

On handling indicator constraints in mixed integer programming

Pietro Belotti¹ · Pierre Bonami² · Matteo Fischetti³ · Andrea Lodi^{4,5} · Michele Monaci⁴ · Amaya Nogales-Gómez⁶ · Domenico Salvagnin^{3,7}

Received: 28 July 2015 / Published online: 4 May 2016 © Springer Science+Business Media New York 2016

Abstract Mixed integer programming (MIP) is commonly used to model indicator constraints, i.e., constraints that either hold or are relaxed depending on the value of a binary variable. Unfortunately, those models tend to lead to weak continuous relaxations and turn out to be unsolvable in practice; this is what happens, for e.g., in the case of Classification problems with Ramp Loss functions that represent an important application in this context. In this paper we show the computational evidence that a relevant class of these Classification instances can be solved far more efficiently if a nonlinear, nonconvex reformulation of the indicator constraints is used instead of the linear one. Inspired by this empirical and surprising observation, we show that aggressive bound tightening is the crucial ingredient for solving this class of instances, and we devise a pair of computationally effective algorithmic approaches that exploit it within MIP. One of these methods is currently part of the arsenal of IBM-Cplex since version 12.6.1. More generally, we argue that aggressive bound tightening is often overlooked in MIP, while it represents a significant building block for enhancing MIP technology when indicator constraints and disjunctive terms are present.

Andrea Lodi andrea.lodi@polymtl.ca

- ³ University of Padova, Padua, Italy
- ⁴ University of Bologna, Bologna, Italy
- ⁵ École Polytechnique de Montréal, Montreal, Canada
- ⁶ Mathematical and Algorithmic Sciences Lab, Huawei France R&D, Paris, France
- 7 IBM, Milano, Italy

¹ FICO, Birmingham, UK

² IBM, Madrid, Spain

Keywords Mixed-integer linear programming · Mixed-integer quadratic programming · Indicator constraints

1 Introduction

Let us consider the linear inequality

$$\alpha^T x \le x_0,\tag{1}$$

in which both $x \in \mathbb{R}^d$ and $x_0 \in \mathbb{R}$ are variables, while α is a given *d*-dimensional vector. It is a very well-known modeling trick in Mixed Integer Linear Programming (MILP) to use a binary variable to control whether linear constraint (1) is active or not depending on other parts of the model or at the price of paying a penalty in the objective function. Then, the constraint is reformulated as the following *big-M* or *indicator* constraint

$$\alpha^T x \le x_0 + Mt,\tag{2}$$

where $t \in \{0, 1\}$ and M is a large-enough value that guarantees that the constraint is inactive if t = 1. As similar tricks may be used in Mixed Integer Quadratic Programming (MIQP), from now on, we will write MIP to denote both classes of problems.

Although they provide a clean and flexible modeling tool to deal with nonlinearities and logical implications by staying within the MIP framework, it is well-known that indicator constraints present the drawback of having a weak continuous relaxation. Indeed, depending on the value M and on the value attained by expression " $\alpha^T x - x_0$ ", very small (fractional) values of t might be sufficient to satisfy the constraint. This leads to quality issues with a continuous relaxation value typically very far away from the mixed integer optimum, but, sometimes even more importantly, might lead to numerical issues, with the MIP solvers being unable to assert if a t value below the integer tolerance is in fact a true solution.

An alternative for logical implications that has been used in the Mixed Integer Nonlinear Programming (MINLP) literature for decades is provided by the *complementary* formulation

$$(\alpha^T x - x_0)\bar{t} \le 0,\tag{3}$$

where $\bar{t} = 1-t$. However, (3) is a nonconvex constraint. Such a source of nonconvexity might not significantly complicate the solution of already nonconvex MINLP models arising, for example, in Chemical Engineering applications, see [19]. In addition, numerical issues on the choice of the value of M do not appear anymore, at least in the formulation. On the contrary, in the cases where those logical constraints were the only sources of nonconvexity, the common approach has always been that of using constraints (2) and MIP techniques.

Before stating the contribution of the present paper it is worth mentioning that the big-M formulation (2) is just the weakest (but easiest and commonly used) disjunctive programming approach (see, e.g., [2,11,19]) to deal with indicator constraints and disjunctions in general. The reader is referred to [6] for a detailed and more theoretical discussion on the topic that is outside the scope of this paper.

Contribution of the paper In this paper we expose a class of convex Mixed Integer Quadratic Programming problems arising in *Supervised Classification* where the Global Optimization (GO) solver Couenne [13] using reformulation (3) is consistently faster than the state-of-the-art commercial MIP solvers IBM-Cplex [21], Gurobi [20] and FICO Xpress [17] on the *big-M* formulation (2). This is quite counter-intuitive because, in general, convex MIPs admit more efficient solution techniques both in theory and in practice, especially by benefiting from virtually all machinery of MIP solvers. Inspired by this empirical and surprising observation, we show that aggressive bound tightening is the crucial ingredient for solving this class of instances, and we devise a pair of computationally effective algorithmic approaches that exploit it within MIP.

Specifically,

- On the one side, we were able with a specialized algorithm to optimally solve in just seconds instances that could not be solved by state-of-the-art MIP solver in hours;
- On the other side, one of the devised methods is currently part of the arsenal of IBM-Cplex since version 12.6.1.

More generally, we argue that aggressive bound tightening is often overlooked in MIP, while it represents a significant building block for enhancing MIP technology when indicator constraints and disjunctive terms are present.

Organization of the paper The remainder of the paper is organized as follows. In Sect. 2 we discuss the application we use as an example. In Sect. 3 we show the initial set of computational results, showing that state-of-the-art commercial MIP solvers have similar performance on the instances in our testbed, and are clearly dominated by a global optimization solver. In Sect. 4 we discuss why those results are surprising while in Sect. 5 we carefully analyze the reasons of the success of Couenne versus IBM-Cplex. In Sect. 6 we present two approaches to enhance IBM-Cplex trying to mimic Couenne's behavior, and present some computational experiments both on our specific testbed and on the general MIPLIB2010 library of instances [23]. Finally, some conclusions are drawn in Sect. 7.

2 Support vector machines with the ramp loss

In *Supervised Classification*, see, e.g., [31], we are given a set of objects Ω partitioned into classes and the aim is to build a procedure for classifying new objects. In its simplest form, each object $i \in \Omega$ has associated a pair (x_i, y_i) , where the predictor vector x_i takes values on a set $X \subseteq \mathbb{R}^d$ and $y_i \in \{-1, 1\}$ is the class membership of object *i*, also known as the label of object *i*.

Support Vector Machines (SVM) methods, see, e.g., [14], have proven to be one of the state-of-the-art methods for Supervised Classification. The SVM aims at separating the two classes by means of a hyperplane, $\omega^{\top}x + b = 0$, found by solving the optimization problem

$$\min_{\omega \in \mathbb{R}^d, b \in \mathbb{R}} \frac{\omega^\top \omega}{2} + \frac{C}{n} \sum_{i=1}^n g((1 - y_i(\omega^\top x_i + b))^+),$$

where *n* is the size of the sample used to build the classifier, $(a)^+ = \max\{a, 0\}, C$ is a nonnegative parameter, and *g* a nondecreasing function in \mathbb{R}_+ , the so-called *loss* function. The reader is referred to [9] for a recent review on Mathematical Optimization and SVMs.

Several recent papers have been devoted to analyzing SVM with the so-called *ramp* loss function, $g(t) = (\min\{t, 2\})^+$, discussed in detail in the next section.

2.1 An MIQP formulation

In this paper, we are interested in the SVM with the ramp loss function, see [12,29]. In this model, objects are penalized in a different way depending on whether or not they fall inside or outside the margin, i.e., if they fall between $\omega^{\top}x + b = -1$ and $\omega^{\top}x + b = 1$. Misclassified objects that fall outside the margin have a fixed loss of 2, while objects that fall inside the margin have a continuous loss between 0 and 2. The state-of-the-art algorithm is given in [8], where the ramp loss model, denoted as RLM, is formulated as the MIQP problem

(RLM)
$$\min_{\omega,b,\xi,z} \frac{1}{2} \sum_{j=1}^{d} \omega_j^2 + \frac{C}{n} \left(\sum_{i=1}^{n} \xi_i + 2 \sum_{i=1}^{n} z_i \right)$$
(4)

s.t.

$$y_i(\omega^{\top} x_i + b) \ge 1 - \xi_i - M z_i \quad \forall i = 1, \dots, n$$
(5)

$$0 \le \xi_i \le 2 \quad \forall i = 1, \dots, n \tag{6}$$

$$\in \{0,1\}^n \tag{7}$$

$$\omega \in \mathbb{R}^d \tag{8}$$

$$b \in \mathbb{R},$$
 (9)

where M > 0 is a big enough constant, $\xi = (\xi_i)$ denotes the vector of deviation/penalty variables and *C* is the tradeoff parameter that calls for tuning. For a given object *i*, the binary variable z_i is equal to 1 if object *i* is misclassified outside the margin and 0 otherwise. The reader is referred to [8] for further details on this formulation, which is denoted as "SVMIP1(ramp)" in [8].

Ζ.

The appeal of model (4)–(9) relies in the fact that it can (potentially) be solved by a black-box MIQP solver, e.g., IBM-Cplex. More precisely, objective function (4) is convex while constraints are linear, thus virtually all the very sophisticated and effective machinery for MIP problems can be applied—the effectiveness of the resulting approach deserving evaluation from a computational point of view. However, the solution method proposed in [8] is able to solve to optimality only a quite limited number of instances, although some problem-specific cutting planes and reductions are used to help the MIP solver, namely, IBM-Cplex. Essentially, this difficulty is due to the *big-M* constraints (5) that make the continuous relaxation of model (4)–(9) very weak. Branching is effective for small problems but the almost-complete enumeration is ineffective for instances of serious size. Cutting planes are not likely to solve the problem, unless they would be specifically designed to face the *big-M* issue, or, more precisely, the disjunctive nature of constraints (5).

2.2 A nonconvex formulation

Motivated by the difficulty of dealing with constraints (5), we analyzed the alternative nonlinear, nonconvex, formulation of the RLM

$$\min_{\omega,b,\xi,\bar{z}} \frac{1}{2} \sum_{j=1}^{d} \omega_j^2 + \frac{C}{n} \left(\sum_{i=1}^n \xi_i + 2 \sum_{i=1}^n (1 - \bar{z}_i) \right)$$
(10)

s.t.

$$(y_i(\omega^{\top} x_i + b) - 1 + \xi_i) \cdot \bar{z}_i \ge 0 \quad \forall i = 1, \dots, n$$
 (11)

$$0 \le \xi_i \le 2 \quad \forall i = 1, \dots, n \tag{12}$$

 $\bar{z} \in \{0, 1\}^n$ (13)

$$\omega \in \mathbb{R}^d \tag{14}$$

$$b \in \mathbb{R}.$$
 (15)

Precisely as in RLM (4)–(9), binary variables are used to disable constraints (11), which replace constraints (5), but are the complemented version of *z* variables (13), i.e., $\bar{z}_i = 1 - z_i$. Namely, $\bar{z}_i = 1$ forces the *i*-th constraint to be active, thus allowing a maximum violation of $\xi_i = 2$, while $\bar{z}_i = 0$ disables the constraint in a classical "complementary" way.

Of course, constraints (11) are responsible for the nonconvexity of the MINLP model (10)–(15). However, its continuous version obtained by simply replacing constraints (13) with $\bar{z} \in [0, 1]^n$ is solved to (local) optimality by the Nonlinear Programming (NLP) solver IPOPT [22] providing a mixed binary solution that is very accurate, and relatively quick to compute. Indeed, it is easy to prove that

Proposition 2.1 *Any local optimal solution of the continuous version of model* (10)–(15) *is mixed binary.*

Proof The proposition is proven by contradiction. By definition, a local optimal solution $(\omega, b, \xi, \bar{z})$ satisfies constraints (11). For any i = 1, ..., n, either $\bar{z}_i = 0$ (integer), or $\bar{z}_i \in (0, 1)$. In the latter case, there exists an equivalent (still feasible) solution $(\omega, b, \xi, \bar{z}_1, ..., \bar{z}_{i-1}, \bar{z}_i + \epsilon, \bar{z}_{i+1}, ..., \bar{z}_n)$ having smaller objective function value for any $\epsilon \in (0, 1 - \bar{z}_i]$.

Proposition 2.1 above implies that the global optimal solution is mixed binary as well, thus solving the continuous version of the problem with a GO solver like Couenne solves the overall problem to optimality. On the other hand, Couenne is an

MINLP solver, hence it can handle integrality constraints on (a subset of) the variables. Thus, in the results presented in the next section we have kept integrality on variables \bar{z} .

3 A raw set of computational results

We have performed an exploratory test for the nonconvex MINLP formulation proposed in Sect. 2. We consider only the artificial datasets proposed by [8], to be able to control both the dataset and the problem size. In this paper we focus on a challenging subset of instances, namely 23 instances of size n = 100, d = 2, TypeB proposed in [8].

Our first set of experiments compares three state-of-the-art commercial MIP solvers, namely FICO Xpress 7.8, Gurobi 6.0.2 and IBM-Cplex 12.6.1, applied to model (4)–(9). These experiments were run on an Intel i5-750 CPU running at 2.67 GHz with 4 threads, imposing a time limit of 1 h per instance. As our model contains *big-M* coefficients, all solvers were run with zero integer feasibility tolerance. Table 1 reports computing times (time), number of nodes (nodes), percentage gap¹ of the upper (ub) and lower (lb) bounds. All computing times are expressed in CPU seconds; in case the time limit is hit, the entry in the column "time" indicates a "×". Gaps associated with instances solved to optimality are reported as "–".

The results in Table 1 clearly show that MIP solvers are quite ineffective on this class of instances: they are typically able to compute the optimal primal solution (especially Gurobi and IBM-Cplex), though in many cases a considerable dual gap may remain at the time limit. We have run the solvers with default algorithmic setting, i.e., a branch-and-cut based on solving a QP relaxation at every node of the branch-and-bound tree. Since there are other alternatives to solve MIQPs (the main ones being variants of outer approximation [16,27], or hybrid approaches, see Bonami et al. [5] for algorithmic approaches to convex MINLPs), it is legitimate to ask if some alternative choice would lead to better results. To answer that question from a methodological standpoint (instead of running quite a large number of tests with different parameter setting for each commercial MIP solver) we ran the open-source convex MINLP solver Bonmin [4,7] by using the three available algorithmic options, namely B-BB, B-OA and B-Hyb. The NLP-based branch and bound B-BB gives by far the best results by consistently computing decent upper bounds but it is unable to solve to optimality any of the instances. Both the other algorithms, namely B-OA and B-Hyb, encounter numerical problems and routinely fail. In view of these results it appears likely that the QP-based branch and bound applied by MIP solvers is indeed the best approach and is faster than Bonmin B-BB because they do exploit the fact that the objective function is quadratic and the constraints linear, while Bonmin does not.

The second set of raw experiments compares $IBM-Cplex^2$ with Couenne executed "out-of-the-box" on model (10)–(15). Quite surprisingly, this latter approach

¹ Percentage gaps are computed with respect to the optimal solution value, which is, instead, reported in Table 2.

² Because the results of the commercial MIP solvers are reasonably aligned, IBM-Cplex being, according to Table 1, better than FICO Xpress and Gurobi, we concentrate in the rest of the paper on IBM-Cplex and we assume the improvements of Sect. 6 apply to FICO Xpress and Gurobi as well.

Table 1 Computational result	s for state-of-the-art commerci	al MIP solvers	
FICO Xpress	Gurobi	IBM-Cplex	

	FICO Xpress				Gurobi				IBM-Cplex			
	Time	Nodes	% gap		Time	Nodes	% g	gap	Time	Nodes	% g	gap
			ub	lb			ub	lb			ub	lb
1	1658.12	5726k	_	_	601.45	8563k	_	_	1510.54	15,928k	_	_
2	×	14,116k	10.81	18.15	×	55,748k	_	6.46	×	28,467k	_	17.20
3	×	10,050k	9.35	37.89	×	24,464k	-	31.70	×	23,375k	_	37.63
4	×	13,159k	-	4.64	×	450k	_	42.66	2294.11	26,952k	-	-
5	×	11,420k	11.00	25.19	×	37,511k	-	17.49	×	23,821k	-	21.31
6	×	13,910k	5.88	17.65	×	106k	-	70.58	×	29,551k	-	16.14
7	×	15,352k	-	17.00	×	136k	-	58.82	×	30,959k	-	13.91
8	2894.31	10,900k	-	-	261.37	26k	-	-	3104.27	34,784k	-	-
9	×	14,541k	5.88	17.58	×	89k	-	83.11	×	27,976k	-	17.37
10	×	14,592k	5.87	17.31	×	106k	-	35.29	×	30,681k	-	16.10
11	×	14,082k	-	7.36	99.70	81k	-	-	3526.42	37,986k	-	-
12	×	14,350k	5.86	17.53	×	210k	-	52.82	×	25,958k	-	17.51
13	×	14,503k	-	16.04	×	943k	-	38.44	×	30,902k	-	15.85
14	×	14,167k	-	8.43	1767.02	15,609k	-	-	3,103.51	35,402k	-	-
15	×	14,598k	5.83	17.02	×	833k	-	37.72	×	29,375k	-	13.93
16	×	14,873k	5.84	16.87	×	1006k	-	38.28	×	30,139k	-	14.63
17	×	14,707k	-	9.02	×	178k	-	50.01	3,368.02	37,750k	-	-
18	×	14,363k	5.80	17.55	×	120k	-	70.61	×	27,444k	-	17.55
19	×	14,578k	0.03	16.40	×	49,620k	-	9.04	×	30,548k	-	16.30
20	×	14,684k	1.68	9.14	×	261k	-	48.54	3,249.70	36,008k	-	-
21	×	14,681k	5.67	17.66	×	1356k	-	36.33	×	25,638k	-	17.73
22	×	14,395k	-	16.04	×	34,996k	-	11.90	×	30,175k	-	16.68
23	×	14,080k	-	9.75	1456.30	16,214k	-	-	3578.31	38,990k	-	_

Instances of TypeB proposed by [8], n = 100, time limit of 1 h, executed on an Intel i5-750 CPU running at 2.67 GHz with 4 threads

performs better than running IBM-Cplex on the MIQP formulation. Table 2 gives this comparison, and provides the same information as in Table 1, plus the optimal value of each instance (for future reference). All these experiments were executed on an Intel Xeon E3-1220V2 with a time limit of 1 h per run. As Couenne can be executed in sequential mode only, IBM-Cplex was also run in single-thread mode on the same hardware. For IBM-Cplex, we also set to zero the integer feasibility tolerance parameters, the remaining parameters being left to their default.

The results of Table 2 are quite straightforward to interpret, with a strict dominance of Couenne with respect to IBM-Cplex: the former solves all instances but one, while the latter can solve only three of them; for these instances, IBM-Cplex requires a considerably larger computing time and number of nodes with respect to Couenne. The above results suggest that a GO solver is intrinsically more suitable for solving this class of problems. In the unique instance Couenne is unable to solve to optimality (instance 3) the issue is probably that the upper bound is not improved enough,

	Optimal value	Couenne				IBM-Cpl			
		Time	Nodes	% gap		Time	Nodes	% g:	ap
				ub	lb			ub	lb
1	157,994.959	151.87	22,574	_	-	2201.36	14,596,206	_	_
2	179,368.534	595.06	108,682	-	-	×	22,055,893	-	18.04
3	220,673.592	×	646,606	4.59	10.60	×	20,502,574	-	37.77
4	5,225.994	142.95	22,934	-	-	2106.66	11,874,373	-	-
5	5,957.083	1180.24	251,678	-	-	×	18,024,207	-	22.97
6	11,409,617.494	1571.70	257,749	-	_	×	19,929,191	_	17.28
7	11,409,058.363	698.72	120,425	-	_	×	25,499,307	_	13.84
8	10,737,725.660	448.38	70,342	-	_	×	24,616,925	_	6.92
9	5,705,364.054	1,433.84	238,028	-	_	×	21,880,932	_	15.83
10	5,704,804.923	578.12	104,368	-	_	×	25,439,344	_	12.60
11	5,369,016.540	348.82	56,161	-	_	3474.71	24,885,164	_	-
12	2,853,237.334	1637.12	267,588	-	-	×	21,428,199	-	17.30
13	2,852,678.203	565.32	101,677	-	-	×	24,928,209	-	16.92
14	2,684,661.980	340.74	54,212	-	-	×	21,016,707	-	7.81
15	1,427,173.974	1508.68	247,860	-	_	×	19,753,234	_	17.15
16	1,426,614.843	525.55	93,268	-	_	×	22,589,650	_	17.41
17	1,342,484.700	394.45	61,247	-	_	×	23,383,761	_	3.41
18	714,142.294	1156.81	186,351	-	-	×	18,182,507	-	15.06
19	713,583.163	513.23	91,329	-	_	×	23,144,030	_	11.56
20	671,396.060	498.46	77,747	-	_	×	26,976,870	_	6.60
21	357,626.454	1084.87	180,408	-	_	×	21,317,307	_	17.85
22	357,067.323	669.17	117,288	-	_	×	23,654,707	_	11.61
23	335,851.740	448.92	71,110	-	-	×	23,722,223	-	7.14

Table 2 Computational results for Couenne and IBM-Cplex

Instances of TypeB proposed by [8], n = 100, time limit of 1 h, executed on an Intel Xeon E3-1220V2

thus being unable to propagate (see next section) and strengthen the formulation. Conversely, IBM-Cplex is always able to find the right upper bound (namely, the optimal solution value) but the lower bound value remains far from the optimal value, thus being unable to prove optimality. Comparison between Tables 1 and 2 shows that a considerable deterioration of IBM-Cplex's performance is experienced, which is mainly due to the fact that the latter refers to single-thread runs.

4 Why are these results surprising?

Although, as anticipated in the introduction, convex MIQP solvers should be more effective than GO ones especially because they can exploit the very sophisticated MIP machinery, one can still argue that a comparison in performance between two different solution methods and computer codes is anyway hard to perform. However, digging

into the way Couenne solves the problem leads to confirm the initial surprise or even increase it.

4.1 McCormick linearization

The first observation is that the way constraints (11) are managed by Couenne is through the classical McCormick linearization [25]. Namely, for i = 1, ..., n, two new auxiliary variables ϑ_i and u_i are introduced that are associated with expressions in (11):

1.
$$\vartheta_i = y_i(\omega^T x_i + b) - 1 + \xi_i$$
, with $\vartheta_i^L \le \vartheta_i \le \vartheta_i^U$
2. $u_i = \vartheta_i \bar{z}_i$.

Then, the product corresponding to each new variable u_i is linearized as

$$u_i \ge 0 \tag{16}$$

$$u_i \ge \vartheta_i^L \, \bar{z}_i \tag{17}$$

$$u_i \ge \vartheta_i + \vartheta_i^U \bar{z}_i - \vartheta_i^U \tag{18}$$

$$u_i \le \vartheta_i + \vartheta_i^L \, \bar{z}_i - \vartheta_i^L \tag{19}$$

$$u_i \le \vartheta_i^U \bar{z}_i, \tag{20}$$

again for i = 1, ..., n, where (16) corresponds precisely to (11), and (17)–(20) are the McCormick envelopes. Essentially, setting $\bar{z}_i = 0$ again deactivates constraint *i* by simply enforcing the loose $\vartheta_i \in [\vartheta_i^L, \vartheta_i^U]$, where ϑ_i^L plays the role of the *big-M*.

In other words, Couenne initially builds a linear big-M formulation itself, with the difference that a specific ϑ_i^L value for each i is computed. Although such an internal computation is not responsible for the higher effectiveness of Couenne (typically Couenne is more conservative than the static values of M used in the literature and especially by [8]) this is, in practice, not a negligible issue because a safe value of M is not trivial to be determined a priori.

In the next section we will extensively discuss how McCormick inequalities are strengthened, as well as the bounds on ϑ variables. This will be shown to be crucial for Couenne but, at first, this similarity confirms the surprise.

4.2 Branching

It is well known that a major component of GO solvers is the iterative tightening of the convex (most of the time linear) relaxation of the nonconvex feasible region by branching on continuous variables (see, e.g., [3]). Another surprising fact here is that the default version of Couenne does not take advantage of this possibility and branches on the binary variables \bar{z} 's. Because everything is linear (after McCormick linearization) and the objective function convex, as soon as all binaries are fixed the problem is solved.

Even better performance for Couenne could be obtained by branching on continuous variables. Namely, instructed to branch preferably on continuous variables,

$ \overline{\text{Time (s)}} \overline{\text{Nodes}} \frac{\% \text{ gap}}{ub} \overline{\text{Iime (s)}} \overline{\text{Iime (s)}} \overline{\text{Nodes}} \frac{\% \text{ gap}}{ub} \overline{\text{Iime (s)}} \overline{\text{Iime (s)}$		Optimal value	Couenne	default			Couenne continuous			
ublbub1 $157,994.959$ 151.87 $22,574$ $ 257.54$ 74.424 $-$ 2 $179,368.534$ 595.06 $108,682$ $ 479.86$ 140.405 $-$ 3 $220,673.592$ \times $646,606$ 4.59 10.60 774.60 $216,118$ $-$ 4 $5,225.994$ 142.95 $22,934$ $ 408.33$ $121,896$ $-$ 5 $5,957.083$ 1180.24 $251,678$ $ 724.94$ $209,346$ $-$ 6 $11,409,617.494$ 1571.70 $257,749$ $ 640.58$ $184,391$ $-$ 7 $11,409,058.363$ 698.72 $120,425$ $ 912.07$ $269,332$ $-$ 8 $10,737,725.660$ 448.38 $70,342$ $ 492.00$ $143,426$ $-$ 9 $5,705,364.054$ 1433.84 $238,028$ $ 609.29$ $177,236$ $-$ 10 $5,704,804.923$ 578.12 $104,368$ $ 921.80$ $276,704$ $-$ 11 $5,369,016.540$ 348.82 $56,161$ $ 505.37$ $145,282$ $-$ 12 $2,853,237.334$ 1637.12 $267,588$ $ 855.99$ $255,654$ $-$ 14 $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $-$ 15 $1,427,173.974$ 1508.68 247			Time (s)	Nodes	% gap		Time (s)	Nodes	% gaj	р
$\begin{array}{cccccccccccccccccccccccccccccccccccc$					ub	lb			ub	lb
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	1	157,994.959	151.87	22,574	_	_	257.54	74,424	_	_
3 $220,673.592$ \times $646,606$ 4.59 10.60 774.60 $216,118$ $-$ 4 $5,225.994$ 142.95 $22,934$ $ 408.33$ $121,896$ $-$ 5 $5,957.083$ 1180.24 $251,678$ $ 724.94$ $209,346$ $-$ 6 $11,409,617.494$ 1571.70 $257,749$ $ 640.58$ $184,391$ $-$ 7 $11,409,058.363$ 698.72 $120,425$ $ 912.07$ $269,332$ $-$ 8 $10,737,725.660$ 448.38 $70,342$ $ 492.00$ $143,426$ $-$ 9 $5,705,364.054$ 1433.84 $238,028$ $ 609.29$ $177,236$ $-$ 10 $5,704,804.923$ 578.12 $104,368$ $ 921.80$ $276,704$ $-$ 11 $5,369,016.540$ 348.82 $56,161$ $ 505.37$ $145,282$ $-$ 12 $2,853,237.334$ 1637.12 $267,588$ $ 624.38$ $176,251$ $-$ 13 $2,852,678.203$ 565.32 $101,677$ $ 855.99$ $255,654$ $-$ 14 $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $-$ 15 $1,427,173.974$ 1508.68 $247,860$ $ 571.29$ $163,552$ $-$ 16 $1,426,614.843$ 525.55 $93,268$ $ 811.76$ <	2	179,368.534	595.06	108,682	_	_	479.86	140,405	_	_
4 $5,225.994$ 142.95 $22,934$ $ 408.33$ $121,896$ $-$ 5 $5,957.083$ 1180.24 $251,678$ $ 724.94$ $209,346$ $-$ 6 $11,409,617.494$ 1571.70 $257,749$ $ 640.58$ $184,391$ $-$ 7 $11,409,058.363$ 698.72 $120,425$ $ 912.07$ $269,332$ $-$ 8 $10,737,725.660$ 448.38 $70,342$ $ 492.00$ $143,426$ $-$ 9 $5,705,364.054$ 1433.84 $238,028$ $ 609.29$ $177,236$ $-$ 10 $5,704,804.923$ 578.12 $104,368$ $ 921.80$ $276,704$ $-$ 11 $5,369,016.540$ 348.82 $56,161$ $ 505.37$ $145,282$ $-$ 12 $2,853,237.334$ 1637.12 $267,588$ $ 624.38$ $176,251$ $-$ 13 $2,852,678.203$ 565.32 $101,677$ $ 855.99$ $255,654$ $-$ 14 $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $-$ 15 $1,427,173.974$ 1508.68 $247,860$ $ 571.29$ $163,552$ $-$ 16 $1,426,614.843$ 525.55 $93,268$ $ 811.76$ $251,205$ $-$ 17 $1,342,484.700$ 394.45 $61,247$ $ -$ <	3	220,673.592	×	646,606	4.59	10.60	774.60	216,118	_	_
$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	4	5,225.994	142.95	22,934	_	_	408.33	121,896	_	_
6 $11,409,617.494$ 1571.70 $257,749$ $ 640.58$ $184,391$ $ 7$ $11,409,058.363$ 698.72 $120,425$ $ 912.07$ $269,332$ $ 8$ $10,737,725.660$ 448.38 $70,342$ $ 492.00$ $143,426$ $ 9$ $5,705,364.054$ 1433.84 $238,028$ $ 609.29$ $177,236$ $ 10$ $5,704,804.923$ 578.12 $104,368$ $ 921.80$ $276,704$ $ 11$ $5,369,016.540$ 348.82 $56,161$ $ 505.37$ $145,282$ $ 12$ $2,853,237.334$ 1637.12 $267,588$ $ 624.38$ $176,251$ $ 13$ $2,852,678.203$ 565.32 $101,677$ $ 855.99$ $255,654$ $ 14$ $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $ 15$ $1,427,173.974$ 1508.68 $247,860$ $ 571.29$ $163,552$ $ 16$ $1,426,614.843$ 525.55 $93,268$ $ 811.76$ $251,205$ $ 17$ $1,342,484.700$ 394.45 $61,247$ $ 513.45$ $148,947$ $ 19$ $713,583.163$ 513.23 $91,329$ $ 722.22$ $219,986$ $ 20$ $671,396.060$ 498.46 $77,74$	5	5,957.083	1180.24	251,678	_	_	724.94	209,346	_	_
7 $11,409,058.363$ 698.72 $120,425$ $ 912.07$ $269,332$ $-$ 8 $10,737,725.660$ 448.38 $70,342$ $ 492.00$ $143,426$ $-$ 9 $5,705,364.054$ 1433.84 $238,028$ $ 609.29$ $177,236$ $-$ 10 $5,704,804.923$ 578.12 $104,368$ $ 921.80$ $276,704$ $-$ 11 $5,369,016.540$ 348.82 $56,161$ $ 505.37$ $145,282$ $-$ 12 $2,853,237.334$ 1637.12 $267,588$ $ 624.38$ $176,251$ $-$ 13 $2,852,678.203$ 565.32 $101,677$ $ 855.99$ $255,654$ $-$ 14 $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $-$ 15 $1,427,173.974$ 1508.68 $247,860$ $ 571.29$ $163,552$ $-$ 16 $1,426,614.843$ 525.55 $93,268$ $ 811.76$ $251,205$ $-$ 17 $1,342,484.700$ 394.45 $61,247$ $ 513.45$ $148,947$ $-$ 19 $713,583.163$ 513.23 $91,329$ $ 722.22$ $219,986$ $-$ 20 $671,396.060$ 498.46 $77,747$ $ 382.18$ $118,811$ $-$ 21 $357,626.454$ 1084.87 $180,408$ $ 459.15$	6	11,409,617.494	1571.70	257,749	_	-	640.58	184,391	_	-
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	7	11,409,058.363	698.72	120,425	_	_	912.07	269,332	_	_
$\begin{array}{rrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrr$	8	10,737,725.660	448.38	70,342	_	-	492.00	143,426	_	-
10 $5,704,804.923$ 578.12 $104,368$ 921.80 $276,704$ -11 $5,369,016.540$ 348.82 $56,161$ 505.37 $145,282$ -12 $2,853,237.334$ 1637.12 $267,588$ 624.38 $176,251$ -13 $2,852,678.203$ 565.32 $101,677$ 855.99 $255,654$ -14 $2,684,661.980$ 340.74 $54,212$ 466.56 $138,252$ -15 $1,427,173.974$ 1508.68 $247,860$ 571.29 $163,552$ -16 $1,426,614.843$ 525.55 $93,268$ 811.76 $251,205$ -17 $1,342,484.700$ 394.45 $61,247$ 361.27 $108,762$ -18 $714,142.294$ 1156.81 $186,351$ 513.45 $148,947$ -19 $713,583.163$ 513.23 $91,329$ 722.22 $219,986$ -20 $671,396.060$ 498.46 $77,747$ 382.18 $118,811$ -21 $357,626.454$ 1084.87 $180,408$ 459.15 $133,964$ -22 $357,067.323$ 669.17 $117,288$ 611.38 $185,459$ -	9	5,705,364.054	1433.84	238,028	_	-	609.29	177,236	_	-
11 $5,369,016.540$ 348.82 $56,161$ $ 505.37$ $145,282$ $-$ 12 $2,853,237.334$ 1637.12 $267,588$ $ 624.38$ $176,251$ $-$ 13 $2,852,678.203$ 565.32 $101,677$ $ 855.99$ $255,654$ $-$ 14 $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $-$ 15 $1,427,173.974$ 1508.68 $247,860$ $ 571.29$ $163,552$ $-$ 16 $1,426,614.843$ 525.55 $93,268$ $ 811.76$ $251,205$ $-$ 17 $1,342,484.700$ 394.45 $61,247$ $ 361.27$ $108,762$ $-$ 18 $714,142.294$ 1156.81 $186,351$ $ 513.45$ $148,947$ $-$ 19 $713,583.163$ 513.23 $91,329$ $ 722.22$ $219,986$ $-$ 20 $671,396.060$ 498.46 $77,747$ $ 382.18$ $118,811$ $-$ 21 $357,626.454$ 1084.87 $180,408$ $ 459.15$ $133,964$ $-$ 22 $357,067.323$ 669.17 $117,288$ $ 611.38$ $185,459$ $-$	10	5,704,804.923	578.12	104,368	_	-	921.80	276,704	_	-
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	11	5,369,016.540	348.82	56,161	-	-	505.37	145,282	-	_
13 $2,852,678.203$ 565.32 $101,677$ $ 855.99$ $255,654$ $-$ 14 $2,684,661.980$ 340.74 $54,212$ $ 466.56$ $138,252$ $-$ 15 $1,427,173.974$ 1508.68 $247,860$ $ 571.29$ $163,552$ $-$ 16 $1,426,614.843$ 525.55 $93,268$ $ 811.76$ $251,205$ $-$ 17 $1,342,484.700$ 394.45 $61,247$ $ 361.27$ $108,762$ $-$ 18 $714,142.294$ 1156.81 $186,351$ $ 513.45$ $148,947$ $-$ 19 $713,583.163$ 513.23 $91,329$ $ 722.22$ $219,986$ $-$ 20 $671,396.060$ 498.46 $77,747$ $ 382.18$ $118,811$ $-$ 21 $357,626.454$ 1084.87 $180,408$ $ 459.15$ $133,964$ $-$ 22 $357,067.323$ 669.17 $117,288$ $ 611.38$ $185,459$ $-$	12	2,853,237.334	1637.12	267,588	_	-	624.38	176,251	_	-
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	13	2,852,678.203	565.32	101,677	_	-	855.99	255,654	_	-
$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	14	2,684,661.980	340.74	54,212	-	-	466.56	138,252	-	-
$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	15	1,427,173.974	1508.68	247,860	-	-	571.29	163,552	-	-
$\begin{array}{cccccccccccccccccccccccccccccccccccc$	16	1,426,614.843	525.55	93,268	_	-	811.76	251,205	_	-
18 714,142.294 1156.81 186,351 - - 513.45 148,947 - 19 713,583.163 513.23 91,329 - - 722.22 219,986 - 20 671,396.060 498.46 77,747 - - 382.18 118,811 - 21 357,626.454 1084.87 180,408 - - 459.15 133,964 - 22 357,067.323 669.17 117,288 - - 611.38 185,459 -	17	1,342,484.700	394.45	61,247	_	-	361.27	108,762	_	-
19 713,583.163 513.23 91,329 - - 722.22 219,986 - 20 671,396.060 498.46 77,747 - - 382.18 118,811 - 21 357,626.454 1084.87 180,408 - - 459.15 133,964 - 22 357,067.323 669.17 117,288 - - 611.38 185,459 -	18	714,142.294	1156.81	186,351	_	-	513.45	148,947	_	-
20 671,396.060 498.46 77,747 - - 382.18 118,811 - 21 357,626.454 1084.87 180,408 - - 459.15 133,964 - 22 357,067.323 669.17 117,288 - - 611.38 185,459 -	19	713,583.163	513.23	91,329	_	-	722.22	219,986	_	-
21 357,626.454 1084.87 180,408 - - 459.15 133,964 - 22 357,067.323 669.17 117,288 - - 611.38 185,459 -	20	671,396.060	498.46	77,747	-	-	382.18	118,811	_	_
22 357,067.323 669.17 117,288 611.38 185,459 -	21	357,626.454	1084.87	180,408	_	-	459.15	133,964	_	-
	22	357,067.323	669.17	117,288	-	_	611.38	185,459	_	-
23 335,851.740 448.92 71,110 404.87 116,966 -	23	335,851.740	448.92	71,110	-	-	404.87	116,966	-	-

 Table 3
 Computational results for Couenne default and Couenne branching emphasis on continuous variables

Instances of TypeB proposed by [8], n = 100, time limit of 1 h, executed on an Intel Xeon E3-1220V2

Couenne always selects ϑ variables, which clearly leads to additional bound tightening with respect to branch on binaries. Indeed, if in a given relaxation we have $\vartheta_i = c$, the two branches $\vartheta_i \leq c \text{ OR } \vartheta_i \geq c$ propagate as follows: (*i*) if c < 0 then $\vartheta_i \leq c$ implies $\overline{z}_i = 0 \text{ OR } (ii)$ if c > 0 then $\vartheta_i \geq c$ implies $\overline{z}_i = 1$. The results are reported in Table 3 and clearly show the computational advantage of this choice. Thus, again, it is surprising that the default branching strategy of Couenne leads to an improvement over the sophisticated branching framework of IBM-Cplex.

4.3 L₁ norm

A natural question is if the results reported in Table 2 are due to the somehow less sophisticated evolution of IBM-Cplex in its MIQP extension with respect to the

MILP one. In order to answer this question we performed an experiment in which the quadratic part of the objective function was replaced by the L_1 norm on ω . More precisely, the sum of the absolute values of ω_i is minimized (and linear constraints to deal with the absolute value are added). This results in a pure MILP once the *big-M* constraints (5) are considered (solved by IBM-Cplex) or a nonconvex MINLP with linear objective function (solved by Couenne) if constraints (11) are used instead. This linear objective variant of the RLM does not result in any change for the comparison and Couenne continues achieving better results than IBM-Cplex.

5 Bound reduction in nonconvex MINLP problems

Bound reduction is a crucial tool in MINLP: it allows one to eliminate portions of the feasible set while guaranteeing that at least one optimal solution is retained. Although its origins can be traced back to Artificial Intelligence (see [15]), it finds wide application in Constraint Programming and in solvers for both Nonlinear Optimization [26] and for MIP problems [1,28].

Consider a generic optimization problem $\min\{f(x) : x \in X, \ell \le x \le u\}$, where $X \subset \mathbb{R}^n$ and $x, \ell, u \in \mathbb{R}^n$. Also, suppose an upper bound $U \in \mathbb{R} \cup \{+\infty\}$ on the objective function value of the optimal solution is available: if $U < +\infty$, then a feasible solution $x \in X \cap [\ell, u]$ is available such that f(x) = U. Bound reduction attempts to find tighter lower bounds $\ell'_i > \ell_i$ and upper bounds $u'_i < u_i$. In general, a good upper bound is often the key to a strong bound reduction.

An ideal bound reduction procedure obtains bounds by exploiting the full problem structure:

$$\ell'_{i} = \max \{ x_{i} : x \in X, \ell \le x \le u, f(x) \le U \}, u'_{i} = \min \{ x_{i} : x \in X, \ell \le x \le u, f(x) \le U \}.$$
(21)

However, the 2n optimization problems above can be as hard as the original optimization problem itself, therefore this approach is impractical.

A fast bound reduction procedure, known as *Feasibility Based Bound Tightening* (FBBT), yields new bounds on a variable x_i using bounds on other variables that are linked to x_i through a constraint or the objective function. For instance, the constraint $x_1x_2 \le 4$ and the bounds $x_1 \ge 1$, $x_2 \ge 1$ yield new upper bounds $x_1 \le 4$, $x_2 \le 4$. An example that is closer to our application is the constraint $x_i^2 \le u$, where $u \ge 0$, which obviously implies $x_i \in [-\sqrt{u}, \sqrt{u}]$.

A specialized version of this procedure applies to affine functions, and is commonly used in MIP solvers, see [1]. Consider the range constraint

$$\ell_0 \le \alpha_0 + \sum_{j=1}^n \alpha_j x_j \le u_0.$$

Define $J^+ = \{j = 1, ..., n : \alpha_j > 0\}$ and $J^- = \{j = 1, ..., n : \alpha_j < 0\}$. Bounds ℓ_0, u_0 on the expression imply new (possibly tighter) bounds ℓ'_j, u'_j on $x_j, j = 1, ..., n : \alpha_j \neq 0$:

$$\forall j : \alpha_{j} > 0, \ \ell_{j}' = \frac{1}{\alpha_{j}} \left(\ell_{0} - \left(\alpha_{0} + \sum_{i \in J^{+} \setminus \{j\}} \alpha_{i} u_{i} + \sum_{i \in J^{-}} \alpha_{i} \ell_{i} \right) \right),$$

$$u_{j}' = \frac{1}{\alpha_{j}} \left(u_{0} - \left(\alpha_{0} + \sum_{i \in J^{+} \setminus \{j\}} \alpha_{i} \ell_{i} + \sum_{i \in J^{-}} \alpha_{i} u_{i} \right) \right);$$

$$\forall j : \alpha_{j} < 0, \ \ell_{j}' = \frac{1}{\alpha_{j}} \left(u_{0} - \left(\alpha_{0} + \sum_{i \in J^{+}} \alpha_{i} \ell_{i} + \sum_{i \in J^{-} \setminus \{j\}} \alpha_{i} u_{i} \right) \right),$$

$$u_{j}' = \frac{1}{\alpha_{j}} \left(\ell_{0} - \left(\alpha_{0} + \sum_{i \in J^{+}} \alpha_{i} u_{i} + \sum_{i \in J^{-} \setminus \{j\}} \alpha_{i} \ell_{i} \right) \right).$$

$$(22)$$

5.1 Applying bound reduction to model (10)–(15)

Couenne is a branch-and-bound solver for MINLP problems that uses, among others, several bound reduction techniques, including FBBT. At the beginning, Couenne runs a greedy rounding procedure to obtain a feasible solution of the problem, and hence an upper bound U. Applying FBBT using two rules mentioned above (for affine functions and for the square operator) yields tight bounds on ω_i at the root node of Couenne's branch-and-bound tree. Consider the objective function of our problem

$$\frac{1}{2}\sum_{j=1}^{d}\omega_{j}^{2} + \frac{C}{n}\left(\sum_{i=1}^{n}\xi_{i} + 2\sum_{i=1}^{n}(1-\bar{z}_{i})\right),$$

which is bounded from above by U. Also, note that $\sum_{i=1}^{n} \xi_i + 2 \sum_{i=1}^{n} (1 - \overline{z}_i)$ is nonnegative. Since $\sum_{j=1}^{d} \omega_j^2 \ge 0$ and $\sum_{i=1}^{n} \xi_i + 2 \sum_{i=1}^{n} (1 - \overline{z}_i) \ge 0$, and also ξ_i and \overline{z}_i are nonnegative, we have

$$\omega_i \in \left[-\sqrt{2U}, \sqrt{2U}\right] \quad \forall i = 1, \dots, d.$$

A related observation concerns the constraints of our model. The tighter bounds on the ω_i 's variables, which are initially unbounded in the definition of the problem, do not seem to have an influence on constraints (11), where the variable *b* remains unbounded. This family of nonlinear constraints can be simplified to $\vartheta_i \bar{z}_i \ge 0$, with $\vartheta_i = (y_i(\omega^T x_i + b) - 1 + \xi_i) \in [-\infty, +\infty]$ and $\bar{z}_i \in \{0, 1\}$, for all i = 1, ..., n. Due to the infinite bounds, this constraint does not admit a linear relaxation, which is useful for any MINLP solver to obtain a lower bound. When imposing fictitious bounds [-M, M] on ϑ_i , with large enough *M*, one gets the constraint $\vartheta_i \ge M(\bar{z}_i - 1)$, which is the *big-M* constraint used in MIP. In these cases, *probing* techniques can be of help. A probing bound reduction algorithm works as follows: impose a fictitious upper bound $\lambda_i \in (\ell_i, u_i)$ on a variable x_i , thereby restricting x_i to $[\ell_i, \lambda_i]$. If the restricted problem can be proven (through FBBT, for example) to be infeasible or to have a lower bound that is above a cutoff *U*, then no feasible solution with better objective function value can be found in the restriction. Therefore, the new lower bound $x_i \ge \lambda_i$ is valid. This procedure can be applied to tighten the upper bound as well, by imposing a fictitious lower bound $\mu_i \in (\ell_i, u_i)$. Although applying it to all variables is time consuming, it is especially useful for unbounded variables. Probing is a common tightening technique in MIP and MINLP solvers (see [28], and [3] and [30], respectively.)

We will now describe the specific reductions that Couenne applied to the RLM, and we will computationally show that these reductions are crucial to obtain the results shown in Sect. 4.

5.2 Strengthening McCormick constraints

The coefficients of the McCormick constraints are the lower and upper bounds on the variables involved. Hence, these constraints can be replaced by stronger ones as soon as tighter bounds on the variables are available. Couenne does this automatically by means of a cut separator that is called at every branch-and-bound node, and only adds tighter McCormick cuts if they are violated by the Linear Programming (LP) solution available at that node.

Note that McCormick cuts are only useful if both variables ϑ_i and \bar{z}_i are not fixed, as, otherwise, the constraint $u_i = \vartheta_i \bar{z}_i$ becomes linear. Of course, new McCormick cuts are making previous ones redundant and Couenne relies on the branch-and-bound manager (specifically, Cbc [10]) to deal with redundancy in the constraint set.

While looking for reasons of Couenne's performance, we have run an experiment where McCormick cuts were only added to the initial LP relaxation but excluded from separation at all nodes. The performance worsened dramatically on all instances, which indicates that the bound on the involved variables u_i , ϑ_i , \bar{z}_i is tightened and should be exploited.

5.3 Bound tightening

Couenne uses several techniques for bound tightening among those described above. In the context of this problem, tightening is based on the following elements:

- the objective function, if a cutoff U is available;
- the definition $u_i = \vartheta_i \bar{z}_i (\geq 0)$ and related constraint $u_i \geq 0$;
- the definition $\vartheta_i = (y_i(\omega^\top x_i + b) 1 + \xi_i).$

Propagation possibly generates new bounds on ω , b, ξ , which are obtained through standard presolve procedures [1]. For the sake of completeness, we add them here by defining the coefficients $\alpha_{ik} = x_{ik}y_i$, where x_{ik} is the k-th feature (k = 1, ..., d) of object i (i = 1, ..., n). Also, we use superscripts "L" (resp. "U") to denote lower (resp. upper) bounds. The bounds are then updated by formulas

$$\begin{aligned} \forall k : \alpha_{ik} > 0 \ \omega_k^L &= \frac{1}{\alpha_{ik}} \left(\vartheta_i^L - \sum_{j:\alpha_{ij} < 0} \alpha_{ij} \omega_j^L - \sum_{j \neq k:\alpha_{ij} > 0} \alpha_{ij} \omega_j^U - (y_i b)^U - \xi^U + 1 \right), \\ \omega_k^U &= \frac{1}{\alpha_{ik}} \left(\vartheta_i^U - \sum_{j:\alpha_{ij} < 0} \alpha_{ij} \omega_j^U - \sum_{j \neq k:\alpha_{ij} > 0} \alpha_{ij} \omega_j^L - (y_i b)^L - \xi^L + 1 \right), \end{aligned}$$

🖉 Springer

$$\begin{aligned} \forall k : \alpha_{ik} < 0 \ \omega_k^L &= \frac{1}{\alpha_{ik}} \left(\vartheta_i^U - \sum_{j \neq k: \alpha_{ij} < 0} \alpha_{ij} \omega_j^U - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^L - (y_i b)^L - \xi^L + 1 \right), \\ \omega_k^U &= \frac{1}{\alpha_{ik}} \left(\vartheta_i^L - \sum_{j \neq k: \alpha_{ij} < 0} \alpha_{ij} \omega_j^L - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^U - (y_i b)^U - \xi^U + 1 \right), \\ \forall i : y_i > 0, \ b^L &= \frac{1}{y_i} \left(\vartheta_i^L - \sum_{j: \alpha_{ij} < 0} \alpha_{ij} \omega_j^L - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^L - \xi^L + 1 \right), \\ b^U &= \frac{1}{y_i} \left(\vartheta_i^U - \sum_{j: \alpha_{ij} < 0} \alpha_{ij} \omega_j^U - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^L - \xi^L + 1 \right), \\ \forall i : y_i < 0, \ b^L &= \frac{1}{y_i} \left(\vartheta_i^U - \sum_{j: \alpha_{ij} < 0} \alpha_{ij} \omega_j^U - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^L - \xi^L + 1 \right), \\ b^U &= \frac{1}{y_i} \left(\vartheta_i^L - \sum_{j: \alpha_{ij} < 0} \alpha_{ij} \omega_j^L - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^U - \xi^U + 1 \right), \\ \xi^L &= \vartheta_i^L - \sum_{j: \alpha_{ij} < 0} \alpha_{ij} \omega_j^L - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^U - (y_i b)^U + 1, \\ \xi^U &= \vartheta_i^U - \sum_{j: \alpha_{ij} < 0} \alpha_{ij} \omega_j^U - \sum_{j: \alpha_{ij} > 0} \alpha_{ij} \omega_j^L - (y_i b)^L + 1, \end{aligned}$$

where, as an example, the value ω_k^L is the maximum lower bound value on ω_k over all i = 1, ..., n of the first formula above. The same holds for all other bounds.

When solving an instance of our problem, tightening typically happens after obtaining a new integer feasible solution or after branching on a variable. We describe here in detail the tightening steps. Note that the sequence of tightening steps is often repeated to ensure "tight enough" bounds on all variables, especially the "critical" ones, i.e., those that can be free and unbounded, like ω 's and b in out case. This loop is of fundamental importance in GO solvers that, generally, cannot rely on tight continuous relaxations.

- After branching on a binary variable \bar{z}_i , if \bar{z}_i is fixed to 0 the lower bound on the objective function is increased, which may allow for extra tightening on the variables appearing in the objective or, if the lower bound is above the cutoff, for pruning the node.
- If \bar{z}_i is instead fixed to 1, Couenne uses the new lower bound on ϑ_i to obtain a better bound on ω and b.
- When a new upper bound U is found (Couenne finds a good one at the beginning), this triggers a tightening on ω . No tightening is done on any \bar{z}_i variable since all of them have the same coefficient in the objective.

Note that in both cases (branching on \bar{z}_i and new bound U) the tightening of ω 's and b above allows one to strengthen the McCormick inequalities, which are necessary

to build a linear relaxation of the problem. To summarize, both branches and a new cutoff value allow to tighten the bounds on ϑ , ω , and b, and this, in turn, allows for strengthening McCormick inequalities at every node.

Needless to say, disabling bound reduction in Couenne leads to a dramatic worsening of the performance. It appears therefore that both bound reduction and McCormick cuts are the responsible for the success of Couenne on this class of instances, and, conversely, the lack of those ingredients, or at least their "too light" use, seems to be crucial in the troubles encountered by MIP solvers. In the next section we will show that enhancing the MIP solvers in this line is possible.

6 Enhancing MIP solvers

Inspired by the outperforming results of Couenne over IBM-Cplex (and virtually all other MIP solvers), and by the analysis in the previous section, we have tried to exploit the successful MINLP tools, namely bound tightening, to deal with the weak MIP relaxations associated with *big-M* constraints. We will show two ways of enhancing the behavior of IBM-Cplex

- either by using online (and cheap) tightening of Sect. 5.3 and locally-valid McCormick constraints,
- or by an a priori (more expensive) strengthening of the formulation, where the bounds on ω's and b are computed by solving MIPs as in (21).

The former approach, outlined in Sect. 6.1, leans towards a real integration of aggressive bound tightening in MIP, while the latter can be seen as a sophisticated MIP algorithm that exploits bound tightening and is described in Sect. 6.2.

6.1 Local cuts

As explained in Section 5.2, one of the main ingredients used by Couenne to solve the classification instances is the iterative strengthening of constraints (5) in the tree. This process can be mimicked in a MIP solver by using locally valid versions of the big-M constraints (5).

Basically, any time one tightens the upper and lower bounds on ω and b through propagation, one can recompute a potentially smaller value of M by recomputing the minimal value the left-hand-side of constraint (5) can take under these new bounds. If this value is indeed smaller, one can add a tighter version of (5). Note that this new constraint strictly dominates the previous one. As explained in Sect. 5, the combination of branching and propagation gives tighter bounds for ω and b at the nodes of the branch-and-bound tree. However, since these bounds on ω and b are computed for a node of the tree, the resulting strengthened version of (5) is only valid for the sub-tree rooted in that node.

In IBM-Cplex, since version 12.6.1, locally valid versions of constraints (5) can be automatically added as *local cuts* within the branch-and-bound tree. These inequalities as referred to as *local implied bound cuts*.

The way in which IBM-Cplex generates these cuts is through the use of socalled *implications*. An implication is the logical description of an indicator constraint: t = 0 (or 1) $\Rightarrow \alpha^{\top} x \le x_0$. Implications can be either automatically found by presolve procedures that analyze the structure of the model or, alternatively, they can be directly given as an input to the solver (see, [17,21]). Implications naturally imply linear cuts by using their *big-M* form (2).

To assess the efficiency of local implied bound cuts on the supervised classification instances, we replaced the constraints (5) with the indicator constraints

$$z = 0 \Rightarrow y_i(\omega^{\top} x_i + b) \ge 1 - \xi_i \quad \forall i = 1, \dots, n.$$

We then solved the 23 instances of Table 1, using the same parameters for IBM-Cplex and imposing a very aggressive setting for separating local implied bound cuts. In this very aggressive settings, IBM-Cplex tries to recompute a smaller *big-M* for each indicator constraint at every node of the branch-and-bound tree. Every time a smaller *big-M* is found, the previous *big-M* constraint is removed and replaced by the tighter one.

Computing times in CPU seconds and number of nodes are reported in Table 4. IBM-Cplex is executed (using 4 threads), in the default settings and with local implied bound cuts set to very aggressive. While default IBM-Cplex is able to solve to optimality only 8 instances within the time limit, the enhanced version solves to optimality 17 out of 23 instances.

To test the effectiveness of the local implied bound cuts in IBM-Cplex 12.6.1 on a larger set of instances we considered the general MIPLIB2010 library and solved each instance without and with local implied bound cuts (default and aggressive settings, respectively). It is worth observing that these instances do not explicitly have indicator constraints, so local implied bound cuts are generated on the implications discovered by IBM-Cplex during the preprocessing phase. In addition, many of the instances might not even have either *big-Ms* or implications and, even when these implications exist, and are discovered by IBM-Cplex, they might be less crucial than in our SVM instances where all constraints are of that kind. Nevertheless, the aim of this experiment is to assert if local implied bound cuts could be beneficial in general as a first step to (i) devise effective heuristics to estimate the interest of generating these cuts and (ii) improve on the heuristics devoted to find effective implications.

The results of the experiment are reported in Table 5, whose first line concerns the entire testbed of 358 instances. The experiments were run on an Intel Xeon E3-1220V2 at 3.10 GHz using 4 threads with a time limit equal to 1 h and a memory limit of 8GB; all other parameters were left to their default values. For each setting of the solver, we give the number of instances solved to proven optimality and the average computing time and number of nodes, both in arithmetic and geometric means. The second line refers to the 200 instances that are solved to proven optimality in both settings. Geometric means are shifted by 1 second and 10 nodes for computing times and number of nodes, respectively.

To possibly reduce the erratic behaviour of MIP solvers (see, e.g., [18,24]) we selected the 18 instances for which adding local implied bound cuts produced a speedup larger than 1.5 (including 3 problems that were solved using local cuts and that were

	Optimal value	imal value Default Local implied bound cuts							
		Time (s)	Nodes	% g	ap	Time (s)	Nodes	% g	ap
				ub	lb			ub	lb
1	157,994.959	1510.54	15,928,227	_	_	264.41	1,249,006	_	_
2	179,368.534	×	28,467,361	_	17.20	×	15,213,623	_	9.43
3	220,673.592	×	23,375,230	_	37.63	×	9,862,566	_	26.05
4	5,225.994	2294.11	26,952,817	_	_	263.90	1,474,124	_	-
5	5,957.083	×	23,821,213	_	21.31	×	16,422,588	_	7.48
6	11,409,617.494	×	29,551,078	_	16.14	1694.17	9,146,159	_	_
7	11,409,058.363	×	30,959,870	_	13.91	577.51	3,408,553	_	_
8	10,737,725.660	3104.27	34,784,855	_	_	776.55	6,193,222	_	_
9	5,705,364.054	×	27,976,050	_	17.37	1682.58	11,230,059	_	-
10	5,704,804.923	×	30,681,973	_	16.10	574.09	4,128,400	_	_
11	5,369,016.540	3526.42	37,986,138	_	_	1026.54	8,189,774	_	_
12	2,853,237.334	×	25,958,736	_	17.51	2,149.25	13,497,192	_	_
13	2,852,678.203	×	30,902,369	_	15.85	485.87	3,069,338	_	-
14	2,684,661.980	3103.51	35,402,901	_	_	836.03	6,461,712	_	-
15	1,427,173.974	×	29,375,413	_	13.93	×	14,170,081	_	9.52
16	1,426,614.843	×	30,139,333	_	14.63	446.72	2,544,316	_	-
17	1,342,484.700	3368.02	37,750,319	_	_	715.83	5,632,229	_	-
18	714,142.294	×	27,444,832	_	17.55	×	13,783,012	_	12.08
19	713,583.163	×	30,548,097	_	16.30	721.45	5,138,203	_	_
20	671,396.060	3249.70	36,008,315	_	_	746.76	5,207,078	_	_
21	357,626.454	×	25,638,748	_	17.73	×	14,254,749	_	11.55
22	357,067.323	×	30,175,822	_	16.68	341.16	1,551,983	_	-
23	335,851.740	3578.31	38,990,037	_	_	1742.39	11,689,501	_	_

Table 4 Computational results for IBM-Cplex default and IBM-Cplex with local implied bound cuts

Instances of Type B proposed by [8], n = 100, executed on an Intel i5-750 CPU running at 2.67 GHz with 4 threads

Table 5 Computational results for IBM-Cplex default and IBM-Cplex with local implied bound cuts

	#inst.	Default					Local implied bound cuts				
	#opt		arit. mean		geom.	geom. mean		arit. mean		geom. mean	
			Time	Nodes	Time	Nodes		Time	Nodes	Time	Nodes
ALL	358	202	1,764	1,334,421	333	12,330	203	1,760	1,032,553	339	11,894
OPT	200	200	323	855,490	50	4,733	200	345	816,893	53	4,726

Instances of MIPLIB2010, executed using 4 threads on an Intel Xeon E3-1220V2 at 3.10 GHz

not solved otherwise) and solved each of them using 10 different random seeds. Table 6 gives the number of optimal solutions and average number of nodes and computing time for each setting of IBM-Cplex and each random seed, whereas Table 7 gives similar information for each setting of IBM-Cplex and each instance.

Seed	Default			Local imp	Local implied bound cuts				
	#opt.	Time	Nodes	#opt.	Time	Nodes			
0	15	1282	1,364,859	18	477	438,924			
1	16	1118	1,252,240	16	1025	1,075,938			
2	16	942	1,351,707	17	824	454,438			
3	16	863	725,954	17	1013	705,262			
4	17	810	1,495,415	15	1124	778,474			
5	16	1241	1,478,112	15	1134	1,276,767			
6	16	967	1,041,608	18	686	510,572			
7	16	758	707,348	16	1001	741,095			
8	17	924	1,462,432	17	925	579,676			
9	17	1004	1,082,111	18	578	365,557			
	162	991	1196,178	167	879	692,670			

Table 6 Computational results for IBM-Cplex default and IBM-Cplex with local implied bound cuts

Selected instances from MIPLIB2010, executed using 4 threads on an Intel Xeon E3-1220V2 at 3.10 GHz

Instance	Default			Local in	plied bound	cuts
	#opt.	Time	Nodes	#opt.	Time	Nodes
atlanta-ip	10	2600	25,299	10	2335	22,868
blp-ic97	10	398	209,959	10	403	205,085
glass4	10	506	1,757,770	10	195	631,617
gmu-35-40	10	37	400,522	10	44	440,854
gmut-77-40	10	515	524,877	10	428	399,682
lectsched-4-obj	10	5	2558	10	4	2795
neos-1440460	10	8	16,712	10	10	17,974
neos-1442119	10	349	562,066	9	667	966,298
neos-1605061	10	1,467	23,486	9	1,390	24,218
neos16	2	3,276	13,634,651	7	1989	7,524,606
ns1702808	10	3	9067	10	1	5962
ns1952667	10	161	10,053	10	109	7003
ns2081729	10	422	2,142,646	10	203	883,369
ns894244	9	2659	41,791	10	3078	53,032
ofi	10	468	9330	10	406	7063
p2m2p1m1p0n100	10	0	19,094	10	1	26,186
reblock166	3	2961	1,719,778	4	2538	867,706
wachplan	8	2,002	421,553	8	2018	381,747
	162	991	1,196,178	167	879	692,670

 Table 7
 Computational results for IBM-Cplex default and IBM-Cplex with local implied bound cuts

Selected instances from MIPLIB2010, executed using 4 threads on an Intel Xeon E3-1220V2 at 3.10 GHz

The results are encouraging: very aggressive local implied bound cuts do not deteriorate the results of IBM-Cplex default (which was not obvious for what discussed above concerning MIPLIB2010 instances) and there is an improvement on certain instances that does not look caused by random noise. Indeed, by concentrating on the 18 instances (roughly 5% of the total) that have shown a notable improvement and by repeating the experiments with 10 different random seeds the results are consistent, which means that IBM-Cplex is able to guess implications and on some cases strengthening the variable bounds at the nodes with locally valid cuts is extremely beneficial. This shows that our encouraging results are not restricted to SVM instances.

6.2 Iterative domain reduction

Iterative domain reduction can be seen as a preprocessing tool to enhance the behavior of a MIP solver. An initial bound tightening is performed by solving a sequence of MIPs to strengthen the lower and upper bounds on the ω variables and on b.

Let us denote by *P* the set of feasible solutions of the RLM, by $Z(\omega, \xi, z)$ the objective value (4) of solution (ω, ξ, z) , and by *U* the value of an upper bound on (4). To simplify notation, let ω_0 denote the *b* variable. Lower (l_i) and upper (u_i) bounds on each ω_i are iteratively tightened by solving the following MIPs:

$$l_i = \min\{\omega_i : (w, \xi, z) \in P, Z(w, \xi, z) \le U\}, \quad \forall i = 0, \dots, d,$$
(23)

$$u_i = \max\{\omega_i : (w, \xi, z) \in P, Z(w, \xi, z) \le U\}, \quad \forall i = 0, \dots, d.$$
(24)

where, at each step, solutions in P must satisfy all the ω -bounds computed in the previous iterations. In order to limit computing time, MIPs (23) and (24) are solved within a node limit. This iterative process is applied in a cyclic way until no bound improvement is obtained.

We have tested this approach on the 23 instances from Table 1. First, an initial upper bound U is computed by solving the RLM with a node limit of 100k (plus 25 polish nodes). Then, for each lower and upper bound tightening, the MIPs problems are solved within a node limit of 10k. When no bound improvement is obtained, the final RLM is solved with all new bounds on ω variables, unless each variable has lower and upper bounds equal each other. The results are reported in Table 8, where computing times and nodes refer to the overall scheme, i.e., they include the computation of the initial solution, the iterative bound tightening and possibly the final run. To our pleasant surprise, all instances were solved to optimality in less than 90 s, and required less that 1 min and 150k nodes on average.

7 Conclusions

We have shown that the nonconvex reformulation of so-called big-M constraints and the consequent use of a general-purpose MINLP solver instead of a MIP solver can lead, surprisingly, to faster computing times for a special class of classification problems. Through a careful analysis of Couenne's features and components we have

	Optimal value	IBM-Cpl	ex default			IBM-Cplex i.d.r.			
		Time (s)	Nodes	% ga	ар	Time (s)	Nodes	% gaj	р
				ub	lb			ub	lb
1	157,994.959	1510.54	15,928,227	_	_	42.50	145,231	_	_
2	179,368.534	×	28,467,361	_	17.20	56.62	148,105	_	_
3	220,673.592	×	23,375,230	_	37.63	49.06	149,472	_	_
4	5,225.994	2294.11	26,952,817	_	_	52.19	145,716	_	-
5	5,957.083	×	23,821,213	_	21.31	40.66	148,682	_	_
6	11,409,617.494	×	29,551,078	_	16.14	53.55	157,290	_	_
7	11,409,058.363	×	30,959,870	_	13.91	43.21	146,247	_	_
8	10,737,725.660	3104.27	34,784,855	_	_	78.81	146,013	_	_
9	5,705,364.054	×	27,976,050	_	17.37	49.49	147,621	_	_
10	5,704,804.923	х	30,681,973	-	16.10	51.25	145,978	-	_
11	5,369,016.540	3526.42	37,986,138	-	-	57.40	147,397	-	_
12	2,853,237.334	×	25,958,736	-	17.51	59.77	147,909	-	_
13	2,852,678.203	х	30,902,369	-	15.85	76.12	145,499	-	_
14	2,684,661.980	3103.51	35,402,901	-	-	48.12	146,772	-	-
15	1,427,173.974	×	29,375,413	-	13.93	50.83	148,504	-	-
16	1,426,614.843	×	30,139,333	-	14.63	41.02	145,243	-	-
17	1,342,484.700	3368.02	37,750,319	-	-	57.32	156,405	-	_
18	714,142.294	×	27,444,832	-	17.55	50.06	148,358	-	-
19	713,583.163	×	30,548,097	-	16.30	79.69	145,439	-	_
20	671,396.060	3249.70	36,008,315	-	-	61.33	146,101	-	-
21	357,626.454	×	25,638,748	-	17.73	50.28	161,660	-	_
22	357,067.323	×	30,175,822	-	16.68	45.28	145,992	-	-
23	335,851.740	3578.31	38,990,037	_	_	48.05	145,657	_	-

 Table 8
 Computational results for IBM-Cplex default and IBM-Cplex enhanced with iterative domain reduction (i.d.r.)

Instances of TypeB proposed by [8], n = 100, executed on an Intel i5-750 CPU running at 2.67 GHz with 4 threads

been able to isolate those that make a difference, namely aggressive bound tightening and iterative strengthening of the McCormick linearization. We have proposed two ways of integrating these ingredients within MIP approaches, both leading to finally being able to computationally solve these classification problems. One of these methods is currently part of the arsenal of IBM-Cplex 12.6.1.

More generally, we argue that aggressive bound tightening is often overlooked in MIP, while it represents a significant building block for enhancing MIP technology when indicator constraints and disjunctive terms are present. Finally, it is also conceivable that other ingredients that are fundamental in MINLP could prove beneficial for MIP.

Acknowledgments The research of Fischetti, Monaci and Salvagnin was supported by the University of Padova (Progetto di Ateneo "Exploiting randomness in Mixed Integer Linear Programming"). The research of Fischetti, Lodi, Monaci and Salvagnin was supported by MiUR, Italy (PRIN Project "Mixed-Integer

Nonlinear Optimization: Approaches and Applications"). The work of Nogales-Gómez was supported by an STSM Grant from COST Action TD1207. The authors are grateful to Andrea Tramontani and Sven Wiese for many interesting discussions on the subject and to two anonymous referees for a careful reading and many useful comments that helped improving the paper.

References

- 1. Andersen, E., Andersen, K.: Presolving in linear programming. Math. Program. 71, 221–245 (1995)
- 2. Balas, E.: Disjunctive programming. Ann. Discret. Math. 5, 3-51 (1979)
- Belotti, P., Lee, J., Liberti, L., Margot, F., Wächter, A.: Branching and bounds tightening techniques for non-convex MINLP. Optim. Methods Softw. 24(4–5), 597–634 (2009)
- Bonami, P., Biegler, L.T., Conn, A.R., Cornuéjols, G., Grossmann, I.E., Laird, C.D., Lee, J., Lodi, A., Margot, F., Sawaya, N., Wächter, A.: An algorithmic framework for convex mixed integer nonlinear programs. Discret. Optim. 5, 186–2004 (2008)
- Bonami, P., Kilinc, M., Linderoth, J.: Algorithms and software for convex mixed integer nonlinear programs. In: Lee, J., Leyffer, S. (eds.) Hot Topics in Mixed Integer Nonlinear Programming, IMA Volumes, pp. 1–40. Springer, Berlin (2012)
- Bonami, P., Lodi, A., Tramontani, A., Wiese, S.: On mathematical programming with indicator constraints. Math. Program. 151, 191–223 (2015)
- 7. Bonmin, v. 1.7.4. https://projects.coin-or.org/Bonmin
- 8. Brooks, J.P.: Support vector machines with the ramp loss and the hard margin loss. Oper. Res. **59**(2), 467–479 (2011)
- Carrizosa, E., Romero Morales, D.: Supervised classification and mathematical optimization. Comput. Oper. Res. 40, 150–165 (2013)
- 10. Cbc, v. 2.9. https://projects.coin-or.org/Cbc
- Ceria, S., Soares, J.: Convex programming for disjunctive convex optimization. Math. Program. 86, 595–614 (1999)
- Collobert, R., Sinz, F., Weston, J., Bottou, L.: Trading convexity for scalability. In: Proceedings of the 23rd International Conference on Machine Learning, pp. 201–208 (2006)
- 13. Couenne, v. branch/CouenneClassifier, r1046. https://projects.coin-or.org/Couenne
- 14. Cristianini, N., Shawe-Taylor, J.: An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods. Cambridge University Press, Cambridge (2000)
- 15. Davis, E.: Constraint propagation with interval labels. Artif. Intell. 32(3), 281-331 (1987)
- Duran, M.A., Grossmann, I.: An outer-approximation algorithm for a class of mixed-integer nonlinear programs. Math. Program. 36, 307–339 (1986)
- 17. FICO Xpress Optimization Suite, v. 7.8. http://www.fico.com/xpress
- 18. Fischetti, M., Monaci, M.: Exploiting erraticism in search. Oper. Res. 62, 114–122 (2014)
- Grossmann, I.E., Trespalacios, F.: Systematic modeling of discrete-continuous optimization models through generalized disjunctive programming. AIChE J. 59(9), 3276–3295 (2013)
- 20. Gurobi, v. 6.0.2. http://www.gurobi.com
- 21. IBM-Cplex, v. 12.6.1. http://www.ibm.com/software/products/en/ibmilogcpleoptistud
- 22. Ipopt, v. 3.9.2. http://projects.coin-or.org/Ipopt
- Koch, T., Achterberg, T., Andersen, E., Bastert, O., Berthold, T., Bixby, R.E., Danna, E., Gamrath, G., Gleixner, A.M., Heinz, S., Lodi, A., Mittelmann, H., Ralphs, T., Salvagnin, D., Steffy, D.E., Wolter, K.: Miplib 2010. Math. Program. Comput. 3, 103–163 (2011)
- Lodi, A., Tramontani, A.: Performance variability in mixed-integer programming. In: Topaloglu, H. (ed.) TutORials in Operations Research: Theory Driven by Influential Applications, pp. 1–12. INFORMS, Catonsville (2013)
- McCormick, G.P.: Computability of global solutions to factorable nonconvex programs: part I—convex underestimating problems. Math. Program. 10, 147–175 (1976)
- Messine, F.: Deterministic global optimization using interval constraint propagation techniques. RAIRO-RO 38(4), 277–294 (2004)
- Quesada, I., Grossmann, I.E.: An LP/NLP based branch-and-bound algorithm for convex MINLP optimization problems. Comput. Chem. Eng. 16, 937–947 (1992)
- Savelsbergh, M.W.P.: Preprocessing and probing techniques for mixed integer programming problems. ORSA J. Comput. 6, 445–454 (1994)

- 29. Shen, X., Tseng, G.C., Zhang, X., Wong, W.H.: On *ψ*-learning. J. Am. Stat. Assoc. 98, 724–734 (2003)
- Tawarmalani, M., Sahinidis, N.V.: Convexification and Global Optimization in Continuous and Mixed-Integer Nonlinear Programming: Theory, Algorithms, Software, and Applications. Kluwer Academic Publishers, Boston (2002)
- Wu, X., Kumar, V., Ross Quinlan, J., Ghosh, J., Yang, Q., Motoda, H., McLachlan, G.J., Ng, A., Liu, B., Yu, P.S., Zhou, Z.-H., Steinbach, M., Hand, D.J., Steinberg, D.: Top 10 algorithms in data mining. Knowl. Inf. Syst. 14, 1–37 (2007)