

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

A Machine Learning-Based Approximation of Strong Branching

Alejandro Marcos Alvarez, Quentin Louveaux, Louis Wehenkel
Université de Liège, Department of EE&CS,
Sart-Tilman B28, Liège, Belgium, {amarcos,q.louveaux,l.wehenkel}@ulg.ac.be

We present in this paper a new generic approach to variable branching in branch-and-bound for mixed-integer linear problems. Our approach consists in imitating the decisions taken by a good branching strategy, namely strong branching, with a fast approximation. This approximated function is created by a machine learning technique from a set of observed branching decisions taken by strong branching. The philosophy of the approach is similar to reliability branching. However, our approach can catch more complex aspects of observed previous branchings in order to take a branching decision. The experiments performed on randomly generated and MIPLIB problems show promising results.

Key words: branch-and-bound; variable branching; strong branching; supervised machine learning

1. Introduction

Most Mixed-Integer Programming (MIP) solvers are based on the branch-and-bound (B&B) algorithm (Land and Doig 1960). Over the years, numerous fundamental features, such as cutting planes, presolve, heuristics or advanced branching strategies, have been added to the solvers in order to improve their performances (Achterberg and Wunderling 2013). However, among those additional features, branching, i.e., the process that divides the feasible region into two or more subproblems, is probably the key component that most affects the efficiency of the solver (Achterberg and Wunderling 2013).

Branching strategies have been extensively studied in the literature, and we briefly review here some key contributions to that field. The simplest criterion, known as *most-infeasible branching*, consists in branching on the variable that has the greatest fractional part, i.e., the variable whose fractional part is closest to 0.5. However, most-infeasible branching is known to perform poorly in practice, and other methods, such as *pseudocost branching*

(Benichou et al. 1971), have been developed later. Pseudocost branching keeps a history of the dual bound increases observed during previous branchings, and uses this information to estimate the dual bound improvements for each candidate variable at the current node. Although pseudocost branching is very efficient in terms of computation time, the branchings performed at the very beginning of the B&B tree might be inefficient, as no reliable history has been recorded at that time. Later, Applegate et al. (1995) proposed a strategy, known as *strong branching*, that overcomes this limitation. Strong branching explicitly evaluates the dual bound increase for each fractional variable by actually computing the LP relaxations resulting from the branching on that variable. The variable that leads to the largest increases is chosen as branching variable for the current node. Despite its apparent simplicity, strong branching is, up to now, the most efficient branching strategy in terms of the number of nodes in the B&B tree. However, this efficiency is achieved at the expense of computation time, and strong branching is unfortunately intractable in practice. More recently, Achterberg et al. (2005) proposed to combine the advantages of both pseudocost and strong branching in a branching strategy called *reliability branching*. Many other branching strategies have been developed for the past 15 years, such as *inference branching* (Li and Anbulagan 1997), *non-chimerical branching* (Fischetti and Monaci 2012b), *active constraint branching* (Patel and Chinneck 2007), and *cloud branching* (Berthold and Salvagnin 2013), but their thorough description is beyond the scope of this paper. Finally, let us mention *hybrid branching* (Achterberg and Berthold 2009), which is probably today’s state-of-the-art branching strategy. Hybrid branching efficiently combines five scores obtained from other common branching strategies (including reliability branching as well as other strategies used in CSP and SAT solving), and is used as the main branching strategy in CPLEX 12.5 (Achterberg and Wunderling 2013). More specifically, hybrid branching uses the considered branching strategies to compute five different scores for each candidate variable. These scores are first normalized and then merged into a single value through a weighted sum. The variable that maximizes that sum is chosen as branching variable.

Following the ideas introduced by pseudocost branching, researchers have recently started investigating branching strategies that rely on information collected through multiple B&B restarts. *Backdoor branching* (Fischetti and Monaci 2012a) and *information-based branching* (Karzan et al. 2009) are two key contributions to this aspect. The mechanism

of these strategies is two-phased. During the first phase, the optimization of the current problem is restarted from the beginning multiple times (using several heuristics to guide the search), and the algorithm collects some information about each run (e.g., the conflict clauses or the most important branching variables). In the second phase, the real optimization starts, and the harvested information is used to take efficient branching decisions. The idea behind those methods is to quickly and superficially explore different parts of the B&B tree and to decide, based on those shallow explorations, which part it is best to focus on. The work presented in this paper relies on the very same basic idea. Finally, let us mention the work of Di Liberto et al. (2013) who recently proposed a method that uses machine learning techniques to switch between several branching heuristics. More specifically, they present an approach that first determines which strategy, among a restricted set of branching heuristics, produces the best results for a given set of problem instances. Then, when a new problem is solved by B&B, the algorithm checks, at several points in time, which strategy should be used at that particular moment, based on the optimal strategies that were found in the first phase. Although the main mechanism of their approach is different, the ideas are very similar, and their method has strong ties with the one developed in this paper. Indeed, their method also relies on the idea of leveraging prior information through principled learning techniques, but, while they use machine learning to switch between existing branching strategies, we use machine learning to provide a quick approximation to an existing branching strategy.

Our goal is to overcome the large computational overhead resulting from a strong branching decision. Speeding up strong branching-like decisions is not a new idea, as it is already behind other branching heuristics, such as reliability branching or non-chimerical branching. In this paper, we propose an alternative approach that uses machine learning techniques in order to imitate the strong branching decisions in an efficient way. More specifically, we propose a two-phased approach that yields a ‘learned’ branching strategy that can be used within B&B as an approximation of strong branching. The first phase involves solving a set of training problems with strong branching as branching heuristic in order to generate a set of branching decisions. During this phase, each branching decision is recorded in a dataset, called training set, that will then be used by a machine learning algorithm to learn a function imitating strong branching decisions. In the second phase, we introduce, as any other branching heuristic, the learned heuristic into B&B, and evaluate

its efficiency on a set of test problems. An important characteristic of our approach is that the first phase needs only to be done once. Indeed, once the branching heuristic is learned from the training set, it can be included directly into B&B, without requiring a new training phase each time a problem needs to be solved. In that sense, we can say that the training phase can be done in an off-line fashion, thus avoiding useless computational overhead at the beginning of each optimization.

In this paper, we address binary Mixed-Integer Linear Programming (MILP) problems of the form

$$\begin{aligned}
&\text{minimize} && c^\top x && (1) \\
&\text{s.t.} && Ax \leq b \\
&&& x_j \in \{0, 1\} \quad \forall j \in I \\
&&& x_j \in \mathbb{R}_+ \quad \forall j \in C,
\end{aligned}$$

where $c \in \mathbb{R}^n$, $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ respectively denote the cost coefficients, the coefficient matrix and the right-hand side. I and C are two sets containing the indices of the integer and continuous variables, respectively. We denote the solution at a given node of the B&B by x^* , and we will call, with a little abuse, the variable x_i , with $i \in I$, a fractional variable if it has a fractional value in the current solution x^* . The set of fractional variables of x^* is denoted F .

2. Preliminaries

In this section, we first present the concept of branching in a general functional form, and then briefly introduce the field of machine learning for the beginner.

2.1. Functional form of branching strategies

Any branching heuristic can be formulated in a generic functional form \mathcal{B} such that

$$\mathcal{B}: (i, \cdot) \mapsto \mathbb{R},$$

where i represents the index of the candidate branching variable, and \cdot represents unspecified arguments of \mathcal{B} . The branching variable i^* is chosen as the one that maximizes¹ the scores given by \mathcal{B} , i.e.,

$$i^* = \arg \max_{i \in F} \mathcal{B}(i, \cdot).$$

¹If ties occur, the variable that arrives first in the lexicographical order, i.e., the one with the smallest i , is chosen.

The functional form \mathcal{B} is different for every branching criterion, and proposing a new branching heuristic merely consists in providing a new \mathcal{B} , including its implementation and the specification of its arguments. For example, in the case of most-infeasible branching (MIB), \mathcal{B}_{mib} only requires the current fractional solution to output a score for a variable. Then, the functional form of MIB is written $\mathcal{B}_{\text{mib}}(i, x^*) = \min(1 - x_i^*, x_i^*)$. Another example is strong branching, which requires more input arguments, as it needs more information to take a decision. The functional form of strong branching (SB) is $\mathcal{B}_{\text{sb}}(i, c, A, b, x^*, l^*, u^*)$, where l^* and u^* represent the lower and upper bounds of the variables at the current node, respectively. The implementation of \mathcal{B}_{sb} consists in creating two subproblems by changing the upper and lower bounds of variable i in the current problem to $\lfloor x_i^* \rfloor$ and $\lceil x_i^* \rceil$, respectively. The LP-relaxations of the subproblems thus created are then solved, and, for each subproblem, the difference between the objective value of the subproblem and the current problem is computed. These differences represent the objective increases observed between the current node and the subproblems when tighter bounds are used for the variable i . The output of \mathcal{B}_{sb} is finally given by the product of the computed differences.

2.2. Learning functions from data

A very popular use of supervised machine learning (ML) concerns the automatic construction, or learning, of functions from data. Basically, ML is used in the context of some task \mathcal{T} , for which a function f mapping the current state of \mathcal{T} to a given output space is needed. Such a function maps inputs from a space Φ to an output space \mathcal{Y} , i.e., $f(\cdot) \in \mathcal{F} : \Phi \mapsto \mathcal{Y}$, where \mathcal{F} is the set of possible mappings from Φ to \mathcal{Y} . Formally, a general machine learning algorithm \mathcal{A} is a procedure of the form $\mathcal{A} : (\Phi \times \mathcal{Y})^N \mapsto \mathcal{F}$ that takes as input a dataset $D = ((\phi_i, y_i))_{i=1}^N \in (\Phi \times \mathcal{Y})^N$, and that outputs a function $f \in \mathcal{F}$ that minimizes some loss function \mathcal{L} on the dataset D , where N represents the number of elements in the dataset. Stated in mathematical terms, the function f^* resulting from the application of algorithm \mathcal{A} to dataset D is given by

$$f^* = \mathcal{A}(D) \in \arg \min_{f \in \mathcal{F}} \sum_{i=1}^N \mathcal{L}(y_i, f(\phi_i)).$$

Ideally, the input of function f should be the state of \mathcal{T} itself. However, representing the complete state is often a difficult problem, e.g., because its dimensionality is too high, or because it contains a lot of irrelevant information. For this reason, in the machine

learning community, the inputs ϕ of the functions f are usually ‘features’ representing a simplified version of the state of \mathcal{T} . Formally, features are characteristics extracted from the state of \mathcal{T} , such that those features represent part of the current state, ideally the part that most influences the output. The features, which are typically designed by human experts, often determine the efficiency of learning methods. As they represent only part of the current state of the task, it is important that the parts described by the features are indeed correlated with the desired output. Note that in this context, ‘correlated’ should not be understood in the traditional sense of the Pearson or Spearman correlation between individual features and the desired output. It rather means that it is crucial, for the learning algorithm to work, that there exists a strong enough statistical dependency, possibly taking very complex forms, between the set of chosen features and the desired output. For this reason, the features need to be carefully designed and tailored to the problem of interest.

3. Learning branching decisions

Since machine learning techniques can be used to approximate (learn) functions, and since branching can be determined by a function, it seems natural to consider machine learning as a reasonable way to build a branching strategy. We must stress here that machine learning does not provide a completely new branching criterion, as it requires the observation of real branching decisions to actually build a branching function. Rather than providing a new branching heuristic, machine learning constructs a branching strategy that imitates the decisions of the branching strategy that generated the dataset D , in our case, strong branching.

Our goal is to create an efficient approximation of strong branching that could be used in practice. In other words, the branching heuristic $\mathcal{B}_{\text{learned}}(i, \phi_i)$ that we propose is such that

$$\mathcal{B}_{\text{learned}}(i, \phi_i) \approx \mathcal{B}_{\text{sb}}(i, c, A, b, x^*, l^*, u^*), \quad (2)$$

where ϕ_i is a feature vector describing the state of the optimization problem at the current node from the perspective of variable i . The feature vector ϕ_i does not describe the current node of the B&B tree, but rather describes variable i in the current node. Those features need to be efficiently computable and have to well represent the problem at the current B&B node from the perspective of variable i . Section 4 explains in more detail how the features are designed.

3.1. Dataset generation

The first step towards the creation of $\mathcal{B}_{\text{learned}}$ with a machine learning technique is to generate a dataset from which the function can be learned.

In order to create a dataset of pairs (ϕ_i, y_i) , we optimize the problems contained in some sets, which we call training problems, with B&B, using strong branching as branching heuristic. At each node explored by B&B during the optimization of those problems, the strong branching score $\mathcal{B}_{\text{sb}}(i, c, A, b, x^*, l^*, u^*) = s_i$ is computed for each fractional variable $i \in F$ together with the features ϕ_i associated with that variable. These features-score pairs are then saved in the dataset D_{sb} , which is then used as input of the learning algorithm to generate a learned branching strategy $\mathcal{B}_{\text{learned}}(i, \phi_i)$.

3.2. Machine learning algorithm

Once the features are designed to correctly describe each variable of the problem, and once a dataset of input-output pairs is available, we can apply a machine learning algorithm to learn a function from the dataset. In this work, we use *Extremely Randomized Trees* (Geurts et al. 2006), or ExtraTrees. ExtraTrees are based on an ensemble of regression trees and are a slightly modified version of the better known random forests (Breiman 2001). Our choice is motivated by the simplicity and the robustness of ExtraTrees. Indeed, the performance of ExtraTrees is very robust against the choice of their parameters, and the default values provided in (Geurts et al. 2006) work very well in practice.

4. Features describing the current subproblem

As mentioned earlier, the features are the key component of our approach, since they critically condition the efficiency of the method. On the one hand, the features need to be complete and precise in order to describe the subproblem as accurately as possible. On the other hand, they need to be efficient to compute. It is important to keep this tradeoff in mind, because there are many good features that could have a very positive impact on the efficiency of the method, but that are too expensive to compute. An example of such features is the objective increase obtained when branching is performed on a variable, i.e., the numbers that are actually used by strong branching to take a decision. Such features cannot be used in our approach, because of the huge computational overhead required by their computation.

Before describing the features we used, we need to emphasize the three properties that these features should have. First, the number of features needs to be independent of the

size of the problem instance. Indeed, the learning algorithm that we selected can cope only with datasets in which all the feature vectors ϕ_i have the same number of elements. If this number depends on the size of the problem, a different branching strategy must be learned for each problem size. This is of course an impractical situation, and enforcing the size-independency is the best way to obtain a single learned branching strategy that can be used for any problem size. This might seem a straightforward requirement, but the size-independency is not trivial to achieve. Elementary features such as c , A or b cannot be used directly in that case. Another important requirement is that the features should be invariant with respect to irrelevant changes in the problem, such as row or column permutation. Finally, the developed features need to be independent of the scale of the problem, i.e., if the parameters (c , A , and b) are multiplied by some factor, the features should remain identical.

The features ϕ_i that we describe assume that the problem is in the canonical form (1). Each feature vector ϕ_i is computed for variable i at the current node, before being fed to $\mathcal{B}_{\text{learned}}$. The features are divided into three subsets representing different aspects of the optimization state, namely ‘static problem features’, ‘dynamic problem features’, and ‘dynamic optimization features’.

Note that, Hutter et al. (2014) recently introduced a certain number of features describing MIP problems that are related to the ones presented in this work. Some of our features are based on similar ideas as theirs. Specifically, features related to A_{ji}/b_j , $A_{ji}/\sum A_{ji}$, and the slack variables are used in both studies. There is one major difference however in the design and use of those features. In Hutter et al. (2014)’s work, the features are aggregated (either over the constraints or over the variables), because they are used to characterize a problem instance as a whole. In our work, on the other hand, the features are not aggregated over the variables, because they are meant to characterize a single variable within the problem. The fact that the design choices are different is not surprising. Indeed, since the use of ML in this paper and in Hutter et al. (2014)’s paper is in pursuit of different goals, the features that are relevant for each task are likely to be different.

4.1. Static problem features

The first set of features are computed from the sole parameters c , A and b . They are calculated once and for all and they represent the static state of the problem. Their goal

is to give an overall description of the problem. These features are designed such that the aforementioned requirements are fulfilled.

The first three of them are devoted to the description of the current variable in terms of the cost function. Besides the sign of the element c_i , we also use $|c_i| / \sum_{j:c_j \geq 0} |c_j|$ and $|c_i| / \sum_{j:c_j < 0} |c_j|$. Distinguishing both is important, because the sign of the coefficient in the cost function is of utmost importance for evaluating the impact of a variable on the objective value.

The second class of static features is meant to represent the influence of the coefficients of variable i in the coefficient matrix A . We develop three measures, namely M_j^1 , M_j^2 and M_j^3 , that describe variable i within the problem in terms of the constraint j . Once the values of the measure M_j are computed, the corresponding features added to the feature vector ϕ_i are given by $\min_j M_j$ and $\max_j M_j$. The rationale behind this choice is that, when it comes to describing the constraints of a given problem, only the extreme values are relevant.

The first measure M_j^1 is composed of two parts: M_j^{1+} computed by $A_{ji}/|b_j|$, $\forall j$ such that $b_j \geq 0$, and M_j^{1-} computed by $A_{ji}/|b_j|$, $\forall j$ such that $b_j < 0$. The minimum and maximum values of M_j^{1+} and M_j^{1-} are used as features, to indicate by how much a variable contributes to the constraint violations.

Measure M_j^2 models the relationship between the cost of a variable and the coefficients of the same variable in the constraints. Similarly to the first measure, M_j^2 is split in $M_j^{2+} = |c_i|/A_{ji}$, $\forall j$ with $c_i \geq 0$, and $M_j^{2-} = |c_i|/A_{ji}$, $\forall j$ with $c_i < 0$. As for the previous measure, the feature vector ϕ_i contains both the minimum and the maximum values of M_j^{2+} and M_j^{2-} .

Finally, the third measure M_j^3 represents the inter-variable relationships within the constraints. The measure is split into $M_j^{3+} = |A_{ji}| / \sum_{k:A_{jk} \geq 0} |A_{jk}|$ and $M_j^{3-} = |A_{ji}| / \sum_{k:A_{jk} < 0} |A_{jk}|$. M_j^{3+} is in turn divided in M_j^{3++} and M_j^{3+-} that are calculated using the formula of M_j^{3+} for $A_{ji} \geq 0$ and $A_{ji} < 0$, respectively. The same splitting is performed for M_j^{3-} . Again, the minimum and maximum of the four M_j^3 computed for all constraints are added to the features.

4.2. Dynamic problem features

The second type of features is related to the solution of the problem at the current B&B node. Those features contain the proportion of fixed variables at the current solution, the

up and down fractionalities of variable i , the up and down Driebeek penalties (Driebeek 1966) corresponding to variable i , normalized by the objective value at the current node, and the sensitivity range of the objective function coefficient of variable i , also normalized by $|c_i|$.

4.3. Dynamic optimization features

The last set of features is meant to represent the overall state of the optimization. These features summarize global information that is not available from the single current node. When branching is performed on a variable, the objective increases are stored for that variable. From these numbers, we extract statistics for each variable: the minimum, the maximum, the mean, the standard deviation, and the quartiles of the objective increases. These statistics are used as features to describe the variable for which they were computed. As those features should be independent of the scale of the problem, we divide each objective increase by the objective value at the current node, such that the computed statistics correspond to the relative objective increase for each variable. Finally, the last feature added to this subset is the number of times variable i has been chosen as branching variable, normalized by the total number of branchings performed.

5. Experiments

This section describes the experimental procedure that we set up in order to assess the efficiency of our approach. It is composed of three steps: (1) we generate a dataset D_{sb} using strong branching, (2) we learn from it a branching heuristic, and (3) we compare the learned heuristic with other branching strategies on various problems. This section describes the different datasets used within our approach as well as the experimental setup.

5.1. Problem sets

We use two types of problem sets, namely randomly generated problem sets and standard benchmark problems from the MIPLIB (Bixby et al. 1996, Achterberg et al. 2006). The random problem sets are used for both training (steps 1 and 2) and assessing (step 3) the heuristics, whereas MIPLIB problems are only used for assessment (step 3). The reasons for using random problems are two-fold.

First, the typical evaluation procedure in machine learning consists in evaluating a function learned from a dataset on a different dataset. If the function is both learned and evaluated on the same dataset, the estimated performance might be too optimistic and

might thus not reflect the overall performance of the learned function. To prevent this, the datasets are separated into two parts: the first part is used for training and the second part is used for assessing the method. Since the MIPLIB is the traditional evaluation benchmark for MIP methods, and in order to comply with the machine learning assessment methodology, we decided to use other problems, i.e., the randomly created problems, to learn our branching heuristic. Note that it would have been possible to use other techniques, such as k-fold cross validation, to correctly evaluate the performance of different branching strategies on the MIPLIB without requiring new problems to be created.

Secondly, the performance of any machine learning procedure increases with the size and the variety of the dataset used by the learning algorithm. In theory, the expected accuracy of the ML procedure over the entire input space of the learned function is a monotonically increasing function of the size of the dataset D . As this dataset is created by optimizing a set of problems, increasing the number of problems in the problem set yields a bigger dataset D . It is thus to our advantage to consider as many problems as possible to create D .

The problems used in this work are rather small. The small problem size used in our experiments is justified by the need to observe the entire B&B tree in the dataset D . Indeed, the learned branching heuristic can only reflect the branching decisions found in the dataset. If the training problems are too big or too difficult to solve, the dataset will only contain branching decisions observed at the beginning of the B&B tree, and will not reflect all the possible branching decisions. It is thus important that the dataset also contains branching decisions from the very bottom of the B&B tree. This consideration has to be taken into account when choosing the problems to include in our problem sets. For this reason, the size of the training problems has to be limited, so that the problems can be solved as much as possible with strong branching in a reasonable amount of time. Similarly, since the size of the training problems is limited, so is the size of the test problems.

5.1.1. Random problems. We randomly generate three sets of binary-integer or mixed binary-integer minimization problems that each contain two different types of constraints. The possible constraints are chosen among set covering (SC), multi-knapsack (MKN), bin packing (BP), and equality constraints (EQ). We generated problems that contain constraints of type BP-EQ, BP-SC, and MKN-SC. The number of variables, the number of constraints, and the values of the elements in the matrices c , A and b are randomly generated. More specifically, we arbitrarily choose some bounds on the number of variables,

Table 1 Randomly generated problem sets. ‘all’, ‘bin’ and ‘cont’ indicate the total number, and the number of binary and continuous variables in the problems, respectively.

	# prob.	Variables					
		all		bin		cont	
		min	max	min	max	min	max
BPEQ_train	25	201	225	150	179	43	54
BPEQ_test	50	193	234	145	185	43	54
BPSC_train	25	109	136	109	136	0	0
BPSC_test	50	109	137	109	137	0	0
MKNSC_train	25	188	358	188	358	0	0
MKNSC_test	50	185	342	185	342	0	0

Table 2 Randomly generated problem sets. ‘all’, ‘EQ’, ‘BP’, ‘SC’ and ‘MKN’ specify the total number, and the number of equality, bin packing, set covering, and multi-knapsack constraints in the problem sets, respectively.

	Constraints									
	all		EQ		BP		SC		MKN	
	min	max	min	max	min	max	min	max	min	max
BPEQ_train	94	138	39	50	55	89	0	0	0	0
BPEQ_test	94	135	39	50	53	89	0	0	0	0
BPSC_train	80	110	0	0	50	73	28	40	0	0
BPSC_test	80	112	0	0	50	75	27	39	0	0
MKNSC_train	108	156	0	0	0	0	61	77	42	84
MKNSC_test	108	160	0	0	0	0	58	77	43	89

Table 3 List of problems from MIPLIB3 and MIPLIB2003.

10teams	aflow30a	aflow40b	air03	air04	air05	cap6000	dcmulti
egout	fiber	fixnet6	harp2	khb05250	l152lav	lseu	mas74
mas76	misc03	misc06	misc07	mitre	mod008	mod010	mod011
modglob	nw04	opt1217	p0033	p0201	p0282	p0548	p2756
pk1	pp08a	pp08aCUTS	qiu	rentacar	rgn	set1ch	stein27
stein45	tr12-30	vpm1	vpm2				

on the number of constraints, and on the elements contained in the matrices c , A and b . Then, we generate the parameters defining the problem with a distribution that is uniform between the chosen bounds.

The number of variables in these problems is of the order of a couple of hundreds, and the number of constraints is of the order of one hundred. As some of those problems are going to be used to generate the training dataset, we randomly split each family into a ‘train’ and a ‘test’ set. In the end, we have six datasets ‘BPEQ_train’, ‘BPEQ_test’, ‘BPSC_train’, ‘BPSC_test’, ‘MKNSC_train’ and ‘MKNSC_test’. The test sets contain 50 problems each, while the train sets each contain 25. Tables 1 and 2 summarize statistics about the randomly generated problems. More specifically, those tables respectively contain the bounds on the number of variables and on the number of constraints of the problems belonging to each randomly generated set. Remember that, as explained in the previous

section, the generated problems are intentionally kept small. Those datasets are available online², or can be provided upon request.

5.1.2. MIPLIB. We also compare different branching strategies on a set composed of problems from MIPLIB3 (Bixby et al. 1996) and MIPLIB2003 (Achterberg et al. 2006). We eliminated the non-binary problems and the problems that were too big in order to match the size of the randomly generated problems. The set finally contains 44 problems. Table 3 lists the test problems taken from MIPLIB3.0 and MIPLIB2003.

5.2. Experimental setup

The computations have been performed on a 16-core computer, equipped with two Intel Xeon E5520 (2.27GHz, 8 cores, and 8MB cache) and 32GB RAM, running CentOS 5.4 and CPLEX 12.2. We ran our experiments on our selection of the MIPLIB problems and on the problem sets ‘BPEQ_test’, ‘BPSC_test’ and ‘MKNSC_test’, described in Section 5.1. The training dataset D_{sb} is generated from the problem sets ‘BPEQ_train’, ‘BPSC_train’ and ‘MKNSC_train’.

To assess only the performance of the different branching strategies, we disable heuristics, cuts, and presolve in CPLEX (except for the last experiment). For each optimization, only one core is made available, so that parallelism is disabled as well. We compare our approach to five other branching strategies, namely random branching (random), most-infeasible branching (MIB), non-chimerical branching (NCB) (Fischetti and Monaci 2012b), full strong branching (FSB), and reliability branching (RB) (Achterberg et al. 2005). Random branching is a branching strategy in which the variable is randomly chosen among the fractional variables. We use the perseverant version of non-chimerical branching (Fischetti and Monaci 2012b), and the default parameter values $\lambda = 4$ and $\eta = 8$ for reliability branching (Achterberg et al. 2005). Moreover, the strong branching LP-relaxations are solved to optimality, and there is no limit on the number of candidate fractional variables at each node.

The dataset D_{sb} contains around 7×10^7 training examples. This number is too large and we thus use only 10^5 of them, randomly selected without replacement from the original dataset D_{sb} . The chosen parameters of ExtraTrees are $N = 100$, $k = |\phi|$, and $n_{\min} = 20$. More details about the parameters of the ExtraTrees can be found in the supplemental

²<http://www.montefiore.ulg.ac.be/~ama/research.php>

material. In this setup and on the considered computers (without parallelization), training the branching heuristic takes around 6 minutes. Naturally, this time can be amortized if the same heuristic is used to optimize several problems and even reduced by slightly changing the parameters of the ExtraTrees (at the potential expense of performance). In this work, we use only a fraction of the dataset that we generate in order to train a model that approximates strong branching. This choice is motivated by the fact that the computational efficiency of ExtraTrees is heavily dependent on the size of the training set. The more samples there are in the set, the longer it takes to predict a value. This conflicts with the speed requirements of our method. There exist other learning techniques that do not suffer from this limitation and for which more data could be used without impacting the prediction speed. Note that including more data is beneficial to the prediction accuracy, but, in our case, the gain in the accuracy of the strong branching estimates is not sufficient to counterbalance the additional time taken to make a prediction.

Two types of experiments are performed: one where the optimization is stopped early based on a limit either on the number of explored nodes or on the time spent, and another one where the problems are solved until optimality however long it takes. For the first experiment, the rationale behind the two considered limits is to evaluate different aspects of the branching strategies. When the optimization is limited by the number of explored nodes, we can compare the branching strategies based on the closed gap³ and on the time spent to actually explore that given number of nodes. This sheds some light on how good a branching strategy is compared to other branching strategies. In these conditions, FSB is usually the best in terms of closed gap and the worst in terms of time spent. On the other hand, the time limit is useful to assess different strategies in practical conditions where the number of nodes matters less than the time required to solve a problem. The closed gap is also used in that experiment to assess how far from the optimum the optimization is after a given amount of time. In this case, FSB is typically outperformed, in terms of closed gap, by other strategies. The second experiment (without limit) is used to assess all branching strategies in a practical situation where the problems need to be solved to optimality.

³The closed gap ($\in [0; 1]$) is the ratio of the difference between the current dual bound and the objective value of the initial LP-relaxation, to the difference between the optimal objective value and the objective value of the initial LP-relaxation. A value close to 1 indicates that the optimization is almost finished.

5.3. Results

We now present a selection of results comparing our approach to other branching strategies (random, MIB, NCB, FSB, and RB). Tables 4, 5, and 7 report these results. In these tables, ‘Cl. Gap’ refers to the closed gap, and ‘S/T’ indicates the number of problems solved within the provided nodes or time limit, versus the total number of problems. ‘Nodes’ and ‘Time’ respectively represent the number of explored nodes and the time spent (in seconds) before the optimization either finds the optimal solution, or stops earlier because of one stopping criterion. Those values are measured separately for each problem and are averaged in the tables.

5.3.1. Experiments with limits. Table 4 first shows the results achieved on the random test problem sets for both stopping criteria. Those results show that our approach succeeds in efficiently imitating FSB. Indeed, the experiments performed with a limit on the number of nodes show that the closed gap is only 9% smaller, while the time spent is reduced by 85% compared to FSB. The experiments with a time limit show that the reduced time required to take a decision allows the learned strategy to explore more nodes, and to thus further close the gap than FSB. While these results are encouraging, they are still slightly worse than the results obtained with RB, which is both closer to FSB and faster than our approach.

In addition to the previously presented branching strategies, Table 4 contains one extra experiment for each family of random problems. The normal learned branching strategy, i.e., $(n_{\min} = 20, \text{all})$, is learned based on a dataset containing samples from the three types of training problems, i.e., BPSC, BPEQ, and MKNSC. We also investigate, for each type of random problem, the effect of training the branching strategy from a dataset generated with training problems of the same type as the target test problems. In other words, when we test this strategy on the BPSC test problems, the branching rule is learned based on samples generated with BPSC train problems only, which is indicated in the table by $(n_{\min} = 20, \text{BPSC only})$. Overall, the results show that the strategies learned only on the type of problem on which they are tested perform a bit better than the strategy learned from a dataset containing a mixture of the three types of problems. This indicates that the approach can benefit from training on a specific type of problem, and that the performance of the learned branching policy improves when the problems to optimize are aligned with the problems that are used to generate the training dataset.

Table 4 Optimization results for the random problems.

	Node limit (10^5 nodes)			Time limit (10 min.)		
	S/T	Cl. Gap	Time (s)	S/T	Cl. Gap	Nodes
BPSC_test problems						
Random	0/50	0.36	6.73	0/50	0.65	879,696
MIB	0/50	0.39	7.02	0/50	0.68	836,675
NCB	0/50	0.54	117.45	0/50	0.66	45,048
FSB	0/50	0.55	243.70	0/50	0.63	27,157
RB	0/50	0.52	28.29	1/50	0.77	223,369
Learned ($n_{\min} = 20$, all)	0/50	0.48	56.36	0/50	0.67	112,918
Learned ($n_{\min} = 20$, BPSC only)	0/50	0.51	60.54	0/50	0.70	109,066
BPEQ_test problems						
Random	0/50	0.33	17.44	0/50	0.55	366,982
MIB	0/50	0.40	17.27	0/50	0.61	368,309
NCB	0/50	0.81	290.49	0/50	0.86	22,605
FSB	0/50	0.83	681.75	0/50	0.82	9,492
RB	0/50	0.80	74.53	10/50	0.95	90,273
Learned ($n_{\min} = 20$, all)	0/50	0.75	77.97	5/50	0.92	106,057
Learned ($n_{\min} = 20$, BPEQ only)	0/50	0.77	85.68	4/50	0.92	86,370
MKNSC_test problems						
Random	0/50	0.56	7.26	24/50	0.95	587,123
MIB	0/50	0.60	7.39	31/50	0.97	496,475
NCB	0/50	0.67	102.23	5/50	0.83	51,749
FSB	0/50	0.68	135.41	5/50	0.83	46,832
RB	0/50	0.65	27.79	18/50	0.94	173,513
Learned ($n_{\min} = 20$, all)	0/50	0.64	28.37	18/50	0.93	177,006
Learned ($n_{\min} = 20$, MKNSC only)	0/50	0.64	34.93	16/50	0.92	165,412

Table 5 then shows the results obtained with a node limit and a time limit on the MIPLIB problems. In those experiments, we separated the problems that were solved by all methods from the problems that were not solved by at least one of the compared methods. Similarly to the results obtained on the random problem sets, the proposed branching strategy compares favorably with strong branching both on the node limit and time limit experiments. Nonetheless, the results obtained with the learned branching strategy are still a little below the results obtained with reliability branching. The results presented here are averaged over all considered problems. The detailed results for all problems in the MIPLIB set are available in the supplemental material.

5.3.2. Experiments without limits. Finally, Tables 6 and 7 report the results from our last set of experiments. We apply all branching heuristics on all MIPLIB problems initially contained in Table 3, and let the computers solve the problem for five days⁴. After this time

⁴The computers are given a time budget of five CPU days, after which the B&B process stops automatically. The only exception is the problem instance `pp08aCUTS`, which is solved in due time by all branching heuristics except for

Table 5 Optimization results for the MIPLIB problems of Table 3.

Node limit = 10^5 nodes	Solved by all methods			Not solved by at least one method			
	S/T	Nodes	Time (s)	S/T	Cl. Gap	Nodes	Time (s)
Random	9/44	1,974	2.24	0/44	0.43	10,000	124.50
MIB	9/44	2,532	6.03	6/44	0.50	9,274	233.19
NCB	9/44	879	10.70	11/44	0.72	7,322	232.74
FSB	9/44	692	14.48	12/44	0.73	7,184	629.87
RB	9/44	1,123	15.78	10/44	0.64	7,806	219.39
Learned ($n_{\min} = 20$, all)	9/44	1,194	2.73	10/44	0.62	8,073	162.87
Time limit = 10 min							
Random	19/44	29,588	30.50	0/44	0.47	867,837	600.01
MIB	19/44	14,931	14.68	3/44	0.52	764,439	561.27
NCB	19/44	7,051	41.55	5/44	0.73	101,408	513.00
FSB	19/44	5,687	70.84	3/44	0.66	49,008	534.65
RB	19/44	6,895	27.38	7/44	0.69	257,375	515.40
Learned ($n_{\min} = 20$, all)	19/44	14,008	34.12	5/44	0.63	130,081	512.72

limit, the problems that are not solved by all branching methods are discarded. We create thus a new set of MIPLIB problems (reported in Table 6) that is used to compare different branching strategies until the end of the optimization procedure. Additionally, in another experiment, we let CPLEX use cuts and heuristics (with default CPLEX parameters) in the course of the optimization in order to observe their impact on the efficiency of each branching strategy. The optimization results are shown in Table 7.

Overall, our method compares favorably to its competitors when cuts and heuristics are used by CPLEX. Indeed, in that case, our learned branching strategy is the fastest (almost three times faster than the second fastest method, i.e., MIB) to solve all the 30 considered problems. Note that the apparent bad results of RB are due to three problems that are especially hard for that branching heuristic (air04, air05, and mod011). If we remove them from the computation of the average optimization time, both RB and the learned branching strategy take 74 sec on average to solve the remaining 27 problems. That average time is still 40% smaller than the average time of the runner-up (MIB). These experiments show that our strategy behaves very well when cuts and heuristics are used by CPLEX to solve the problems. The detailed results are available in the supplemental material.

Things appear to be different when cuts and heuristics are not used. Indeed, based on the results of Table 7, our method seems to be very slow, but the large number of nodes

the random branching strategy. We therefore granted additional budget for that configuration (problem + branching heuristic) in order to maximize the number of problems considered in our experiments.

Table 6 Updated list of problems from MIPLIB3 and MIPLIB2003. This list contains the problems from Table 3 that are solved to optimality with each branching heuristic in less than five days.

aflow30a	air03	air04	air05	cap6000	dcmulti
egout	khb05250	l152lav	lseu	mas76	misc03
misc06	misc07	mitre	mod008	mod010	mod011
nw04	p0033	p0201	pk1	pp08aCUTS	qiu
rentacar	rgn	stein27	stein45	vpm1	vpm2

Table 7 Optimization results (until termination) for the updated list of the MIPLIB problems (Table 6).

	w/o cuts and heuristics		w/ cuts and heuristics	
	Nodes	Time (s)	Nodes	Time (s)
Random	7,809,341	29,377.10	152,564	503.38
MIB	3,472,431	7,387.09	105,692	356.52
NCB	145,244	1,136.34	34,500	1,451.74
FSB	129,047	1,597.12	25,941	895.36
RB	318,384	886.12	51,913	2,836.93
Learned ($n_{\min} = 20$, all)	1,037,055	3,023.34	57,652	124.94

and the large amount of time is actually due to a small number of problems for which the method does not work well. These problems artificially increase the average number of nodes and average amount of time reported in the table. A finer analysis can be conducted by interpreting the detailed results reported in the supplemental material. When done, we see that our method is faster than RB in 11/30 cases and faster than FSB in 21/30 cases, thus alleviating the a priori bad performance of the learned branching strategy. A possible explanation for why our method does not perform well on those problems can be that these problems, because too large, are not well represented in the dataset that we use in order to learn the branching strategy. This again shows the importance of considering a large and diverse dataset for the learning of an efficient branching strategy.

6. General observations about the proposed method

The goal of this section is to give some general observations about the method that we propose in this paper, and to try to draw a parallel with other existing branching strategies.

6.1. Comparison with other branching strategies

The idea behind our approach can be seen as a combination of the ideas behind reliability branching and information-based branching. On the one hand, our approach is similar to reliability branching in the sense that we want to find a fast proxy to strong branching. However, the approaches differ in the means by which the information is collected, and in how the information is used. On the other hand, our approach is similar to information-based branching in the sense that it tries to use general information collected during a

preliminary phase in order to speed up the current optimization. In both cases, the idea is to obtain a broad overview of the B&B process and to use this overview to take sensible branching decisions. The main difference lies in the way information is collected: information-based branching harvests information through multiple restarts on the same problem, while our method collects optimization information obtained through the optimization of a set of different optimization problems. Another important difference is that, in the case of information-based branching, the collection phase needs to be performed for each problem to optimize, while our approach only requires the information to be collected once. In the light of Table 4, it is to be noted, however, that customizing the branching strategy to the problem being optimized could supposedly yield better results, which would be in favor of the approach that involves collecting some data for each problem to solve.

6.2. Another branching strategy?

The underlying mechanism of our approach is not different from other branching strategies. Indeed, in all cases, features are computed from the current state of the problem (in some way or another) and then used to decide which variable to branch on. Those features can be, e.g., the fractionality of the variable in the current solution (used in most-infeasible and reliability branching), the pseudocosts of the variables (used in reliability branching), or the objective increases observed when the branching is performed (used in strong branching). In our approach, we include many types of features, including some of the aforementioned. The difference lies in how these features are used. While, in traditional methods, the features are assumed to explain everything and are all used to take the branching decision, we let the learning algorithm decide which features are relevant and which are not. The hope is that the machine will make a better choice than the humans when they decide to choose some features and to leave others aside. In that sense, we can see our method as a very general branching strategy that can imitate any other heuristic, as long as the appropriate features are provided.

For example, in its current implementation, our approach can be tied to pseudocost and reliability branching. Indeed, our features include the value of the fractionality of the variables and their pseudocosts (see Section 4). This means that all the information that pseudocost branching uses to take a decision is provided to the learning algorithm. Our approach can thus, in principle, use only the pseudocosts and the variable fractionalities to take a branching decision, if the learning algorithm decides that these features best explain

the desired output. But the learning algorithm can also produce a different branching criterion, using other features together with the pseudocosts, hopefully better approximating the desired output. Note that the desired output that we are focusing on is the strong branching score, which is different from the score used by pseudocost branching to take a branching decision.

6.3. Connections with the no free lunch theorems

The *no free lunch* theorems (NFL) (Wolpert and Macready 1997) state that, for certain types of mathematical problems, the performance of any optimization algorithm, averaged over all problems in the class, is equivalent. Additional NFL results indicate that matching, or aligning, algorithms to problems is a way to achieve better performance. This suggests that incorporating knowledge about the problem into the optimization algorithm has the potential to improve its efficiency. Although NFL theorems do not apply to MIP solving and branch-and-bound (Wolpert and Macready 1997), a similar behavior is often observed in practice. For example, while strong branching is very effective (in terms of the number of explored nodes) on general MIP problems, this strategy is not optimal for Constraint Satisfaction Problems (CSP), where other branching strategies are preferred. This indicates that, even if the NFL theorems do not apply as is, it is realistic to imagine that incorporating prior knowledge about the problem could help to improve the performance of traditional MIP solving approaches. The results published in prior work in similar contexts further support this conclusion (see, e.g., Hutter et al. 2010).

Surprisingly, in the MIP optimization area, hybrid branching (Achterberg and Berthold 2009) is the only strategy that takes this fact into consideration. Hybrid branching combines, in a weighted sum, several criteria known to be effective for different types of problems. This probably explains why hybrid branching outperforms other branching strategies across multiple problem sets (Achterberg and Berthold 2009, Achterberg and Wunderling 2013). Despite its simplicity, hybrid branching is a state-of-the-art strategy used in the last CPLEX release (Achterberg and Wunderling 2013). However, one might be interested in what would happen if the criteria were combined in a more elaborated way.

Although our work is completely devoted to the imitation of strong branching, we believe that the same framework can be applied to adapt the branching heuristic to the problem being optimized. In that aspect, the proposed method has a great potential to achieve better performance across large sets of problems, as the branching strategy generated using

machine learning can, in principle, imitate decisions made by different branching strategies at the same time.

7. Conclusion

In this paper, we proposed a new approach to design branching strategies for MILP problems. It consists in observing branching decisions taken by a supposedly good strategy, FSB in our case, and to imitate those decisions with a strategy obtained by machine learning. To this end, we develop a set of features that are used to characterize the current state of the problem in the B&B tree from the perspective of a particular variable. This set of features is then used as the input of the learned branching heuristic in order to predict the expected increase of the objective function that a branching on this variable would produce. The experiments show promising results and suggest that further research in this direction may lead to a practical approach to tune branching strategies to particular subclasses of MILP problems.

The underlying mechanism of our approach is not different from other popular branching heuristics. Indeed, in both cases, features are computed from the current state of the problem, and then used to decide which variable to branch on. In our approach, however, we may include many types of features, including those used by popular strategies. The approach is able to sort out which of those features are useful, and to automatically determine how to combine them to rank branching decisions. In that sense, we can see our method as a very general branching strategy that can imitate any other heuristic, as long as the appropriate features are provided. Our method can also discover novel heuristics by combining the features used by several popular methods with novel ones.

Beyond the imitation of existing strategies, another interesting research direction could be to design learning procedures that can distinguish, in a dataset of decisions obtained from different heuristics, the decisions that are productive, from those that are counter-productive.

Acknowledgments

This work was funded by the Dysco IUAP network of the Belgian Science Policy Office and the Pascal2 network of excellence of the EC. AMA's thesis is funded by a FRIA scholarship from the F.R.S.-FNRS. The scientific responsibility rests with the authors.

References

- Achterberg, T., T. Berthold. 2009. Hybrid branching. *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Springer, 309–311.
- Achterberg, T., T. Koch, A. Martin. 2005. Branching rules revisited. *Operations Research Letters* **33** 42–54.
- Achterberg, T., T. Koch, A. Martin. 2006. MIPLIB 2003. *Operations Research Letters* **34** 361–372.
- Achterberg, T., R. Wunderling. 2013. Mixed integer programming: analyzing 12 years of progress. *Facets of Combinatorial Optimization*. Springer, 449–481.
- Applegate, D., R.E. Bixby, V. Chvátal, W. Cook. 1995. Finding cuts in the tsp (a preliminary report). Tech. Rep. 05, DIMACS.
- Benichou, M., J.M. Gauthier, P. Girodet, G. Hentges, G. Ribiere, O. Vincent. 1971. Experiments in mixed-integer linear programming. *Mathematical Programming* **1** 76–94.
- Berthold, T., D. Salvagnin. 2013. Cloud branching. *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Springer, 28–43.
- Bixby, R.E., S. Ceria, C. McZeal, M.W.P. Savelsbergh. 1996. An updated mixed integer programming library: MIPLIB 3.0.
- Breiman, L. 2001. Random forests. *Machine learning* **45** 5–32.
- Di Liberto, G., S. Kadioglu, K. Leo, Y. Malitsky. 2013. Dash: Dynamic approach for switching heuristics. *arXiv preprint arXiv:1307.4689* .
- Driebeek, N.J. 1966. An algorithm for the solution of mixed integer programming problems. *Management Science* **12** 576–587.
- Fischetti, M., M. Monaci. 2012a. Backdoor branching. *INFORMS Journal on Computing* **25** 693–700.
- Fischetti, M., M. Monaci. 2012b. Branching on nonchimerical fractionalities. *Operations Research Letters* **40** 159–164.
- Geurts, P., D. Ernst, L. Wehenkel. 2006. Extremely randomized trees. *Machine learning* **63** 3–42.
- Hutter, F., H. H. Hoos, K. Leyton-Brown. 2010. Automated configuration of mixed integer programming solvers. *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Springer, 186–202.
- Hutter, F., L. Xu, H. H. Hoos, K. Leyton-Brown. 2014. Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence* **206** 79–111.
- Karzan, F.K., G.L. Nemhauser, M.W.P. Savelsbergh. 2009. Information-based branching schemes for binary linear mixed integer problems. *Mathematical Programming Computation* **1** 249–293.
- Land, A.H., A.G. Doig. 1960. An automatic method of solving discrete programming problems. *Econometrica: Journal of the Econometric Society* 497–520.

- Li, C.M., Anbulagan. 1997. Look-ahead versus look-back for satisfiability problems. *Principles and Practice of Constraint Programming-CP97*. Springer, 341–355.
- Patel, J., J.W. Chinneck. 2007. Active-constraint variable ordering for faster feasibility of mixed integer linear programs. *Mathematical Programming* **110** 445–474.
- Wolpert, D.H., W.G. Macready. 1997. No free lunch theorems for optimization. *Evolutionary Computation, IEEE Transactions on* **1** 67–82.

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

A Machine Learning-Based Approximation of Strong Branching Supplemental material

Alejandro Marcos Alvarez, Quentin Louveaux, Louis Wehenkel
Université de Liège, Department of EE&CS,
Sart-Tilman B28, Liège, Belgium, {amarcos,q.louveaux,l.wehenkel}@ulg.ac.be

These appendices constitute the supplemental material of the paper entitled *A Machine Learning-Based Approximation of Strong Branching* made available to the reader.

Appendix A: About ExtraTrees parameters

As mentioned in Section 3.2 of the main document, the performances of the ExtraTrees are very robust with respect to the choice of their parameters. The ExtraTrees actually have three parameters: N , which is the number of trees in the method, k , which is the number of features evaluated at each node during the creation of the trees, and n_{\min} , which is the number of learning samples contained in a node below which that node becomes a leaf. The number of trees is set to the default value of $N = 100$ in our experiments. The parameter k , which represents the number of features that are considered for the creation of the next node in the ExtraTrees, is also set to a default value of $k = |\phi|$. The exact understanding of these parameters is beyond the scope of this appendix, and we refer the reader to (Geurts et al. 2006) for a deeper explanation.

Tables 1, 2, 3 and 4 compare the influence of the parameter n_{\min} on the performance of the method. The experiments are the same as those shown in Section 5.3 of the main document. The main observation that can be made from those tables is that the parameter n_{\min} influences the computational time, but not the accuracy of the taken decisions. Indeed, the greater the n_{\min} , the faster the method. This behavior was expected, as a large n_{\min} produces smaller trees, generally yielding a reduced computational time required to take a branching decision. On the other hand, the third column of the tables 1 and 4, and the sixth column of the table 2, all of which correspond to the gap closed after the node limit has been reached, show that

Table 1 Results for the problems of BPEQ_test, BPSC_test and MKNSC_test.

	Node limit (10^5 nodes)			Time limit (10 min.)		
	S/T	Cl. Gap	Time (s)	S/T	Cl. Gap	Nodes
Learned - $n_{\min} = 1$	0/150	0.62	72.23	16/150	0.81	104,090
Learned - $n_{\min} = 5$	0/150	0.62	63.28	19/150	0.82	114,389
Learned - $n_{\min} = 10$	0/150	0.62	61.91	21/150	0.83	122,038
Learned - $n_{\min} = 20$	0/150	0.62	54.23	23/150	0.84	131,994

Table 2 Results for the MIPLIB problems. Node limit = 10^5 nodes.

	Solved by all methods			Not solved by at least one method			
	S/T	Nodes	Time (s)	S/T	Cl. Gap	Nodes	Time (s)
Learned - $n_{\min} = 1$	9/44	1,436	3.42	11/44	0.63	8,084	110.70
Learned - $n_{\min} = 5$	9/44	1,229	3.25	10/44	0.62	8,176	103.34
Learned - $n_{\min} = 10$	9/44	1,437	5.45	10/44	0.63	8,083	84.75
Learned - $n_{\min} = 20$	9/44	1,194	2.73	10/44	0.62	8,073	162.87

Table 3 Results for the MIPLIB problems. Time limit = 10 min.

	Solved by all methods			Not solved by at least one method			
	S/T	Nodes	Time (s)	S/T	Cl. Gap	Nodes	Time (s)
Learned - $n_{\min} = 1$	19/44	13,887	36.19	5/44	0.64	112,810	510.72
Learned - $n_{\min} = 5$	19/44	15,310	38.46	5/44	0.64	116,493	525.46
Learned - $n_{\min} = 10$	19/44	14,647	37.42	7/44	0.65	124,346	499.55
Learned - $n_{\min} = 20$	19/44	14,008	34.12	5/44	0.63	130,081	512.72

Table 4 Results for the problems from BPSC_test. Comparison between the strategy learned on the entire dataset and the strategy learned only from BPSC_train examples.

	Node limit (10^5 nodes)			Time limit (10 min.)		
	S/T	Cl. Gap	Time (s)	S/T	Cl. Gap	Nodes
Learned - $n_{\min} = 1$	0/50	0.47	76.20	0/50	0.63	83,938
Learned - $n_{\min} = 5$	0/50	0.47	67.31	0/50	0.64	94,826
Learned - $n_{\min} = 10$	0/50	0.47	68.58	0/50	0.65	102,941
Learned - $n_{\min} = 20$	0/50	0.48	56.36	0/50	0.67	112,918
Learned - $n_{\min} = 1$ - BPSC only	0/50	0.51	90.69	0/50	0.67	77,174
Learned - $n_{\min} = 5$ - BPSC only	0/50	0.51	79.53	0/50	0.67	82,535
Learned - $n_{\min} = 10$ - BPSC only	0/50	0.51	72.79	0/50	0.68	92,586
Learned - $n_{\min} = 20$ - BPSC only	0/50	0.51	60.54	0/50	0.70	109,066

the accuracy of the taken decision is not influenced by n_{\min} in the range of tested values. These observations illustrate that the method is actually robust to the choice of n_{\min} .

Besides the experiments included in this appendix, further work should focus on a more detailed study of the influence of the different parameters of the ExtraTrees on the performance of the optimization procedure.

Table 5 List of problems from MIPLIB3 and MIPLIB2003.

10teams	aflow30a	aflow40b	air03	air04	air05	cap6000	dcmulti
egout	fiber	fixnet6	harp2	khh05250	l152lav	lseu	mas74
mas76	misc03	misc06	misc07	mitre	mod008	mod010	mod011
modglob	nw04	opt1217	p0033	p0201	p0282	p0548	p2756
pk1	pp08a	pp08aCUTS	qiu	rentacar	rgn	set1ch	stein27
stein45	tr12-30	vpm1	vpm2				

Table 6 Updated list of problems from MIPLIB3 and MIPLIB2003. This list contains the problems from Table 5 that are solved to optimality with each branching heuristic in less than five days.

aflow30a	air03	air04	air05	cap6000	dcmulti
egout	khh05250	l152lav	lseu	mas76	misc03
misc06	misc07	mitre	mod008	mod010	mod011
nw04	p0033	p0201	pk1	pp08aCUTS	qiu
rentacar	rgn	stein27	stein45	vpm1	vpm2

Appendix B: Detailed experimental results

This appendix contains the detailed experimental results for the MIPLIB problems used in our experiments. Averaging the following results over all problems gives the aggregated results shown in the main paper. The detailed results are given in Tables 7 through 20.

The first set of tables, i.e., Tables 7 through 14, reports the results for the MIPLIB problems contained in Table 5 when limits are set either on the number of nodes or on the time spent. The B&B version that is used for these experiments is pure, i.e., no cuts nor heuristics are used.

The second set of tables, i.e., Tables 15 through 20, reports the optimization results for the problems contained in Table 6. These problems are obtained by keeping from the initial list the problems that are solved within a 5 days time limit with all considered branching strategies. Tables 15 through 20 then report the optimization results with no time (or node) limit on the second list of MIPLIB problems. Additionally, the second half of these tables contain the results when CPLEX's cuts and heuristics are used in the course of the optimization. The first half still corresponds to the pure version of B&B.

Table 7 Detailed results for the MIPLIB problems (1/4). Node limit = 10^5 nodes.

Problem names		10teams	aflow30a	aflow40b	air03	air04	air05	cap6000	dcmulti	egout	fiber	fixnet6
LP Obj.		917	983.17	1,005.66	338,864	55,535.40	25,877.60	-2,450,000	183,976	149.59	156,083	1,200.88
True Obj.		924	1,158	1,170	340,160	56,100	26,374	-2,450,000	188,182	568.10	405,935	3,983
Fin. by all		0	0	0	1	0	0	0	0	0	0	0
Method												
Random	Fin.	4	5	5	0	4	4	5	5	4	5	5
	Obj.	917	1,051.18	1,043.31	340,160	55,946.50	26,234.50	-2,450,000	187,298	562.60	189,746	1,573.11
	Nodes	10,000	10,000	10,000	25	10,000	10,000	10,000	10,000	10,000	10,000	10,000
	Time	364.40	9.35	44.47	1.45	2,219.56	1,173.48	30.70	5.61	1.12	5.39	4
	Cl. Gap	0	0.39	0.23	1	0.68	0.72	0.41	0.79	0.99	0.13	0.13
MIB	Fin.	4	5	5	0	4	5	5	0	0	5	5
	Obj.	920	1,059.97	1,038.70	340,160	55,980.80	26,270.30	-2,450,000	188,182	568.10	184,391	2,004.30
	Nodes	10,000	10,000	10,000	7	10,000	10,000	10,000	9,493	7,015	10,000	10,000
	Time	251.55	10.44	46.41	0.60	5,965.28	1,516.76	29.02	4.44	0.79	5.91	4.12
	Cl. Gap	0.43	0.44	0.20	1	0.74	0.79	0.41	1	1	0.11	0.29
NCB	Fin.	0	5	5	0	0	0	5	0	0	5	5
	Obj.	924	1,117.99	1,099.12	340,160	56,137	26,374	-2,450,000	188,182	568.10	359,328	2,528.38
	Nodes	1,453	10,000	10,000	3	131	215	10,000	1,065	5,713	10,000	10,000
	Time	456.18	98.16	354.31	2.04	1,146.13	772.90	147.50	2.99	1.94	52.82	74.06
	Cl. Gap	1	0.77	0.58	1	1	1	0.44	1	1	0.81	0.48
FSB	Fin.	0	5	5	0	0	0	5	0	0	5	5
	Obj.	924	1,120.14	1,099.87	340,160	56,137	26,400	-2,450,000	188,182	568	362,329	2,482.51
	Nodes	259	10,000	10,000	3	111	177	10,000	927	4,730	10,000	10,000
	Time	1,215.41	217.34	979.50	2.75	3,249.42	3,840	159	4.83	2.16	267.72	246.11
	Cl. Gap	1	0.78	0.58	1	1	1	0.44	1	1	0.83	0.46
RB	Fin.	5	5	5	0	0	0	5	0	0	5	5
	Obj.	917	1,103.69	1,089.99	340,160	56,137	26,374	-2,450,000	188,182	568.10	287,517	2,322.90
	Nodes	10,000	10,000	10,000	3	3,601	1,489	10,000	1,013	9,143	10,000	10,000
	Time	1,200.43	37.74	136.16	0.53	2,786.84	1,461.74	118	1.79	2.37	18.77	18.36
	Cl. Gap	0	0.69	0.52	1	1	1	0.44	1	1	0.53	0.40
Learned	Fin.	0	5	5	0	0	4	5	0	0	5	5
	Obj.	924	1,064.88	1,039.41	340,160	56,137	26,284.40	-2,450,000	188,182	568.10	313,565	2,239.68
	Nodes	1,739	10,000	10,000	3	3,833	10,000	10,000	1,565	3,007	10,000	10,000
	Time	108.16	57.10	146.12	0.48	1,648.75	2,237.31	45.80	2.76	1.31	42.17	61.90
	Cl. Gap	1	0.47	0.21	1	1	1	0.82	0.44	1	1	0.63

The row ‘Cl. Gap’ refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row ‘Fin. by all’ indicates whether all methods were able to solve this problem to optimality. For each method, the ‘Fin.’ rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 8 Detailed results for the MIPLIB problems (2/4). Node limit = 10^5 nodes.

Problem names		harp2	khb05250	1152lav	lseu	mas74	mas76	misc03	misc06	misc07	mitre	mod008
LP Obj.		-74,300,000	95,900,000	4,656.36	834.68	10,482.80	38,893.90	1,910	12,841.70	1,415	114,741	290.93
True Obj.		-73,900,000	107,000,000	4,722	1,120	11,801.20	40,005.10	3,360	12,850.90	2,810	115,155	307
Fin. by all		0	0	0	0	0	0	1	1	0	0	0
Method												
Random	Fin.	5	4	4	5	5	5	0	0	4	5	4
	Obj.	-74,300,000	107,000,000	4,700.21	1,005.77	10,860.40	39,280.80	3,360	12,851.10	2,294.29	115,094	305.69
	Nodes	10,000	10,000	10,000	10,000	10,000	10,000	1,179	527	10,000	10,000	10,000
	Time	16.92	6.17	32.30	1.10	3.37	3.04	0.51	0.69	10.44	182.61	1.50
	Cl. Gap	0.16	0.98	0.67	0.60	0.29	0.35	1	1.02	0.63	0.85	0.92
MIB	Fin.	5	0	4	5	5	5	0	0	4	0	0
	Obj.	-74,300,000	107,000,000	4,705.10	1,062.44	10,877	39,301.10	3,360	12,850.90	2,574.11	115,155	307
	Nodes	10,000	7,477	10,000	10,000	10,000	10,000	629	551	10,000	986	9,091
	Time	15.67	4.95	35.04	1.11	3.42	3.08	0.32	0.76	9	30.06	1.49
	Cl. Gap	0.14	1	0.74	0.80	0.30	0.37	1	1	0.83	1	1
NCB	Fin.	5	0	0	4	5	5	0	0	4	0	0
	Obj.	-74,100,000	107,000,000	4,722	1,093.29	11,041.20	39,371.80	3,360	12,850.90	2,487.50	115,155	307
	Nodes	10,000	1,527	213	10,000	10,000	10,000	579	56	10,000	617	4,605
	Time	95.59	4.76	4.53	5.16	28.23	23.72	2.65	0.58	152.62	49.18	2.85
	Cl. Gap	0.63	1	1	0.91	0.42	0.43	1	1	0.77	1	1
FSB	Fin.	5	0	0	5	5	5	0	0	4	0	0
	Obj.	-74,100,000	107,000,000	4,722	1,088.06	11,041.80	39,410.90	3,360	12,850.90	2,537.50	115,000	307
	Nodes	10,000	1,502	238	10,000	10,000	10,000	371	55	10,000	709	3,333
	Time	336	7.27	47.67	8.37	44.55	35.20	2.78	1.20	378.26	139	1.77
	Cl. Gap	0.65	1	1	0.89	0.42	0.47	1	1	0.80	1	1
RB	Fin.	5	0	0	4	5	5	0	0	4	0	0
	Obj.	-74,100,000	107,000,000	4,722	1,095.03	11,008	39,368	3,360	12,850.90	2,668.33	115,155	307
	Nodes	10,000	1,685	611	10,000	10,000	10,000	641	62	10,000	1,900	2,927
	Time	43.14	2.21	6.71	2.68	9.43	8.12	1.89	0.54	48.44	130.35	1.24
	Cl. Gap	0.56	1	1	0.91	0.40	0.43	1	1	0.90	1	1
Learned	Fin.	5	0	0	4	5	5	0	0	5	0	0
	Obj.	-74,100,000	107,000,000	4,722	1,083.23	10,914.30	39,321	3,360	12,850.90	2,410	115,155	307
	Nodes	10,000	5,790	1,061	10,000	10,000	10,000	1,865	161	10,000	470	6,365
	Time	59.84	8.29	8.77	6.90	21.87	16.38	2.29	0.53	35.01	26.31	3.81
	Cl. Gap	0.41	1	1	0.87	0.33	0.38	1	1	0.71	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 9 Detailed results for the MIPLIB problems (3/4). Node limit = 10^5 nodes.

Problem names		mod010	mod011	modglob	nw04	opt1217	p0033	p0201	p0282	p0548	p2756	pk1
LP Obj.		6,532.08	-62,100,000	20,400,000	16,310.70	-20.02	2,520.57	6,875	176,868	429.68	2,698.95	0
True Obj.		6,548	-54,600,000	20,700,000	16,862	-16	3,089	7,615	258,411	8,691	3,124	11
Fin. by all		0	0	0	1	0	1	1	0	0	0	0
Method												
Random	Fin.	4	5	5	0	5	0	0	5	5	5	5
	Obj.	6,543	-57,400,000	20,500,000	16,862	-20.02	3,089	7,615	184,159	1,422.30	2,713.04	2.38
	Nodes	10,000	10,000	10,000	259	10,000	4,859	649	10,000	10,000	10,000	10,000
	Time	21.46	67.20	3.31	10.80	6.83	0.31	0.45	2.36	4.10	12.33	2.72
	Cl. Gap	0.69	0.62	0.33	1	0	1	1	0.09	0.12	0.03	0.22
MIB	Fin.	0	5	5	0	5	0	0	5	5	5	5
	Obj.	6,548	-55,700,000	20,600,000	16,862	-20.02	3,089	7,615	181,564	512.38	2,702.63	3.05
	Nodes	532	10,000	10,000	1,737	10,000	6,265	4,747	10,000	10,000	10,000	10,000
	Time	3.74	77.08	3.34	45.13	6.82	0.35	2.37	2.56	3.16	9.92	2.66
	Cl. Gap	1	0.84	0.57	1	0	1	1	0.06	0.01	0.01	0.28
NCB	Fin.	0	5	5	0	5	0	0	0	5	5	5
	Obj.	6,548	-55,000,000	20,600,000	16,862	-19.80	3,089	7,615	258,411	8,678.53	2,925.78	4.79
	Nodes	96	10,000	10,000	355	10,000	745	173	624	10,000	10,000	10,000
	Time	2.47	2,764.36	24.84	62.49	17	0.10	1.16	1.12	24.01	142.83	33.08
	Cl. Gap	1	0.94	0.66	1	0.06	1	1	1	1	0.53	0.44
FSB	Fin.	0	5	5	0	5	0	0	0	0	5	5
	Obj.	6,548	-55,000,000	20,600,000	16,862	-19.77	3,089	7,615	258,411	8,690	2,942.55	4.54
	Nodes	31	10,000	10,000	233	10,000	363	186	502	8,920	10,000	10,000
	Time	3.33	3,815.40	63.53	82.35	150.78	0.07	3.71	1.29	24.80	483.98	53.50
	Cl. Gap	1	0.94	0.70	1	0.06	1	1	1	1	0.57	0.41
RB	Fin.	0	5	5	0	5	0	0	0	5	5	5
	Obj.	6,548	-55,300,000	20,600,000	16,862	-19.92	3,089	7,615	258,411	4,881.29	2,715.33	4.07
	Nodes	83	10,000	10,000	1,201	10,000	1,150	359	767	10,000	10,000	10,000
	Time	3.26	991.97	9.07	116.97	16.27	0.15	1.51	0.97	20.97	64.81	9.59
	Cl. Gap	1	0.90	0.58	1	0.03	1	1	1	0.54	0.04	0.37
Learned	Fin.	0	5	5	0	5	0	0	0	5	5	5
	Obj.	6,548	-55,900,000	20,600,000	16,862	-19.98	3,090	7,615	258,411	8,640.78	2,721.47	2.96
	Nodes	123	10,000	10,000	235	10,000	291	612	8,614	10,000	10,000	10,000
	Time	1.58	127.51	22.36	10.80	39.96	0.07	1.78	7.55	23.56	110.99	17.70
	Cl. Gap	1	0.83	0.45	1	0.01	1	1	1	0.99	0.05	0.27

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 10 Detailed results for the MIPLIB problems (4/4). Node limit = 10^5 nodes.

Problem names		pp08a	pp08aCUTS	qiu	rentacar	rgn	set1ch	stein27	stein45	tr12-30	vpm1	vpm2
LP Obj.		2,748.35	5,480.61	-931.64	28,800,000	48.80	32,007.70	13	22	14,210.40	15.42	9.89
True Obj.		7,350	7,350	-132.87	30,400,000	82.20	54,537.80	18	30	131,000	20	13.75
Fin. by all		0	0	0	1	1	0	1	0	0	0	0
Method												
Random	Fin.	5	5	5	0	0	5	0	5	5	5	5
	Obj.	4,270.24	6,265.34	-350.96	30,400,000	82.20	36,346.60	18	27.50	21,683.40	16.63	11.33
	Nodes	10,000	10,000	10,000	41	5,947	10,000	4,283	10,000	10,000	10,000	10,000
	Time	1.77	4.44	97.92	4.54	0.83	3.73	0.60	3.99	4.40	2.48	3.01
MIB	Cl. Gap	0.33	0.42	0.73	1	1	0.19	1	0.69	0.06	0.27	0.37
	Fin.	5	5	5	0	0	5	0	5	5	5	5
	Obj.	4,601.49	6,269.34	-344.50	30,400,000	82.20	35,370.60	18	27.33	23,785.10	16.97	11.32
	Nodes	10,000	10,000	10,000	26	4,147	10,000	4,681	10,000	10,000	10,000	10,000
NCB	Time	1.77	5.07	88.89	3.48	0.58	3.87	0.64	3.75	4.73	2.58	3.17
	Cl. Gap	0.40	0.42	0.74	1	1	0.15	1	0.67	0.08	0.34	0.37
	Fin.	5	5	5	0	0	5	0	5	5	5	5
	Obj.	5,078.04	6,729.63	-156.37	30,400,000	82.20	39,908.60	18	27.50	26,409.90	18.23	12.53
FSB	Nodes	10,000	10,000	10,000	26	2,735	10,000	3,240	10,000	10,000	10,000	10,000
	Time	18.01	58.62	1,403.40	24.52	1.03	37.06	1.77	47.48	50.81	10.54	36.48
	Cl. Gap	0.51	0.67	0.97	1	1	0.35	1	0.69	0.10	0.61	0.69
	Fin.	5	5	5	0	0	5	0	5	5	5	5
RB	Obj.	5,174.11	6,697.22	-216.31	30,400,000	82.20	40,156.80	18	28.07	27,009.20	18.05	12.50
	Nodes	10,000	10,000	10,000	26	2,849	10,000	2,141	10,000	10,000	10,000	10,000
	Time	73.86	199.13	3,404.77	32.19	1.74	592	3.51	181.32	1,724.58	28.28	66.18
	Cl. Gap	0.53	0.65	0.90	1	1	0.36	1	0.76	0.11	0.57	0.67
Learned	Fin.	5	5	4	0	0	5	0	5	5	5	5
	Obj.	4,659.35	6,567.33	-143.30	30,400,000	82.20	40,319	18	27.50	26,487	17.78	11.89
	Nodes	10,000	10,000	10,000	21	2,701	10,000	3,980	10,000	10,000	10,000	10,000
	Time	6.23	15.49	436.22	18.23	0.77	9.60	1.41	21.27	13.61	5.67	19
Learned	Cl. Gap	0.42	0.58	0.99	1	1	0.37	1	0.69	0.11	0.52	0.52
	Fin.	5	5	5	0	0	5	0	5	5	5	5
	Obj.	4,730.36	6,490.81	-279.63	30,400,000	82.20	39,600	18	27.50	24,065.90	16.95	11.35
	Nodes	10,000	10,000	10,000	36	3,401	10,000	4,140	10,000	10,000	10,000	10,000
Learned	Time	42.04	51.12	152.68	4.32	1.20	120	3.13	31.06	366.50	16.07	31.62
	Cl. Gap	0.43	0.54	0.82	1	1	0.34	1	0.69	0.08	0.33	0.38

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 11 Detailed results for the MIPLIB problems (1/4). Time limit = 600 seconds.

Problem names		10teams	aflow30a	aflow40b	air03	air04	air05	cap6000	dcmulti	egout	fiber	fixnet6
LP Obj.		917	983.17	1,005.66	338,864	55,535.40	25,877.60	-2,451,540	183,976	149.59	156,083	1,200.88
True Obj.		924	1,158	1,168	340,160	56,137	26,374	-2,451,200	188,182	568.10	405,935	3,983
Fin. by all		0	0	0	1	0	0	1	1	1	0	0
Method												
Random	Fin.	4	5	5	0	5	4	0	0	0	5	5
	Obj.	917	1,105.30	1,061.77	340,160	55,798.10	26,181.90	-2,451,380	188,182	568.10	217,681	1,837.57
	Nodes	20,337	603,128	141,714	25	762	3,742	81,127	87,985	11,855	1,076,478	1,304,269
	Time	600.01	600.01	600.01	1.45	600.02	600.02	263.52	46.17	1.31	600.01	600.01
	Cl. Gap	0	0.70	0.35	1	0.44	0.61	0.47	1	1	0.25	0.23
MIB	Fin.	4	5	5	0	5	5	0	0	0	5	5
	Obj.	920.27	1,114.89	1,057.98	340,160	55,849.90	26,221	-2,451,340	188,182	568.10	216,400	2,481.05
	Nodes	28,721	565,662	133,573	7	872	3,877	22,637	9,493	7,015	966,294	1,250,501
	Time	600.01	600.01	600.01	0.60	600.10	600.02	64.25	4.48	0.80	600.01	600.01
	Cl. Gap	0.47	0.75	0.32	1	0.52	0.69	0.59	1	1	0.24	0.46
NCB	Fin.	0	5	4	0	5	5	0	0	0	0	5
	Obj.	924	1,103.67	1,141.45	340,160	55,787.30	26,285.90	-2,451,340	188,182	568.10	405,935	2,766.69
	Nodes	1,453	16,178	52,727	3	12	104	17,214	1,065	5,713	88,143	69,931
	Time	456.95	600.01	600.01	2.04	600.06	600.01	261.12	3.04	1.99	275.17	600.01
	Cl. Gap	1	0.69	0.84	1	0.42	0.82	0.59	1	1	1	0.56
FSB	Fin.	5	4	5	0	5	5	0	0	0	4	5
	Obj.	923	1,135.33	1,095.06	340,160	55,632.60	26,139.10	-2,451,340	188,182	568.10	380,621	2,593.05
	Nodes	102	27,814	6,023	3	6	23	15,732	927	4,725	24,604	25,867
	Time	600.10	600.01	600.01	2.76	600.18	600.16	254.21	4.87	2.23	600.01	600.01
	Cl. Gap	0.86	0.87	0.55	1	0.16	0.53	0.59	1	1	0.90	0.50
RB	Fin.	5	4	5	0	5	5	0	0	0	4	5
	Obj.	917	1,149.05	1,102.58	340,160	55,823.60	26,227.20	-2,451,340	188,182	568.10	382,623	2,773.55
	Nodes	4,018	147,308	44,452	3	247	612	18,989	1,013	9,143	289,900	343,403
	Time	600.01	600.01	600.01	0.52	600.03	600.02	225.13	1.80	2.39	600.01	600.01
	Cl. Gap	0	0.95	0.60	1	0.48	0.70	0.59	1	1	0.91	0.57
Learned	Fin.	0	5	5	0	5	5	0	0	0	5	5
	Obj.	924	1,104.13	1,046.45	340,160	55,992.20	26,224.70	-2,451,380	188,182	568.10	372,086	2,411.64
	Nodes	1,739	120,872	43,281	3	553	2,083	66,587	1,565	3,007	150,084	104,029
	Time	108.28	600.01	600.01	0.49	600.01	600.01	314.71	2.74	1.31	600.01	600.01
	Cl. Gap	1	0.69	0.25	1	0.76	0.70	0.47	1	1	0.86	0.44

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 12 Detailed results for the MIPLIB problems (2/4). Time limit = 600 seconds.

Problem names		harp2	khb05250	l152lav	lseu	mas74	mas76	misc03	misc06	misc07	mitre	mod008
	LP Obj.	-74,325,200	95,919,500	4,656.36	834.68	10,482.80	38,893.90	1,910	12,841.70	1,415	114,741	290.93
	True Obj.	-73,899,300	106,940,000	4,722	1,120	11,801.20	40,005.10	3,360	12,850.90	2,810	115,155	307
	Fin. by all	0	1	1	1	0	0	1	1	1	0	1
Method												
Random	Fin.	5	0	0	0	5	4	0	0	0	5	0
	Obj.	-74,231,300	106,940,000	4,722	1,120	11,227.40	39,899.60	3,360	12,851.10	2,810	115,131	307
	Nodes	219,404	11,555	57,281	112,035	1,092,445	1,662,754	1,179	527	82,081	63,300	10,307
	Time	600.03	7.15	107.70	10.23	600.01	600.01	0.52	0.70	53.43	600.01	1.54
	Cl. Gap	0.22	1	1	1	0.56	0.91	1	1.02	1	0.94	1
MIB	Fin.	5	0	0	0	5	4	0	0	0	0	0
	Obj.	-74,236,200	106,940,000	4,722	1,120	11,245.90	39,888.20	3,360	12,850.90	2,810	115,155	307
	Nodes	260,188	7,477	42,037	52,413	1,046,064	1,230,909	629	551	24,005	986	9,091
	Time	600.03	4.96	96.91	6.10	600.01	600.01	0.33	0.77	15.88	30.04	1.51
	Cl. Gap	0.21	1	1	1	0.58	0.89	1	1	1	1	1
NCB	Fin.	5	0	0	0	5	4	0	0	0	0	0
	Obj.	-74,019,800	106,940,000	4,722	1,120	11,349.60	39,784.90	3,360	12,850.90	2,810	115,155	307
	Nodes	52,988	1,527	213	18,533	171,360	234,705	579	56	32,873	617	4,605
	Time	600.01	4.79	4.57	7.24	600.01	600.01	2.71	0.59	254.98	49.25	2.91
	Cl. Gap	0.72	1	1	1	0.66	0.80	1	1	1	1	1
FSB	Fin.	5	0	0	0	5	4	0	0	0	0	0
	Obj.	-74,040,700	106,940,000	4,722	1,120	11,310.80	39,786.60	3,360	12,850.90	2,810	115,155	307
	Nodes	17,259	1,502	238	24,957	132,769	178,936	371	55	25,551	709	3,333
	Time	600.02	7.40	48.17	16.65	600.01	600.01	2.83	1.22	591.65	139.18	1.81
	Cl. Gap	0.67	1	1	1	0.63	0.80	1	1	1	1	1
RB	Fin.	5	0	0	0	5	0	0	0	0	0	0
	Obj.	-74,034,800	106,940,000	4,722	1,120	11,481.90	40,005.10	3,360	12,850.90	2,810	115,155	307
	Nodes	115,774	1,685	611	16,933	547,986	621,181	641	62	19,179	1,900	2,927
	Time	600.03	2.24	6.76	3.99	600.01	496.77	1.92	0.53	59.87	130.55	1.25
	Cl. Gap	0.68	1	1	1	0.76	1	1	1	1	1	1
Learned	Fin.	5	0	0	0	5	4	0	0	0	0	0
	Obj.	-74,115,700	106,940,000	4,722	1,120	11,230.50	39,753	3,360	12,850.90	2,810	115,155	307
	Nodes	90,166	5,790	1,061	23,707	301,782	403,888	1,865	161	97,197	470	6,365
	Time	600.01	8.24	8.66	11.39	600.01	600.01	2.24	0.53	174.27	26.21	3.59
	Cl. Gap	0.49	1	1	1	0.57	0.77	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 13 Detailed results for the MIPLIB problems (3/4). Time limit = 600 seconds.

Problem names		mod010	mod011	modglob	nw04	opt1217	p0033	p0201	p0282	p0548	p2756	pk1
LP Obj.		6,532.08	-62,122,000	20,430,900	16,310.70	-20.02	2,520.57	6,875	176,868	429.68	2,698.95	0
True Obj.		6,548	-54,558,500	20,740,500	16,862	-16	3,089	7,615	258,411	8,691	3,124	11
Fin. by all		1	0	0	1	0	1	1	0	0	0	0
Method												
Random	Fin.	0	5	5	0	5	0	0	5	5	5	4
	Obj.	6,548	-55,956,000	20,607,600	16,862	-20.02	3,089	7,615	192,124	2,085.24	2,733.15	10.92
	Nodes	27,055	82,469	1,293,324	259	687,665	4,859	649	1,712,728	1,285,332	529,728	1,346,758
	Time	47.98	600.01	600.01	10.81	600.01	0.32	0.45	600.01	600.01	600.01	600.01
Cl. Gap		1	0.82	0.57	1	0	1	1	0.19	0.20	0.08	0.99
MIB	Fin.	0	0	5	0	5	0	0	5	5	5	0
	Obj.	6,548	-54,558,500	20,706,800	16,862	-20.02	3,089	7,615	184,470	524.95	2,705.50	11
	Nodes	532	49,533	989,483	1,737	670,264	6,265	4,747	1,504,144	1,121,079	634,588	927,663
	Time	3.76	419.22	600.01	44.77	600.01	0.37	2.43	600.01	600.01	600.01	382.10
Cl. Gap		1	1	0.89	1	0	1	1	0.09	0.01	0.02	1
NCB	Fin.	0	5	5	0	5	0	0	0	0	5	5
	Obj.	6,548	-56,157,900	20,682,800	16,862	-19.65	3,089	7,615	258,411	8,691	2,984.96	9.95
	Nodes	96	2,518	176,901	355	260,272	745	173	624	18,782	34,910	151,768
	Time	2.50	600.05	600.01	62.35	600.01	0.12	1.17	1.14	42.31	600.01	600.01
Cl. Gap		1	0.79	0.81	1	0.09	1	1	1	1	0.67	0.90
FSB	Fin.	0	5	5	0	5	0	0	0	0	5	5
	Obj.	6,548	-56,661,700	20,686,000	16,862	-19.66	3,089	7,615	258,411	8,691	2,950.75	9.23
	Nodes	31	1,260	101,090	233	38,913	363	186	502	8,920	12,364	125,603
	Time	3.36	600.03	600.01	82.23	600.01	0.07	3.78	1.31	25.15	600.01	600.01
Cl. Gap		1	0.72	0.82	1	0.09	1	1	1	1	0.59	0.84
RB	Fin.	0	5	5	0	5	0	0	0	0	5	0
	Obj.	6,548	-55,719,100	20,686,900	16,862	-19.76	3,089	7,615	258,411	8,691	2,763.77	11
	Nodes	83	6,180	493,119	1,201	389,380	1,145	359	767	145,693	92,008	366,151
	Time	3.27	600.06	600.01	117.90	600.01	0.15	1.52	0.97	174.10	600.01	358.78
Cl. Gap		1	0.85	0.83	1	0.07	1	1	1	1	0.15	1
Learned	Fin.	0	0	5	0	5	0	0	0	0	5	5
	Obj.	6,548	-54,558,500	20,600,400	16,862	-19.93	3,089	7,615	258,411	8,691	2,728.89	8.60
	Nodes	123	46,907	259,394	235	135,183	291	612	8,614	86,603	53,100	362,383
	Time	1.59	564.95	600.01	10.68	600.01	0.07	1.73	7.56	110.72	600.01	600.01
Cl. Gap		1	1	0.55	1	0.02	1	1	1	1	0.07	0.78

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 14 Detailed results for the MIPLIB problems (4/4). Time limit = 600 seconds.

Problem names		pp08a	pp08aCUTS	qiu	rentacar	rgn	set1ch	stein27	stein45	tr12-30	vpm1	vpm2
LP Obj.		2,748.35	5,480.61	-931.64	28,806,100	48.80	32,007.70	13	22	14,210.40	15.42	9.89
True Obj.		7,350	7,350	-132.87	30,356,800	82.20	54,537.80	18	30	130,596	20	13.75
Fin. by all		0	0	0	1	1	0	1	1	0	0	0
Method												
Random	Fin.	5	5	4	0	0	5	0	0	5	5	5
	Obj.	5,198.76	6,756.75	-226.26	30,356,800	82.20	37,522.60	18	30	24,059.40	18.25	12.60
	Nodes	1,725,290	1,103,681	74,246	41	5,947	1,351,176	4,283	63,115	1,251,955	1,608,231	1,455,001
	Time	600.01	600.01	600.01	4.54	0.85	600.01	0.60	20.25	600.01	600.01	600.05
	Cl. Gap	0.53	0.68	0.88	1	1	0.24	1	1	0.08	0.62	0.70
MIB	Fin.	5	5	4	0	0	5	0	0	5	5	5
	Obj.	5,625.42	6,817.68	-211.52	30,356,800	82.20	35,814.40	18	30	26,437.20	18.55	12.60
	Nodes	1,525,690	1,001,129	80,056	26	4,147	1,174,846	4,681	86,199	1,049,709	1,563,746	1,331,407
	Time	600.01	600.01	600.01	3.49	0.60	600.01	0.66	26.17	600.02	600.01	600
	Cl. Gap	0.63	0.72	0.90	1	1	0.17	1	1	0.11	0.68	0.70
NCB	Fin.	5	5	5	0	0	5	0	0	5	5	5
	Obj.	5,800.71	6,988.93	-229.11	30,356,800	82.20	41,291.10	18	30	28,983.30	19.25	13.16
	Nodes	223,945	86,240	4,592	26	2,735	122,404	3,239	44,219	120,874	501,139	142,007
	Time	600.01	600.01	600.01	24.52	1.06	600.01	1.80	149.92	600.01	600.01	600
	Cl. Gap	0.66	0.81	0.88	1	1	0.41	1	1	0.13	0.84	0.85
FSB	Fin.	5	5	5	0	0	5	0	0	5	5	5
	Obj.	5,666.07	6,864.18	-399.43	30,356,800	82.20	40,157.30	18	30	25,847.50	19	13.06
	Nodes	87,459	32,806	1,360	26	2,849	10,008	2,141	24,835	3,496	282,230	105,079
	Time	600.01	600.01	600.01	32.21	1.80	600.02	3.61	285.03	600.02	600.01	600.01
	Cl. Gap	0.63	0.74	0.67	1	1	0.36	1	1	0.10	0.78	0.82
RB	Fin.	5	5	0	0	0	5	0	0	5	0	5
	Obj.	5,649.88	7,082.74	-132.87	30,356,800	82.20	42,397.10	18	30	30,376.10	20	12.79
	Nodes	643,995	305,396	14,923	21	2,701	492,097	3,975	50,335	418,805	649,531	299,538
	Time	600.01	600.01	477.28	18.20	0.78	600.01	1.44	70.49	600.01	446.18	600.01
	Cl. Gap	0.63	0.86	1	1	1	0.46	1	1	0.14	1	0.75
Learned	Fin.	5	5	5	0	0	5	0	0	5	5	5
	Obj.	5,318.51	6,828.05	-179.74	30,356,800	82.20	40,650.20	18	30	24,425.10	18.18	12.24
	Nodes	166,483	129,426	47,914	36	3,401	52,779	4,141	50,013	16,200	448,168	219,927
	Time	600.01	600.01	600.01	4.32	1.20	600.01	3.12	97.31	600.03	600.01	600
	Cl. Gap	0.56	0.72	0.94	1	1	0.38	1	1	0.09	0.60	0.61

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 15 Detailed results for the updated list of MIPLIB problems (1/3). Time limit = none.

Problem names		aflow30a	air03	air04	air05	cap6000	dcmulti	egout	khb05250	l152lav	lseu
	LP Obj.	983.17	338,864	55,535.40	25,877.60	-2,451,540	183,976	149.59	95,919,500	4,656.36	834.68
	True Obj.	1,158	340,160	56,137	26,374	-2,451,200	188,182	568.10	106,940,000	4,722	1,120
	Fin. by all	1	1	1	1	1	1	1	1	1	1
Method											
Random	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,450,000	188,182	568.10	107,000,000	4,722	1,120
	Nodes	15,800,801	25	130,471	88,241	81,127	87,985	11,855	11,555	57,281	112,035
	Time	19,276.30	1.44	9,228.73	4,613	259.32	45.45	1.27	6.85	106.19	9.82
	Cl. Gap	1	1	1	1	4.53	1	1	1.01	1	1
MIB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,340	188,182	568.10	106,940,000	4,722	1,120
	Nodes	7,691,239	7	105,821	66,991	22,637	9,493	7,015	7,477	42,037	52,413
	Time	9,109.39	0.59	14,008.90	5,642.01	63.34	4.44	0.77	4.77	95.61	5.78
	Cl. Gap	1	1	1	1	0.59	1	1	1	1	1
NCB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,340	188,182	568.10	106,940,000	4,722	1,120
	Nodes	137,895	3	131	215	17,214	1,065	5,713	1,527	213	18,533
	Time	1,296.89	2.02	1,133.76	768.24	256.47	2.98	1.93	4.71	4.51	6.93
	Cl. Gap	1	1	1	1	0.59	1	1	1	1	1
FSB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,340	188,182	568.10	106,940,000	4,722	1,120
	Nodes	118,057	3	111	177	15,732	927	4,725	1,502	238	24,957
	Time	2,073.98	2.73	3,224.13	3,812.73	250.59	4.81	2.13	7.22	47.27	15.98
	Cl. Gap	1	1	1	1	0.59	1	1	1	1	1
RB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,340	188,182	568.10	106,940,000	4,722	1,120
	Nodes	230,038	3	3,601	1,489	18,989	1,013	9,143	1,685	611	16,933
	Time	898.36	0.52	2,760.38	1,449.85	222.89	1.77	2.31	2.17	6.66	3.83
	Cl. Gap	1	1	1	1	0.59	1	1	1	1	1
Learned	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,380	188,182	568.10	106,940,000	4,722	1,120
	Nodes	2,825,981	3	3,833	65,305	66,587	1,565	3,007	5,790	1,061	23,707
	Time	8,755.66	0.48	1,613.37	5,141.27	310.41	2.30	1.07	7.13	7.88	8.99
	Cl. Gap	1	1	1	1	0.47	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 16 Detailed results for the updated list of MIPLIB problems (2/3). Time limit = none.

Problem names		mas76	misc03	misc06	misc07	mitre	mod008	mod010	mod011	nw04	p0033
	LP Obj.	38,893.90	1,910	12,841.70	1,415	114,741	290.93	6,532.08	-62,122,000	16,310.70	2,520.57
	True Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Fin. by all	1	1	1	1	1	1	1	1	1	1
Method											
Random	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,851.10	2,810	115,155	307	6,548	-54,600,000	16,862	3,089
	Nodes	2,306,195	1,179	527	82,081	288,265	10,307	27,055	517,237	259	4,859
	Time	693.85	0.50	0.70	52.66	1,606.37	1.49	47.67	4,072.82	10.63	0.30
	Cl. Gap	1	1	1.02	1	1	1	1	0.99	1	1
MIB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	2,097,725	629	551	24,005	986	9,091	532	49,533	1,737	6,265
	Time	760.55	0.32	0.77	15.67	29.70	1.45	3.70	410.90	44.66	0.35
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
NCB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	783,139	579	56	32,873	617	4,605	96	19,617	355	745
	Time	1,491.35	2.65	0.58	250.44	48.30	2.82	2.47	5,422.18	61.96	0.10
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
FSB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	597,109	371	55	25,551	709	3,333	31	18,283	233	363
	Time	1,466.68	2.78	1.19	581.55	137.07	1.75	3.28	6,409.24	81.40	0.06
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
RB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	621,181	641	62	19,179	1,900	2,927	83	25,873	1,201	1,145
	Time	477.96	1.89	0.53	59.32	128.88	1.22	3.23	2,715.15	116.54	0.13
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
Learned	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	1,722,679	1,865	161	97,197	470	6,365	123	46,907	235	291
	Time	2,232.96	1.90	0.50	153.05	24.88	2.94	1.51	544.07	10.78	0.05
	Cl. Gap	1	1	1	1	1	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 17 Detailed results for the updated list of MIPLIB problems (3/3). Time limit = none.

Problem names		p0201	pk1	pp08aCUTS	qiu	rentacar	rgn	stein27	stein45	vpm1	vpm2
	LP Obj.	6,875	0	5,480.61	-931.64	28,806,100	48.80	13	22	15.42	9.89
	True Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Fin. by all	1	1	1	1	1	1	1	1	1	1
Method											
Random	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,400,000	82.20	18	30	20	13.75
	Nodes	649	1,400,487	172,918,239	411,879	41	5,947	4,283	63,115	14,838,614	25,017,647
	Time	0.44	585.92	793,141	2,523.43	4.48	0.81	0.58	19.91	17,681.30	27,319.80
	Cl. Gap	1	1	1	1	1.03	1	1	1	1	1
MIB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	4,747	927,663	63,144,303	337,813	26	4,147	4,681	86,199	8,953,981	20,513,193
	Time	2.36	365.08	161,744	2,058.33	3.44	0.56	0.63	25.79	7,629.32	19,579.40
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
NCB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	173	257,925	1,744,935	16,865	26	2,735	3,239	44,219	621,984	640,044
	Time	1.15	957.74	17,663.10	1,610.02	24.15	1.02	1.76	146.05	735.40	2,188.64
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
FSB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	186	306,805	1,561,719	49,693	26	2,849	2,141	24,835	465,770	644,925
	Time	3.69	1,041.17	17,645.50	7,136.80	31.84	1.73	3.49	276.55	954.07	2,692.21
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
RB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	359	366,151	2,668,423	14,923	21	2,701	3,975	50,335	649,531	4,837,431
	Time	1.49	344.92	4,815.89	476.17	17.96	0.75	1.39	69.08	431.22	11,571.10
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
Learned	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	612	1,545,753	7,272,423	116,593	36	3,401	4,141	50,013	6,606,414	10,639,147
	Time	1.43	2,187.65	22,281.40	1,074.32	4.26	1.01	2.42	76.64	10,218.50	36,031.50
	Cl. Gap	1	1	1	1	1	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 18 Detailed results for the updated list of MIPLIB problems (1/3). Time limit = none and CPLEX's cuts and heuristics applied.

Problem names		aflow30a	air03	air04	air05	cap6000	dcmulti	egout	khb05250	l152lav	lseu
	LP Obj.	983.17	338,864	55,535.40	25,877.60	-2,451,540	183,976	149.59	95,919,500	4,656.36	834.68
	True Obj.	1,158	340,160	56,137	26,374	-2,451,200	188,182	568.10	106,940,000	4,722	1,120
	Fin. by all	1	1	1	1	1	1	1	1	1	1
Method											
Random	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,450,000	188,188	568.10	107,000,000	4,722	1,120
	Nodes	58,369	0	50,847	29,149	180	641	3	5	45,433	185
	Time	256.88	0.59	3,956.43	1,839.59	4.06	1.58	0.01	0.41	324.38	0.13
	Cl. Gap	1	1	1	1	4.53	1	1	1.01	1	1
MIB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,350	188,188	568.10	106,940,000	4,722	1,120
	Nodes	35,205	0	46,139	38,793	60	125	3	3	29,545	97
	Time	203.06	0.60	4,403.65	2,607.30	1.30	0.71	0.01	0.42	227.48	0.11
	Cl. Gap	1	1	1	1	0.56	1	1	1	1	1
NCB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,350	188,182	568.10	106,940,000	4,722	1,120
	Nodes	115,987	0	365	487	20	8,135	7	9	187	223
	Time	2,043.61	0.59	3,488.17	1,844.01	0.87	33.28	0.02	0.46	4.74	0.22
	Cl. Gap	1	1	1	1	0.56	1	1	1	1	1
FSB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,180	188,186	568.10	106,940,000	4,722	1,120
	Nodes	31,323	0	231	211	49	40	7	5	169	133
	Time	1,158.81	0.60	10,229.90	3,941.68	1.53	0.77	0.02	0.42	18.73	0.27
	Cl. Gap	1	1	1	1	1.06	1	1	1	1	1
RB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,350	188,187	568.10	106,940,000	4,722	1,120
	Nodes	84,014	0	336,711	78,761	20	429	7	7	459	99
	Time	590.75	0.59	57,934.10	17,733.40	0.93	1.67	0.02	0.45	6.45	0.15
	Cl. Gap	1	1	1	1	0.56	1	1	1	1	1
Learned	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	1,158	340,160	56,137	26,374	-2,451,350	188,195	568.10	106,940,000	4,722	1,120
	Nodes	42,389	0	553	5,785	20	111	5	7	799	75
	Time	296.91	0.60	237.80	969.80	0.74	0.87	0.02	0.44	7.27	0.17
	Closed gap	1	1	1	1	0.56	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 19 Detailed results for the updated list of MIPLIB problems (2/3). Time limit = none and CPLEX's cuts and heuristics applied.

Problem names		mas76	misc03	misc06	misc07	mitre	mod008	mod010	mod011	nw04	p0033
LP Obj.		38,893.90	1,910	12,841.70	1,415	114,741	290.93	6,532.08	-62,122,000	16,310.70	2,520.57
True Obj.		40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
Fin. by all		1	1	1	1	1	1	1	1	1	1
Method											
Random	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,600,000	16,862	3,089
	Nodes	2,869,383	1,009	7	55,815	15	1,589	1,505	58,297	255	1
	Time	1,004.72	0.93	0.55	48.61	4.33	0.47	3.96	3,726.60	46.53	0.01
Cl. Gap		1	1	1	1	1	1	1	0.99	1	1
MIB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	1,742,607	279	11	13,707	29	1,537	39	4,027	337	1
	Time	518.99	0.47	0.56	13.28	4.27	0.47	0.65	333.87	51.31	0.01
Cl. Gap		1	1	1	1	1	1	1	1	1	1
NCB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	400,115	455	7	26,769	20	783	13	160,639	39	1
	Time	311.09	1.99	0.59	222.85	4.26	0.47	0.64	34,024.60	34.68	0.01
Cl. Gap		1	1	1	1	1	1	1	1	1	1
FSB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	355,507	301	7	17,299	20	423	11	19,789	55	1
	Time	394.68	3.60	0.61	410.78	4.44	0.52	0.91	6,564.28	47.72	0.01
Cl. Gap		1	1	1	1	1	1	1	1	1	1
RB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	561,667	305	9	19,231	15	1,755	115	64,870	1,213	1
	Time	188.54	1.02	0.61	61.48	4.30	0.78	1.24	7,436.06	143.11	0.01
Cl. Gap		1	1	1	1	1	1	1	1	1	1
Learned	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	40,005.10	3,360	12,850.90	2,810	115,155	307	6,548	-54,558,500	16,862	3,089
	Nodes	564,679	1,979	9	110,751	10	993	99	5,064	55	1
	Time	307.45	3.43	0.56	212.24	4.05	0.67	1.03	541.92	34.02	0.01
	Closed gap		1	1	1	1	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

Table 20 Detailed results for the updated list of MIPLIB problems (3/3). Time limit = none and CPLEX's cuts and heuristics applied.

Problem names		p0201	pk1	pp08aCUTS	qiu	rentacar	rgn	stein27	stein45	vpm1	vpm2
	LP Obj.	6,875	0	5,480.61	-931.64	28,806,100	48.80	13	22	15.42	9.89
	True Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Fin. by all	1	1	1	1	1	1	1	1	1	1
Method											
Random	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,400,000	82.20	18	30	20	13.75
	Nodes	40	861,991	4,965	462,433	9	2,605	4,499	61,825	0	5,869
	Time	0.73	265	8.08	3,565.69	12.75	0.60	0.78	22.17	0.02	4.95
	Cl. Gap	1	1	1	1	1.03	1	1	1	1	1
MIB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	57	892,187	5,157	265,859	7	2,233	4,633	82,377	0	5,721
	Time	0.83	272.01	7.98	1,999.22	12.93	0.51	0.80	27.61	0.01	5.15
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
NCB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	67	243,581	12,497	12,541	7	3,007	3,499	39,787	0	5,781
	Time	1.31	430.17	68.16	913.26	14.27	1.80	1.79	90.07	0.01	14.19
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
FSB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	67	284,851	7,523	27,779	7	1,489	2,153	24,949	0	3,847
	Time	2.60	651.97	70.38	3,074.14	14.23	0.82	3.46	244.19	0.02	18.59
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
RB	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	113	308,491	9,515	26,277	11	2,271	4,231	53,623	0	3,197
	Time	1.13	153.21	22.80	745.02	14.98	0.83	1.58	58.58	0.01	4.18
	Cl. Gap	1	1	1	1	1	1	1	1	1	1
Learned	Fin.	0	0	0	0	0	0	0	0	0	0
	Obj.	7,615	11	7,350	-132.87	30,356,800	82.20	18	30	20	13.75
	Nodes	43	897,129	3,131	25,065	7	1,619	4,449	54,541	0	10,207
	Time	0.88	651.99	8.18	363.04	13.03	0.65	2.78	69.59	0.01	18.03
	Closed gap	1	1	1	1	1	1	1	1	1	1

The row 'Cl. Gap' refers to the gap closed at the end of the optimization whether it achieved optimality or not. The row 'Fin. by all' indicates whether all methods were able to solve this problem to optimality. For each method, the 'Fin.' rows indicate the termination status of the optimization: 0 for optimality, 1 for unfeasibility, 2 for unboundedness, 3 for unfeasibility or unboundedness, 4 for another stopping criterion with a feasible solution found, and 5 for another stopping criterion with no feasible solution found.

References

Geurts, P., D. Ernst, L. Wehenkel. 2006. Extremely randomized trees. *Machine learning* **63** 3–42.