v0.23.2 [60] to design its ML approximation. To demonstrate the non-trivial nature of our nonconvex test instances, we also solve them to global optimality using BARON 22.11.3 via BARON.jl v0.8.0 and provide BARON with the option of using CPLEX 22.1.0 as an MILP solver.

All of our experiments were run on nodes of the Darwin cluster at LANL with dual socket Intel Broadwell 18-core processors (E5-2695\_v4 CPUs, base clock rate at 2.1GHz), EDR InfiniBand, and 125GB of memory. Each instance was run exclusively on a single node and different solution approaches were run in sequence to limit the impact of variability in machine performance. All Alpine and BARON runs were given a time limit of 2 hours with target relative and absolute optimality gaps of  $10^{-4}$  and  $10^{-9}$ , respectively<sup>4</sup>. No time limit was specified for solving the max-min problem (SP). The rest of BARON’s options, including range reduction options, were kept to default. We deactivate bounds tightening techniques within Alpine because it is largely ineffective for our medium and large-scale instances (our approaches are easily adapted to the setting where bounds tightening is employed). We partition the domains of all variables participating in nonconvex terms within Alpine, and set the rest of Alpine’s options to default, including the partition scaling factor to  $\Delta = 10$ .

## 5.1 Test instances

We describe how we generate homogeneous families of random QCQPs, including instances of the pooling problem, based on the literature. Scripts for generating the different families of instances can be found at [https://github.com/lanl-ansi/Alpine.jl/tree/master/examples/random\\_QCQPs](https://github.com/lanl-ansi/Alpine.jl/tree/master/examples/random_QCQPs).

### 5.1.1 Random bilinear programs

We consider parametric bilinear programs of the form [7]:

$$\begin{aligned} v(\theta) &:= \min_{x \in [0,1]^n} x^T Q^0(\theta)x + (r^0(\theta))^T x \\ \text{s.t.} \quad &x^T Q^i(\theta)x + (r^i(\theta))^T x \leq b_i, \quad \forall i \in [m_I], \\ &(a^j)^T x = d_j, \quad \forall j \in [m_E], \end{aligned}$$

where  $\theta \in [-1, 1]^{d_\theta}$  are parameters,  $r^k(\theta) \in \mathbb{R}^n$ ,  $k \in \{0\} \cup [m_I]$ ,  $Q^k(\theta) \in \mathbb{R}^{n \times n}$ ,  $k \in \{0\} \cup [m_I]$ , are symmetric but not necessarily positive semi-definite,  $a^j \in \mathbb{R}^n$ ,  $j \in [m_E]$ ,  $b \in \mathbb{R}^{m_I}$ , and  $d \in \mathbb{R}^{m_E}$ .

We generate 1000 instances for each of  $n \in \{10, 20, 50\}$  variables with  $|\mathcal{B}| = \min\{5n, \binom{n}{2}\}$  bilinear terms (we count  $x_i x_j$  and  $x_j x_i$  as the same bilinear term; all instances for a fixed dimension  $n$  have the same set of  $|\mathcal{B}|$  bilinear terms),  $|\mathcal{Q}| = 0$  quadratic terms,  $m_I = n$  bilinear inequalities, and  $m_E = 0.2n$  linear equalities [7]. We let the dimension  $d_\theta = 3 \times (0.2m_I + 1)$  (see below for why we make this choice). The problem data is generated as follows (cf. [7]). All entries of the vectors  $a^j$  and  $d$  are generated i.i.d. from the uniform distribution  $U(-1, 1)$ , and all entries of the vector  $b$  are generated i.i.d. from  $U(0, 100)$ . The components of  $\theta$  are generated i.i.d. from  $U(-1, 1)$ . Each  $Q^k$  and  $r^k$ ,  $k \in \{0, 1, \dots, 0.2m_I\}$ , are of the form:

$$Q^k(\theta) = \bar{Q}^k + \sum_{l=3k+1}^{3k+3} \theta_l \tilde{Q}^{k,l-3k}, \quad r^k(\theta) = \bar{r}^k + \sum_{l=3k+1}^{3k+3} \theta_l \tilde{r}^{k,l-3k}.$$

<sup>2</sup>We switch Alpine’s local solver between Ipopt for the random bilinear and QCQP instances and Knitro for the random pooling instances because Ipopt is less effective for the pooling instances.

<sup>3</sup><https://github.com/milzj/MPBNGCInterface.jl>

<sup>4</sup>Alpine’s definition of relative gap differs slightly from BARON’s definition, see Section 5.3.1.

The nonzero entries of the “nominal matrices”  $\bar{Q}^k$  and “nominal vectors”  $\bar{r}^k$  are generated i.i.d. from  $U(-1, 1)$ . For each tuple  $(i, j) \in \mathcal{B}$ , indices  $k \in \{0, 1, \dots, 0.2m_I\}$  and  $l \in \{1, 2, 3\}$ , we set  $\tilde{Q}_{ij}^{k,l} := \gamma_{ij}^{k,l} \bar{Q}_{ij}^k$ , where  $\gamma_{ij}^{k,l}$  are generated i.i.d. from  $U(0, 0.5)$ . Similarly, for each index  $i \in \{1, \dots, n\}$ ,  $k \in \{0, 1, \dots, 0.2m_I\}$ , and  $l \in \{1, 2, 3\}$ , we set  $\tilde{r}_i^{k,l} := \delta_i^{k,l} \bar{r}_i^k$ , where  $\delta_i^{k,l}$  are generated i.i.d. from  $U(0, 0.5)$ . Since each  $\tilde{Q}^{k,l}$  and  $\tilde{r}^{k,l}$  is a different perturbation of  $\bar{Q}^k$  and  $\bar{r}^k$ , the expansions of  $Q^k$  and  $r^k$  may be motivated using principal components analysis. The nonzero entries of the remaining matrices  $Q^k$  and vectors  $r^k$ ,  $k \in \{0.2m_I + 1, \dots, m_I\}$ , are the same across all 1000 instances and generated i.i.d. from  $U(-1, 1)$ . Finally, the constraint coefficients are re-scaled such that the vectors  $b = d = 1$ . Note that for a fixed dimension  $n$ , each instance is uniquely specified by the parameters  $\theta$  of dimension  $d_\theta = 3(0.2n + 1)$ .

### 5.1.2 Random QCQPs with bilinear *and* univariate quadratic terms

We also generate 1000 random QCQPs with  $|\mathcal{B}| = \min\{5n, \binom{n}{2}\}$  bilinear terms and  $|\mathcal{Q}| = \lfloor 0.25n \rfloor$  univariate quadratic terms for each of  $n \in \{10, 20, 50\}$  variables (all instances for a fixed  $n$  have the same set of bilinear and univariate quadratic terms). The coefficients of quadratic terms in the objective and constraints are generated similarly to the coefficients of bilinear terms in Section 5.1.1. The rest of the model parameters and problem data (including  $\theta$ ) are also generated similarly as in Section 5.1.1.

### 5.1.3 The pooling problem

The pooling problem is a classical example of a bilinear program introduced by Haverly [31]. It has several important applications in process systems engineering, including petroleum refining [34, 69], natural gas production [34, 44], and water treatment network design [11, 53, 62]. Its goal is to blend inputs of differing qualities at intermediate pools to produce outputs that meet quality specifications while satisfying capacity constraints at inputs, pools, and outputs. Solving the pooling problem is in general NP-hard [3].

We consider instances of the pooling problem with 45 inputs, 15 pools, 30 outputs, and a single quality. Each instance has 116 input-output arcs, 71 input-pool arcs, and 53 pool-output arcs, yielding 572 variables and 621 constraints, including 360 linear constraints and 261 bilinear equations (with 124 variables involved in bilinear terms). We use the  $pq$ -formulation of the pooling problem outlined in Section 2 of [49]. Note that unlike the random bilinear instances in Section 5.1.1 where all of the original “ $x$  variables” participate in bilinear terms, only 124 out of the 311 original variables in the pooling model participate in bilinear terms.

We first generate a nominal instance using the “random Haverly” instance generation approach<sup>5</sup> in [49] that puts together 15 perturbed copies of one of Haverly’s pooling instances [31] and adds 150 edges to it. We modify the target output quality concentrations generated by [49] to construct harder instances. For each output  $j$ , we compute the minimum  $c_j^{\min}$  and maximum  $c_j^{\max}$  input concentrations of the quality over the subset of inputs from which there exists a path to output  $j$ . We then specify the lower and upper bound on the quality concentration at output  $j$  to be  $c_j^{\min} + \alpha_j(c_j^{\max} - c_j^{\min})$  and  $c_j^{\min} + \beta_j(c_j^{\max} - c_j^{\min})$ , respectively, where  $\alpha_j \sim U(0.2, 0.4)$  and  $\beta_j \sim U(0.6, 0.8)$  are generated independently. We also rescale the capacities of the inputs, pools, and outputs and the costs of the arcs for better numerical performance. Note that while all variables in the formulation are non-negative, upper bounds on the variables are not necessarily equal to one after rescaling. After constructing a nominal instance using the above procedure, we use it to generate 1000 random pooling instances by randomly perturbing each input’s quality concentration (parameters  $\theta$  for this problem family) by up to 20%, uniformly and independently.

## 5.2 Machine learning approximation of strong partitioning

Because solving the max-min problem (SP) even to local optimality may be time consuming, we propose to learn the SP strategy for homogeneous QCQP instances. We detail our off-the-shelf ML approximation of strong partitioning in this section. Although our ultimate goal is to optimize the ML model so that its predictions yield good performance when they are used to inform Alpine’s partitions at the first iteration, we

<sup>5</sup><https://github.com/poolinginstances/poolinginstances>



Scaled MAE	% Partitioning Points				
	< 0.01	< 0.02	< 0.05	< 0.1	< 0.2
Bilinear $n = 10$	60	75	80	95	100
Bilinear $n = 20$	15	22.5	60	87.5	97.5
Bilinear $n = 50$	31	39	70	94	100
QCQP $n = 10$	65	80	95	100	100
QCQP $n = 20$	35	37.5	77.5	92.5	100
QCQP $n = 50$	56	66	85	99	100
Pooling	65.7	70.9	78.6	89.5	97.2

Table 1: Statistics of scaled MAEs of the out-of-sample predictions of the ML model.

instead choose our ML model solely based on its accuracy of predicting the strong partitioning points. We do so because tuning the hyperparameters of the ML model directly for good performance within Alpine incurs a huge computational expense due to the need to re-evaluate the performance of the ML predictions within Alpine for each choice of the hyperparameters. While the choice of the ML model can have a significant impact on the performance of its predictions within Alpine, we leave the design of more sophisticated ML architectures for future work.

We use scikit-learn’s AdaBoost regressor<sup>6</sup> [27] that implements the “AdaBoost.R2 algorithm” [24] to learn a mapping from each QCQP instance to the strong partitioning points. Our base estimator is a scikit-learn regression tree<sup>7</sup> [16] with maximum depth equal to 25, and we set the maximum number of weak learners for the boosting algorithm to 1000. The rest of scikit-learn’s `AdaBoostRegressor` options are set to default. We use 10-fold *cross-validation* to generate out-of-sample ML predictions for all 1000 QCQP instances in each problem family. Specifically, we randomly split the 1000 instances in each family into 10 folds, use 9 out of 10 folds for training the ML model, predict the strong partitioning points for the omitted fold, and loop through different choices of the omitted fold to generate predictions for all 1000 instances. We *emphasize* that we fit our ML model for prediction accuracy and do not perform much hyperparameter tuning since our ultimate goal is good performance of the ML predictions when used within Alpine.

**ML model inputs and outputs.** The choice of features for the ML model can greatly impact its performance. We use the following problem features as *inputs* to the ML model: i. parameters  $\theta$ , which uniquely parametrize each nonconvex QCQP instance, ii. the best found feasible solution during Alpine’s presolve step (which involves a single local solve), and iii. the McCormick lower bounding solution (obtained by solving a single convex program). Although it is theoretically sufficient to use only the parameters  $\theta$  as features because they uniquely identify each QCQP instance, we also use the features (ii) and (iii) since they are relatively cheap to compute and intuitively can help inform the partitioning strategy. These additional features are also complicated transformations of the instance parameters  $\theta$  that may otherwise be challenging to uncover. The *outputs* of our ML model are the  $d$  partitioning points (excluding variable bounds) for each of the  $n$  partitioned variables, resulting in an output dimension of  $d \times n$ . In contrast with much of the literature on learning for MILPs, we train separate ML models for each family of 1000 instances since both the feature and output dimensions of our ML models depend on the problem dimensions. While we plan to design more advanced ML architectures that can accommodate variable feature and output dimensions as part of future work, we do not consider the need to train a different ML model for each problem family to be a major limitation. This is because decision-makers often care about solving instances of the *same* problem family with only a few varying parameters, which means they only need to train a single ML model with fixed feature and output dimensions for their application.

<sup>6</sup><https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.AdaBoostRegressor.html>

<sup>7</sup><https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeRegressor.html>

Problem Family	BARON Solution Time (seconds)				# TLE	TLE Gap (GM)
	Shifted GM	Median	Min	Max		
Bilinear $n = 10$	0.2	0.2	0.1	0.4	0	—
Bilinear $n = 20$	3.5	3.6	1.4	7.0	0	—
Bilinear $n = 50$	257.1	260.6	54.3	4637.4	0	—
QCQP $n = 10$	0.3	0.3	0.1	0.6	0	—
QCQP $n = 20$	4.8	4.6	1.6	10.8	0	—
QCQP $n = 50$	268.2	246.2	14.7	6897.9	19	$2.3 \times 10^{-2}$
Pooling	441.4	422.1	16.6	7114.5	432	$2.7 \times 10^{-2}$

Table 2: Statistics of BARON solution times, including the shifted geometric mean, median, minimum, and maximum times over the subset of 1000 instances for which BARON does not hit the time limit. The last two columns denote the number of instances for which BARON hits the time limit and the corresponding geometric mean of residual optimality gaps at termination, respectively.

Solution Method	Solution Time (seconds)				# TLE
	Shifted GM	Median	Min	Max	
Alpine (default)	30.7	14.1	0.4	4020.9	2
Alpine+SP2	9.4	1.7	0.2	3864.9	1
Alpine+SP4	5.8	1.4	0.2	3871.3	0

Table 3: (**Benchmark QCQPs**) Statistics of solution times. Columns correspond to the shifted geometric mean, median, minimum, and maximum times over the subset of 140 instances that do not hit the time limit. The last column denotes the number of instances for which each method hits the time limit.

We now summarize the out-of-sample prediction errors of our trained ML models when they are used to predict two strong partitioning points per partitioned variable (excluding variable bounds). Table 1 provides statistics of the scaled mean absolute errors (MAEs) of the out-of-sample predictions of the  $2n$  partitioning points (248 points for the pooling problem) produced by the ML model for each problem family. The MAEs of the predicted partitioning points are averaged over the 1000 instances in each family and scaled by the upper bounds of the corresponding  $x$  variables—these upper bounds are simply equal to one for the random bilinear and QCQP instances, but are greater than one for some of the partitioned variables in the pooling instances. Roughly 90% or higher of the partitioning points predicted using ML have a scaled MAE of less than 10% for each problem family, which indicates that the same underlying ML model is able to generate reasonable predictions of the strong partitioning points across these different problem families.

### 5.3 Results and discussion

We begin by benchmarking the hardness of our instances using BARON. We then compare the performance of default Alpine with the use of strong partitioning and its ML approximation (described in Section 5.2) within Alpine through a few metrics. All reported times are in seconds and do not include the time for solving the max-min problem (SP) or training the ML model.

#### 5.3.1 Benchmarking using BARON

To illustrate the non-trivial nature of our instances, we present statistics of their run times using BARON in Table 2. BARON solves the 10 variable and 20 variable random bilinear and QCQP instances within seconds, but takes over 4 minutes on average to solve the 50 variable instances and times out on 19/1000 of the 50 variable instances with univariate quadratic terms. BARON finds the random pooling instances to be

Problem Family	Solution Method	Solution Time (seconds)				# TLE
		Shifted GM	Median	Min	Max	
Bilinear $n = 10$	Alpine (default)	0.51	0.47	0.14	2.41	0
	Alpine+SP2	0.11	0.10	0.06	0.28	0
	Alpine+ML2	0.15	0.10	0.06	1.64	0
Bilinear $n = 20$	Alpine (default)	21.4	21.9	5.1	161.5	0
	Alpine+SP2	4.2	2.0	0.8	132.6	0
	Alpine+ML2	10.0	7.8	1.1	116.0	0
	Alpine+SP4	2.4	1.9	0.8	94.2	0
	Alpine+ML4	9.3	7.2	1.0	117.4	0
Bilinear $n = 50$	Alpine (default)	405.9	336.2	48.0	7135.9	24
	Alpine+SP2	52.8	34.9	4.2	5705.1	4
	Alpine+ML2	101.6	83.6	6.6	7071.7	5
QCQP $n = 10$	Alpine (default)	0.85	0.81	0.62	2.29	0
	Alpine+SP2	0.10	0.09	0.07	0.27	0
	Alpine+ML2	0.27	0.12	0.07	2.89	0
QCQP $n = 20$	Alpine (default)	40.1	35.6	4.6	241.1	0
	Alpine+SP2	7.7	1.7	0.8	135.4	0
	Alpine+ML2	13.0	9.5	1.0	180.1	0
	Alpine+SP4	2.4	1.5	0.7	125.7	0
	Alpine+ML4	9.4	6.4	0.9	101.2	0
QCQP $n = 50$	Alpine (default)	391.5	289.1	36.6	7198.2	0
	Alpine+SP2	63.3	51.9	4.2	6055.2	0
	Alpine+ML2	100.5	118.2	5.3	6514.2	0
Pooling	Alpine (default)	242.8	212.5	25.9	7091.9	7
	Alpine+SP2	66.7	49.7	1.6	6127.1	5
	Alpine+ML2	117.1	101.9	11.4	6097.0	1

Table 4: (**Solution Times**) Statistics of solution times. Columns correspond to the shifted geometric mean, median, minimum, and maximum times over the subset of 1000 instances that did not hit the time limit. The last time column denotes the number of instances for which each method hits the time limit.

significantly harder, timing out on 432/1000 instances<sup>8</sup> and taking roughly 7 minutes to solve the remaining 568/1000 instances on average. As suggested in the literature, we use the shifted geometric mean to compare the solution times of different algorithms on a family of test instances. The shifted geometric mean (shifted GM) of a positive vector  $t \in \mathbb{R}_+^N$  is defined as<sup>9</sup>:

$$\text{Shifted GM}(t) = \exp\left(\frac{1}{N} \sum_{i=1}^N \ln(\max(1, t_i + \text{shift}))\right) - \text{shift},$$

where we set  $\text{shift} = 10$  when comparing solution times in seconds. The last column in Table 2 notes the GM of the relative optimality gap at termination for instances where BARON hits the time limit. Following the definition of relative optimality gap in Alpine, this residual optimality gap is defined as  $\frac{\text{UB}-\text{LB}}{10^{-6}+|\text{UB}|}$ , where UB and LB are the upper and lower bounds returned by BARON at termination. We *emphasize* that our goal is not to compare the different versions of Alpine with BARON but rather to illustrate that our instances

<sup>8</sup>BARON finds global solutions but is unable to prove global optimality within the time limit.

<sup>9</sup><http://plato.asu.edu/ftp/shgeom.html>

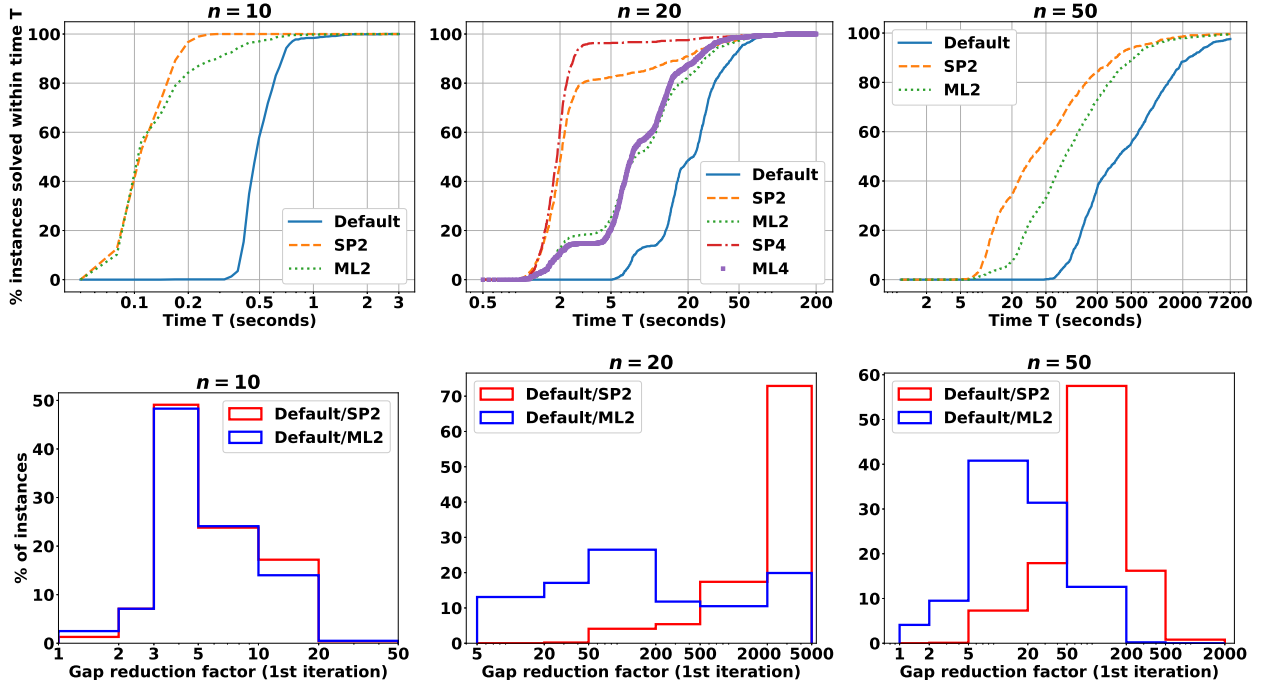


Figure 2: (**Bilinear Instances**) Top row: solution profiles indicating the percentage of instances solved by the different methods within time  $T$  seconds (higher is better). Bottom row: histogram plots of the ratios of the effective optimality gaps (10) of default Alpine with Alpine+SP2 and with Alpine+ML2 after one iteration (larger gap reduction factors are better).

and accelerations of Alpine are non-trivial.

### 5.3.2 Evaluating strong partitioning on benchmark QCQPs

We compare Alpine’s default partitioning strategy with the use of two/four strong partitioning points excluding the bounds (Alpine+SP2 and Alpine+SP4, respectively, see Algorithm 3) on a subset of BARON’s QCQP test library<sup>10</sup> [7]. Specifically, we only consider the 140 QCQP instances from Bao et al. [7] with 20 variables in order to keep the time for solving the max-min problem manageable. Table 3 presents statistics of the run times of default Alpine, Alpine+SP2, and Alpine+SP4 on these 140 instances. Alpine+SP2 and Alpine+SP4 are able to reduce the shifted GM of Alpine’s solution time by factors<sup>11</sup> of 3.3 and 5.3, respectively, which indicates that strong partitioning has the potential to result in significant speedups on broad families of QCQPs. Table 8 reports statistics of the solution times for the max-min problem over these 140 instances.

### 5.3.3 Evaluating the performance of strong partitioning and its ML approximation

We compare Alpine’s default partitioning strategy with the use of two strong partitioning points (excluding variable bounds) per partitioned variable (Alpine+SP2) and its ML approximation (Alpine+ML2) in Alpine’s first iteration. For the cases with  $n = 20$ , we also compare the above approaches with the use of four strong partitioning points (excluding the bounds) per partitioned variable (Alpine+SP4) and its ML approximation (Alpine+ML4) at Alpine’s first iteration. We compare these methods for each family of instances using two metrics: i. statistics of solution times, and ii. statistics of the effective optimality gap after Alpine’s first

<sup>10</sup>These 140 “qcqp2” instances are from <https://minlp.com/nlp-and-minlp-test-problems>

<sup>11</sup>These factors correspond to 69.7% and 81.1% average reductions in Alpine’s solution times.

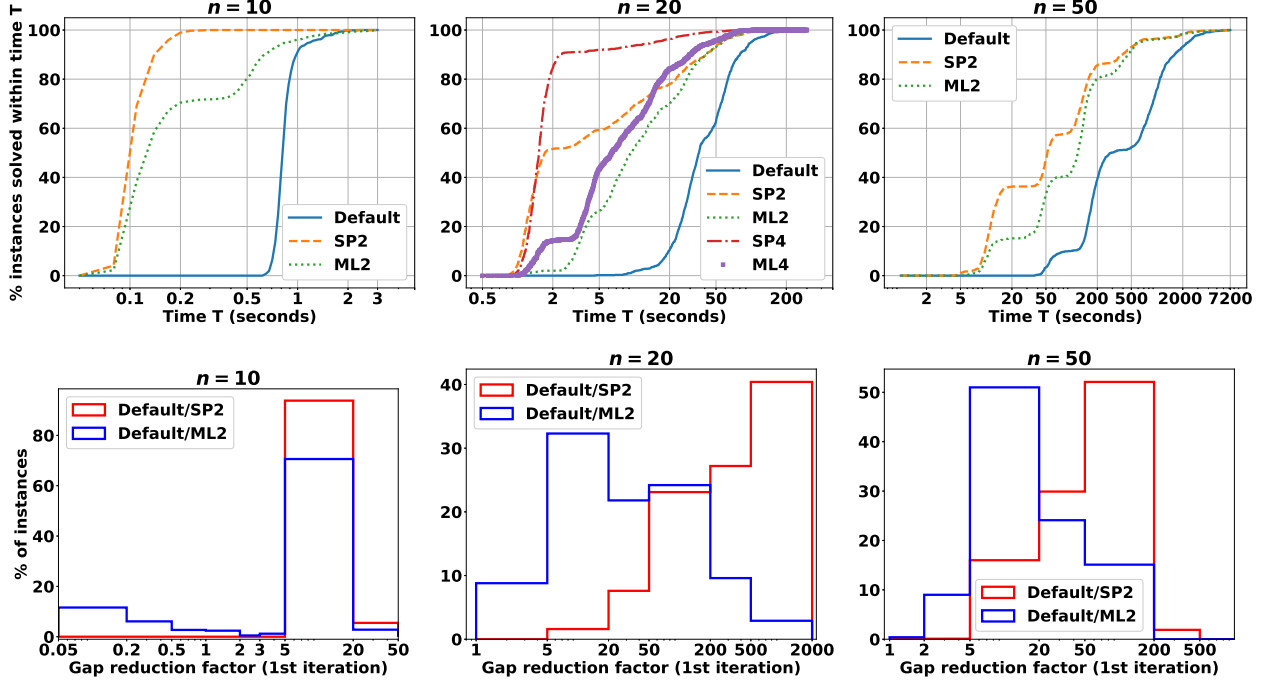


Figure 3: (**QCQP Instances**) Top row: solution profiles indicating the percentage of instances solved by the different methods within time  $T$  seconds (higher is better). Bottom row: histogram plots of the ratios of the effective optimality gaps (10) of default Alpine with Alpine+SP2 and with Alpine+ML2 after one iteration (larger gap reduction factors are better).

iteration. We define the *effective* relative optimality gap as

$$\text{Effective Optimality Gap} = \max \left\{ 10^{-4}, \frac{v^* - v^{\text{LBD}}}{10^{-6} + |v^*|} \right\}, \quad (10)$$

where  $v^*$  is the optimal objective value,  $v^{\text{LBD}}$  is Alpine’s lower bound after one iteration (using any one of the different approaches for specifying partitions), and  $10^{-4}$  is the target optimality gap. By measuring the gap of  $v^{\text{LBD}}$  relative to the optimal objective value  $v^*$  instead of the best found feasible solution, we do not let the performance of the local solver impact our evaluation of the different partitioning methods. Thresholding the optimality gap at  $10^{-4}$  also lends equal importance to all optimality gaps less than the target since all such gaps are sufficient for Alpine to converge.

Table 4 presents statistics of run times of default Alpine, Alpine with the different versions of strong partitioning at the first iteration, and Alpine with the different ML approximations of strong partitioning at the first iteration for the different problem families. Table 5 records the speedup/slowdown of the different versions of Alpine+SP and Alpine+ML over default Alpine. Table 6 presents statistics of the effective optimality gaps (10) of the different approaches after one iteration, whereas Table 7 notes the GM of the residual effective optimality gaps on instances for which the different approaches hit the time limit. Table 8 reports statistics of the solution times for the max-min problem for the different problem families. Figures 2, 3, and 4 plot solution profiles and histograms of the factor improvements of the effective optimality gaps for the bilinear, QCQP, and pooling families. We do not plot performance profiles due to their known issues (see <http://plato.asu.edu/bench.html>).

**Bilinear instances.** Table 4 implies Alpine+SP2 is able to reduce the shifted GM of default Alpine’s solution time by factors of 4.5, 5.1, and 7.7, respectively, for  $n = 10$ ,  $n = 20$ , and  $n = 50$  over 1000 instances.




Figure 4: (**Pooling Instances**) Left plot: solution profile indicating the percentage of instances solved by the different methods within time  $T$  seconds (higher is better). Right plot: histogram plots of the ratios of the effective optimality gaps (10) of default Alpine with Alpine+SP2 and with Alpine+ML2 after one iteration (larger gap reduction factors are better).

Alpine+ML2 is able to generate a moderate approximation of Alpine+SP2 overall, reducing the shifted GM of default Alpine’s solution time by factors of 3.5, 2.1, and 4, respectively, for  $n = 10$ ,  $n = 20$ , and  $n = 50$  over the same 1000 instances. For the  $n = 20$  instances, Alpine+SP4 and Alpine+ML4 reduce the shifted GM of default Alpine’s solution time by factors of 9 and 2.3, respectively. Table 5 implies Alpine+SP2 results in at least  $5\times$  speedup over default Alpine on 41.3% of the  $n = 10$  instances, and results in at least  $10\times$  speedup on 39.9% and 46.1% of the  $n = 20$  and  $n = 50$  instances, respectively. On the other hand, Alpine+ML2 yields at least  $5\times$  speedup over default Alpine on 40.1%, 22.2%, and 45.2% of the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances. Alpine+SP4 results in at least  $10\times$  speedup over default Alpine on 51.7% of the  $n = 20$  instances. Finally, Alpine+SP2 results in a maximum speedup of  $15\times$ ,  $49\times$ , and  $685\times$  for the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances, whereas Alpine+ML2 results in a maximum speedup of  $13\times$ ,  $38\times$ , and  $197\times$  for the same sets of instances.

Table 6 implies Alpine+SP2 reduces the GM of default Alpine’s effective optimality gap (10) after the first iteration by factors of 5.5, 2200, and 80, respectively, for  $n = 10$ ,  $n = 20$ , and  $n = 50$ . Alpine+ML2 reduces the GM of default Alpine’s effective gap after the first iteration by factors of 4.6, 180, and 15, respectively, for the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances. Interestingly, Alpine+SP2 is able to close the effective gap in the first iteration for 100%, 82.3%, and 46% of the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances, whereas default Alpine is able to close the gap in the first iteration for at most 0.1% of the instances for these different problem families, which demonstrates the effectiveness of the strong partitioning strategy. Finally, Table 7 shows that Alpine+SP2 and Alpine+ML2 terminate with smaller average optimality gaps on the  $n = 50$  instances where they time out compared to Alpine.

**QCQP instances.** Table 4 implies Alpine+SP2 is able to reduce the shifted GM of default Alpine’s solution time by factors of 8.4, 5.2, and 6.2, respectively, for  $n = 10$ ,  $n = 20$ , and  $n = 50$ . Alpine+ML2 is able to generate a moderate approximation of Alpine+SP2, reducing the shifted GM of default Alpine’s solution time by factors of 3.1, 3.1, and 3.9, respectively, for  $n = 10$ ,  $n = 20$ , and  $n = 50$  over the same 1000 instances. For the  $n = 20$  instances, Alpine+SP4 and Alpine+ML4 reduce the shifted GM of default Alpine’s solution time by factors of 16.4 and 4.3, respectively. Table 5 implies Alpine+SP2 results in at least  $10\times$  speedup over default Alpine on 20.5%, 54.6%, and 42.1% of the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances, respectively. On the other hand, Alpine+ML2 yields at least  $5\times$  speedup over default Alpine on 65.7%, 34.7%, and 42.7% of the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances. Alpine+SP4 results in at least  $20\times$  speedup over default Alpine on 61% of the  $n = 20$  instances. Finally, Alpine+SP2 results in a maximum speedup of  $22\times$ ,  $87\times$ , and  $98\times$  for the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances, whereas Alpine+ML2 results in a maximum speedup of  $19\times$ ,  $56\times$ , and  $32\times$  for the same sets of instances.

Table 6 implies Alpine+SP2 reduces the GM of default Alpine’s effective optimality gap (10) after the

Problem Family	Solution Method	Speedup/Slowdown Factor							
		< 0.5	0.5 – 1	1 – 2	2 – 5	5 – 10	10 – 20	20 – 50	> 50
Bilinear $n = 10$	% Alpine+SP2 inst.	—	—	1.1	57.6	40.1	1.2	0	0
	% Alpine+ML2 inst.	0.2	2.1	7.7	49.9	40.0	0.1	0	0
Bilinear $n = 20$	% Alpine+SP2 inst.	0.2	3.3	7.2	18.2	31.2	29.9	10.0	0.0
	% Alpine+ML2 inst.	3.3	9.8	25.5	39.2	15.3	6.0	0.9	0.0
	% Alpine+SP4 inst.	0.2	0.7	1.3	13.4	32.7	37.1	14.5	0.1
	% Alpine+ML4 inst.	2.8	10.5	23.3	41.4	15.2	5.9	0.9	0.0
Bilinear $n = 50$	% Alpine+SP2 inst.	0.4	1.3	7.2	18.7	26.3	24.3	14.9	6.9
	% Alpine+ML2 inst.	0.7	4.7	16.9	32.5	25.3	13.7	5.4	0.8
QCQP $n = 10$	% Alpine+SP2 inst.	—	—	0.1	3.3	76.1	20.4	0.1	0
	% Alpine+ML2 inst.	1.0	3.9	20.9	8.5	53.4	12.3	0	0
QCQP $n = 20$	% Alpine+SP2 inst.	0.1	3.2	12.2	18.4	11.5	19.4	32.6	2.6
	% Alpine+ML2 inst.	0.5	5.1	19.0	40.7	23.1	9.6	1.9	0.1
	% Alpine+SP4 inst.	0	0.2	1.3	3.8	5.5	28.2	53.7	7.3
	% Alpine+ML4 inst.	0	2.9	11.6	33.3	27.0	17.2	7.6	0.4
QCQP $n = 50$	% Alpine+SP2 inst.	0.9	1.3	10.7	22.0	23.0	32.5	7.2	2.4
	% Alpine+ML2 inst.	1.4	4.0	19.5	32.4	22.7	16.6	3.4	0
Pooling	% Alpine+SP2 inst.	2.2	6.4	19.7	26.0	21.8	16.8	6.7	0.4
	% Alpine+ML2 inst.	2.1	11.5	34.5	40.4	9.8	1.4	0.3	0

Table 5: (**Speedup/Slowdown**) Statistics of the speedup/slowdown of the different versions of Alpine with SP and its ML approximation (relative to default Alpine).

first iteration by factors of 13, 300, and 50, respectively, for the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances. On the other hand, Alpine+ML2 reduces the GM of default Alpine’s effective gap after the first iteration by factors of 4.3, 31, and 17, respectively, for  $n = 10$ ,  $n = 20$ , and  $n = 50$ . Note that Alpine+SP2 is able to close the effective gap in the first iteration for 100%, 52.2%, and 39% of the  $n = 10$ ,  $n = 20$ , and  $n = 50$  instances, whereas default Alpine is unable to close the gap in the first iteration for any instance in these problem families.

**Pooling instances.** Table 4 implies Alpine+SP2 and Alpine+ML2 reduce the shifted GM of default Alpine’s solution time by factors of 3.6 and 2.1 over the 1000 instances. Table 5 implies Alpine+SP2 and Alpine+ML2 result in at least  $5\times$  speedup over default Alpine on 45.7% and 11.5% of the instances, respectively. Table 6 implies Alpine+SP2 and Alpine+ML2 reduce the GM of default Alpine’s effective optimality gap (10) after the first iteration by factors of 28 and 4.5, respectively. After the first iteration, Alpine+SP2 closes the effective optimality gap for 45.2% of the instances, whereas default Alpine is unable to close the gap for any of the 1000 instances. Finally, Alpine+SP2 and Alpine+ML2 result in maximum speedups of  $120\times$  and  $41\times$ .

**Summary.** Tables 4 to 6 and Figures 2 to 4 clearly show the benefits of strong partitioning and its ML approximation over Alpine’s default partitioning strategy. They also demonstrate that Alpine+SP and Alpine+ML are able to match or even outperform (particularly on the pooling instances) the performance of the state-of-the-art solver BARON (with default options) on average over the different problem families. While our off-the-shelf ML model is able to yield a moderate approximation of SP across these different problem families, there is a clear scope for significant improvement with tailored ML approaches.



Problem Family	Solution Method	Effective Optimality Gap after 1 iteration				% Instances Gap Closed
		GM	Median	Min	Max	
Bilinear $n = 10$	Alpine (default)	$5.5 \times 10^{-4}$	$4.5 \times 10^{-4}$	$10^{-4}$	$3.4 \times 10^{-2}$	0.1
	Alpine+SP2	$10^{-4}$	$10^{-4}$	$10^{-4}$	$10^{-4}$	100
	Alpine+ML2	$1.2 \times 10^{-4}$	$10^{-4}$	$10^{-4}$	$4.4 \times 10^{-2}$	88.3
Bilinear $n = 20$	Alpine (default)	$2.9 \times 10^{-1}$	$3.3 \times 10^{-1}$	$7.3 \times 10^{-2}$	$4.8 \times 10^{-1}$	0
	Alpine+SP2	$1.3 \times 10^{-4}$	$10^{-4}$	$10^{-4}$	$6 \times 10^{-3}$	82.3
	Alpine+ML2	$1.6 \times 10^{-3}$	$1.9 \times 10^{-3}$	$10^{-4}$	$1.4 \times 10^{-1}$	18.9
	Alpine+SP4	$1.0 \times 10^{-4}$	$10^{-4}$	$10^{-4}$	$4.7 \times 10^{-4}$	96.0
	Alpine+ML4	$2.2 \times 10^{-3}$	$3.6 \times 10^{-3}$	$10^{-4}$	$9.9 \times 10^{-2}$	14.5
Bilinear $n = 50$	Alpine (default)	$1.4 \times 10^{-2}$	$1.7 \times 10^{-2}$	$10^{-4}$	$6.9 \times 10^{-2}$	0.1
	Alpine+SP2	$1.7 \times 10^{-4}$	$1.2 \times 10^{-4}$	$10^{-4}$	$5.4 \times 10^{-1}$	46.0
	Alpine+ML2	$9.5 \times 10^{-4}$	$9.4 \times 10^{-4}$	$10^{-4}$	$4.9 \times 10^{-1}$	5.6
QCQP $n = 10$	Alpine (default)	$1.3 \times 10^{-3}$	$1.2 \times 10^{-3}$	$7.5 \times 10^{-4}$	$1.9 \times 10^{-2}$	0
	Alpine+SP2	$10^{-4}$	$10^{-4}$	$10^{-4}$	$10^{-4}$	100
	Alpine+ML2	$3.0 \times 10^{-4}$	$10^{-4}$	$10^{-4}$	$1.3 \times 10^{-1}$	71.8
QCQP $n = 20$	Alpine (default)	$6.3 \times 10^{-2}$	$7.8 \times 10^{-2}$	$3.0 \times 10^{-3}$	$2.1 \times 10^{-1}$	0
	Alpine+SP2	$2.1 \times 10^{-4}$	$10^{-4}$	$10^{-4}$	$6.6 \times 10^{-3}$	52.2
	Alpine+ML2	$2.0 \times 10^{-3}$	$2.5 \times 10^{-3}$	$10^{-4}$	$5.8 \times 10^{-2}$	2.0
	Alpine+SP4	$1.1 \times 10^{-4}$	$10^{-4}$	$10^{-4}$	$3.6 \times 10^{-3}$	92.6
	Alpine+ML4	$1.5 \times 10^{-3}$	$1.7 \times 10^{-3}$	$10^{-4}$	$6.7 \times 10^{-2}$	14.7
QCQP $n = 50$	Alpine (default)	$8.1 \times 10^{-3}$	$1.0 \times 10^{-2}$	$6.3 \times 10^{-4}$	$2.8 \times 10^{-2}$	0
	Alpine+SP2	$1.6 \times 10^{-4}$	$1.3 \times 10^{-4}$	$10^{-4}$	$1.0 \times 10^{-3}$	39.0
	Alpine+ML2	$4.8 \times 10^{-4}$	$5.3 \times 10^{-4}$	$10^{-4}$	$1.5 \times 10^{-2}$	14.9
Pooling	Alpine (default)	$6.8 \times 10^{-3}$	$6.4 \times 10^{-3}$	$1.2 \times 10^{-3}$	$4.4 \times 10^{-2}$	0
	Alpine+SP2	$2.4 \times 10^{-4}$	$1.4 \times 10^{-4}$	$10^{-4}$	$3.1 \times 10^{-3}$	45.2
	Alpine+ML2	$1.5 \times 10^{-3}$	$1.6 \times 10^{-3}$	$10^{-4}$	$6.3 \times 10^{-3}$	0.1

Table 6: (**Effective Optimality Gaps**) Statistics of effective optimality gaps (10) after one iteration (note: minimum possible value =  $10^{-4}$ , the target gap). Columns record the geometric mean, median, minimum, and maximum effective gaps over 1000 instances. The last column is the percentage of instances for which each method results in the minimum possible effective optimality gap of  $10^{-4}$  after one iteration.

## 6 Future work

There are several interesting avenues for future work.

First, instead of prespecifying the number of partitioning points per variable for SP, we could allocate a different number of partitions per variable based on their relative impact on the lower bound. Suppose we wish to specify *at most*  $d + 2$  partitioning points for each variable and are given a budget  $B \in [d \times n]$  for the total number of partitioning points across all variables (excluding variable bounds). We can solve the following max-min problem to determine both the optimal allocation of partitions *and* the optimal specification of partitioning points across the partitioned variables:

$$\max_{(P,Z) \in \mathfrak{P}_z} v(P),$$

where  $P := (p^1, p^2, \dots, p^n)$  denotes the (potential) partitioning points,  $v(P)$  is defined in (PMR-OA), and  $Z := (z^1, z^2, \dots, z^n)$  is a  $d \times n$  matrix of binary decisions. The partitioning point  $p_j^i$  is added to the partition

Problem Family Method	Bilinear $n = 50$			Pooling		
	Alpine (default)	Alpine+SP2	Alpine+ML2	Alpine (default)	Alpine+SP2	Alpine+ML2
TLE Gap (GM)	$4.4 \times 10^{-4}$	$1.6 \times 10^{-4}$	$1.8 \times 10^{-4}$	$2.9 \times 10^{-4}$	$2.1 \times 10^{-4}$	$2.8 \times 10^{-4}$

Table 7: (**Effective TLE Optimality Gaps**) Geometric mean of residual effective optimality gaps (target =  $10^{-4}$ ) on instances for which methods hit the time limit.

Problem Family	Solution Method	Max-Min Solution Time (seconds)				
		Shifted GM	Median	Min	Max	Std. Dev.
Bilinear $n = 10$	SP2	16	14	6	96	13
Bilinear $n = 20$	SP2	528	445	136	2389	544
	SP4	1244	1117	374	4360	893
Bilinear $n = 50$	SP2	7070	7404	1271	23166	3268
QCQP $n = 10$	SP2	8	8	6	53	3
QCQP $n = 20$	SP2	1731	1826	171	4244	654
	SP4	2152	2740	471	5965	961
QCQP $n = 50$	SP2	16964	17074	8626	23551	2319
Pooling	SP2	15658	15148	1088	77029	8657
Benchmark QCQPs	SP2	413	364	7	27907	4432
	SP4	895	651	12	136320	15444

Table 8: (**Max-Min Solution Times**) Statistics of max-min solution times. Columns correspond to the shifted geometric mean, median, minimum, maximum, and standard deviation of times for solving the max-min problem (SP).

$\mathcal{P}_i$  of  $x_i$  only if the variable  $z_j^i = 1$ . Finally, the MILP representable set  $\mathfrak{P}_Z$  is defined as

$$\mathfrak{P}_Z := \left\{ (P, Z) \in \mathfrak{P} \times \{0, 1\}^{d \times n} : \sum_{i=1}^n \sum_{j=1}^d z_j^i = B, z_j^i = 0 \implies p_j^i = 0, \forall (i, j) \in [n] \times [d] \right\}.$$

If  $z_j^i = 0$ , then the partitioning point  $p_j^i$  is made redundant by forcing it to 0. Note that the above outer-maximization problem involves binary decision variables  $Z$  unlike the strong partitioning problem (SP), which necessitates new techniques for its solution.

Second, designing more efficient approaches for solving the strong partitioning problem (SP) would help scale our approach to larger problem dimensions (and also make it easier to generating more training data). Third, designing tailored ML architectures that can achieve similar speedups as strong partitioning and can accommodate variable feature and output dimensions merits investigation. Fourth, motivated by the cluster problem [35, 36], it would be interesting to explore variants that choose a different subset of variables to be partitioned at each iteration within Alpine. Finally, using strong partitioning to choose Alpine’s partitions at the second iteration and beyond can help promote convergence of its bounds in fewer iterations.

## Acknowledgments

The authors gratefully acknowledge funding from Los Alamos National Laboratory’s (LANL’s) “Center for Nonlinear Studies” and the U.S. Department of Energy’s “Laboratory Directed Research and Development (LDRD)” program under the projects “20230091ER: Learning to Accelerate Global Solutions for Non-convex Optimization” and “20210078DR: The Optimization of Machine Learning: Imposing Requirements on Artificial Intelligence.” This research used resources provided by LANL’s Darwin testbed, which is funded by

the Computational Systems and Software Environments subprogram of LANL’s Advanced Simulation and Computing program (NNSA/DOE).

## References

- [1] T. Achterberg. *Constraint integer programming*. PhD thesis, Technischen Universität Berlin, 2007.
- [2] F. A. Al-Khayyal and J. E. Falk. Jointly constrained biconvex programming. *Mathematics of Operations Research*, 8(2): 273–286, 1983.
- [3] M. Alfaki and D. Haugland. A multi-commodity flow formulation for the generalized pooling problem. *Journal of Global Optimization*, 56(3):917–937, 2013.
- [4] A. M. Alvarez, Q. Louveaux, and L. Wehenkel. A machine learning-based approximation of strong branching. *INFORMS Journal on Computing*, 29(1):185–195, 2017.
- [5] M.-F. Balcan, T. Dick, T. Sandholm, and E. Vitercik. Learning to branch. In *International Conference on Machine Learning*, pages 344–353. PMLR, 2018.
- [6] R. Baltean-Lugojan, P. Bonami, R. Misener, and A. Tramontani. Scoring positive semidefinite cutting planes for quadratic optimization via trained neural networks. *Optimization Online*, 2019. URL <https://optimization-online.org/2018/11/6943/>.
- [7] X. Bao, N. V. Sahinidis, and M. Tawarmalani. Semidefinite relaxations for quadratically constrained quadratic programming: A review and comparisons. *Mathematical Programming*, 129(1):129–157, 2011.
- [8] P. Belotti, J. Lee, L. Liberti, F. Margot, and A. Wächter. Branching and bounds tightening techniques for non-convex MINLP. *Optimization Methods & Software*, 24(4-5):597–634, 2009.
- [9] Y. Bengio, A. Lodi, and A. Prouvost. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.
- [10] M. L. Bergamini, P. Aguirre, and I. Grossmann. Logic-based outer approximation for globally optimal synthesis of process networks. *Computers & Chemical Engineering*, 29(9):1914–1933, 2005.
- [11] M. L. Bergamini, I. Grossmann, N. Scenna, and P. Aguirre. An improved piecewise outer-approximation algorithm for the global optimization of MINLP models involving concave and bilinear terms. *Computers & Chemical Engineering*, 32(3): 477–493, 2008.
- [12] K. Bestuzheva, M. Besançon, W.-K. Chen, A. Chmiela, T. Donkiewicz, J. van Doornmalen, L. Eifler, O. Gaul, G. Gamrath, A. Gleixner, et al. The SCIP optimization suite 8.0. *arXiv preprint: 2112.08872*, 2021.
- [13] D. Bienstock, M. Escobar, C. Gentile, and L. Liberti. Mathematical programming formulations for the alternating current optimal power flow problem. *Annals of Operations Research*, pages 1–39, 2022.
- [14] A. Billionnet, S. Elloumi, and A. Lambert. Extending the QCR method to general mixed-integer programs. *Mathematical Programming*, 131(1):381–401, 2012.
- [15] P. Bonami, A. Lodi, and G. Zarpellon. Learning a classification of mixed-integer quadratic programming problems. In *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 595–604. Springer, 2018.
- [16] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and regression trees*. Routledge, 2017.
- [17] S. Burer and D. Vandenbussche. A finite branch-and-bound algorithm for nonconvex quadratic programming via semidefinite relaxations. *Mathematical Programming*, 113(2):259–282, 2008.
- [18] Q. Cappart, D. Chételat, E. Khalil, A. Lodi, C. Morris, and P. Veličković. Combinatorial optimization and reasoning with graph neural networks. *arXiv preprint: 2102.09544*, 2021.
- [19] P. M. Castro. Normalized multiparametric disaggregation: an efficient relaxation for mixed-integer bilinear problems. *Journal of Global Optimization*, 64(4):765–784, 2016.
- [20] F. Cengil, H. Nagarajan, R. Bent, S. Eksioğlu, and B. Eksioğlu. Learning to accelerate globally optimal solutions to the AC optimal power flow problem. *Electric Power Systems Research*, 212:108275, 2022.
- [21] A. Costa, P. Hansen, and L. Liberti. On the impact of symmetry-breaking constraints on spatial branch-and-bound for circle packing in a square. *Discrete Applied Mathematics*, 161(1-2):96–106, 2013.
- [22] D. De Wolf and Y. Smeers. Generalized derivatives of the optimal value of a linear program with respect to matrix coefficients. *European Journal of Operational Research*, 291(2):491–496, 2021.
- [23] G. Di Liberto, S. Kadioglu, K. Leo, and Y. Malitsky. DASH: Dynamic approach for switching heuristics. *European Journal of Operational Research*, 248(3):943–953, 2016.
- [24] H. Drucker. Improving regressors using boosting techniques. In *ICML*, volume 97, pages 107–115. Citeseer, 1997.
- [25] M. Etheve, Z. Alès, C. Bissuel, O. Juan, and S. Kedad-Sidhoum. Reinforcement learning for variable selection in a branch and bound algorithm. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 176–185. Springer, 2020.

- [26] R. M. Freund. Postoptimal analysis of a linear program under simultaneous changes in matrix coefficients. In *Mathematical Programming Essays in Honor of George B. Dantzig Part I*, pages 1–13. Springer, 1985.
- [27] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [28] M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi. Exact combinatorial optimization with graph convolutional neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [29] B. Ghaddar, I. Gómez-Casares, J. González-Díaz, B. González-Rodríguez, B. Pateiro-López, and S. Rodríguez-Ballesteros. Learning for spatial branching: An algorithm selection approach. *arXiv preprint arXiv:2204.10834*, 2022.
- [30] B. González-Rodríguez, R. Alvite-Pazó, S. Alvite-Pazó, B. Ghaddar, and J. González-Díaz. Polynomial optimization: Enhancing RLT relaxations with conic constraints. *arXiv preprint arXiv:2208.05608*, 2022.
- [31] C. A. Haverly. Studies of the behavior of recursion for the pooling problem. *ACM SIGMAP Bulletin*, 25:19–28, 1978.
- [32] H. He, H. Daume III, and J. M. Eisner. Learning to search in branch and bound algorithms. *Advances in Neural Information Processing Systems*, 27:3293–3301, 2014.
- [33] J. Im. Sensitivity analysis and robust optimization: A geometric approach for the special case of linear optimization. Master’s thesis, University of Waterloo, 2018.
- [34] R. Kannan. *Algorithms, analysis and software for the global optimization of two-stage stochastic programs*. PhD thesis, Massachusetts Institute of Technology, 2018.
- [35] R. Kannan and P. I. Barton. The cluster problem in constrained global optimization. *Journal of Global Optimization*, 69(3):629–676, 2017.
- [36] R. Kannan and P. I. Barton. Convergence-order analysis of branch-and-bound algorithms for constrained problems. *Journal of Global Optimization*, 71(4):753–813, 2018.
- [37] R. Karuppiah and I. E. Grossmann. Global optimization for the synthesis of integrated water systems in chemical processes. *Computers & Chemical Engineering*, 30(4):650–673, 2006.
- [38] E. Khalil, P. Le Bodic, L. Song, G. Nemhauser, and B. Dilkina. Learning to branch in mixed integer programming. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [39] J. Kim, J.-P. P. Richard, and M. Tawarmalani. Piecewise polyhedral relaxations of multilinear optimization. *Optimization Online*, 2022. URL [http://www.optimization-online.org/DB\\_HTML/2022/07/8974.html](http://www.optimization-online.org/DB_HTML/2022/07/8974.html).
- [40] S. Kolodziej, P. M. Castro, and I. E. Grossmann. Global optimization of bilinear programs with a multiparametric disaggregation technique. *Journal of Global Optimization*, 57(4):1039–1063, 2013.
- [41] T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica: Journal of the Econometric Society*, pages 53–76, 1957.
- [42] J. Kotary, F. Fioretto, P. Van Hentenryck, and B. Wilder. End-to-end constrained optimization learning: A survey. *arXiv preprint: 2103.16378*, 2021.
- [43] M. Lee, N. Ma, G. Yu, and H. Dai. Accelerating generalized Benders decomposition for wireless resource allocation. *IEEE Transactions on Wireless Communications*, 2020.
- [44] X. Li, E. Armagan, A. Tomagard, and P. I. Barton. Stochastic pooling problem for natural gas production network design and operation under uncertainty. *AIChE Journal*, 57(8):2120–2135, 2011.
- [45] Y. Lin and L. Schrage. The global solver in the LINDO API. *Optimization Methods & Software*, 24(4-5):657–668, 2009.
- [46] J. Liu, N. Ploshkas, and N. V. Sahinidis. Tuning BARON using derivative-free optimization algorithms. *Journal of Global Optimization*, 74(4):611–637, 2019.
- [47] A. Lodi and G. Zarpellon. On learning and branching: a survey. *Top*, 25(2):207–236, 2017.
- [48] M. Lu, H. Nagarajan, R. Bent, S. D. Eksioglu, and S. J. Mason. Tight piecewise convex relaxations for global optimization of optimal power flow. In *2018 Power Systems Computation Conference*, pages 1–7. IEEE, 2018.
- [49] J. Luedtke, C. d’Ambrosio, J. Linderoth, and J. Schweiger. Strong convex nonlinear relaxations of the pooling problem. *SIAM Journal on Optimization*, 30(2):1582–1609, 2020.
- [50] M. M. Mäkelä. Multiobjective proximal bundle method for nonconvex nonsmooth optimization: Fortran subroutine mpbngc 2.0. url: [http://napsu.karmita.fi/publications/pbncgc\\_report.pdf](http://napsu.karmita.fi/publications/pbncgc_report.pdf). *Reports of the Department of Mathematical Information Technology, Series B. Scientific Computing, B*, 13, 2003.
- [51] G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I—convex underestimating problems. *Mathematical Programming*, 10(1):147–175, 1976.
- [52] C. A. Meyer and C. A. Floudas. Global optimization of a combinatorially complex generalized pooling problem. *AIChE Journal*, 52(3):1027–1037, 2006.
- [53] R. Misener and C. A. Floudas. GloMIQO: Global mixed-integer quadratic optimizer. *Journal of Global Optimization*, 57(1):3–50, 2013.
- [54] R. Misener and C. A. Floudas. ANTIGONE: algorithms for continuous/integer global optimization of nonlinear equations. *Journal of Global Optimization*, 59(2):503–526, 2014.

- [55] H. Nagarajan, M. Lu, E. Yamangil, and R. Bent. Tightening McCormick relaxations for nonlinear programs via dynamic multivariate partitioning. In *International Conference on Principles and Practice of Constraint Programming*, pages 369–387. Springer, 2016.
- [56] H. Nagarajan, M. Lu, S. Wang, R. Bent, and K. Sundar. An adaptive, multivariate partitioning algorithm for global optimization of nonconvex programs. *Journal of Global Optimization*, 74(4):639–675, 2019.
- [57] V. Nair, S. Bartunov, F. Gimeno, I. von Glehn, P. Lichocki, I. Lobov, B. O’Donoghue, N. Sonnerat, C. Tjandraatmadja, P. Wang, R. Addanki, T. Hapuarachchi, T. Keck, J. Keeling, P. Kohli, I. Ktena, Y. Li, O. Vinyals, and Y. Zwols. Solving mixed integer programs using neural networks. *arXiv preprint: 2012.13349*, 2020.
- [58] G. Nannicini, P. Belotti, J. Lee, J. Linderoth, F. Margot, and A. Wächter. A probing algorithm for MINLP with failure prediction by SVM. In *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pages 154–169. Springer, 2011.
- [59] C. J. Nohra, A. U. Raghunathan, and N. Sahinidis. Spectral relaxations and branching strategies for global optimization of mixed-integer quadratic programs. *SIAM Journal on Optimization*, 31(1):142–171, 2021.
- [60] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [61] N. V. Sahinidis. BARON: A general purpose global optimization software package. *Journal of Global Optimization*, 8(2): 201–205, 1996.
- [62] Y. Saif, A. Elkamel, and M. Pritzker. Global optimization of reverse osmosis network for wastewater treatment and minimization. *Industrial & Engineering Chemistry Research*, 47(9):3060–3070, 2008.
- [63] H. D. Sherali and W. P. Adams. *A reformulation-linearization technique for solving discrete and continuous nonconvex problems*, volume 31. Springer Science & Business Media, 2013.
- [64] G. Still. Lectures on parametric optimization: An introduction. *Optimization Online*, 2018. URL <https://optimization-online.org/2018/04/6587/>.
- [65] K. Sundar, H. Nagarajan, J. Linderoth, S. Wang, and R. Bent. Piecewise polyhedral formulations for a multilinear term. *Operations Research Letters*, 49(1):144–149, 2021.
- [66] S. A. Vavasis. Quadratic programming is in NP. *Information Processing Letters*, 36(2):73–77, 1990.
- [67] F. G. Vázquez, J.-J. Rückmann, O. Stein, and G. Still. Generalized semi-infinite programming: a tutorial. *Journal of Computational and Applied Mathematics*, 217(2):394–419, 2008.
- [68] D. S. Wicaksono and I. A. Karimi. Piecewise MILP under- and overestimators for global optimization of bilinear programs. *AIChE Journal*, 54(4):991–1008, 2008.
- [69] Y. Yang and P. I. Barton. Integrated crude selection and refinery optimization under uncertainty. *AIChE journal*, 62(4): 1038–1053, 2016.
- [70] G. Zarpellon, J. Jo, A. Lodi, and Y. Bengio. Parameterizing branch-and-bound search trees to learn branching policies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3931–3939, 2021.