

Learning for Spatial Branching: An Algorithm Selection Approach

Bissan Ghaddar

Corresponding author. Ivey Business School, Western University, London, Ontario, Canada, bghaddar@ivey.ca

Ignacio Gómez-Casares, Julio González-Díaz, Brais González-Rodríguez, Beatriz Pateiro-López
CITMAga (Galician Center for Mathematical Research and Technology). Department of Statistics, Mathematical Analysis
and Optimization and MODESTYA Research Group, University of Santiago de Compostela.

Sofía Rodríguez-Ballesteros

CITMAga (Galician Center for Mathematical Research and Technology)

The use of machine learning techniques to improve the performance of branch-and-bound optimization algorithms is a very active area in the context of mixed integer linear problems, but little has been done for non-linear optimization. To bridge this gap, we develop a learning framework for spatial branching and show its efficacy in the context of the Reformulation-Linearization Technique for polynomial optimization problems. The proposed learning is performed offline, based on instance-specific features and with no computational overhead when solving new instances. Novel graph-based features are introduced, which turn out to play an important role for the learning. Experiments on different benchmark instances from the literature show that the learning-based branching rule significantly outperforms the standard rules.

Key words: Spatial branching, non-linear optimization, polynomial optimization, machine learning, statistical learning.

1. Introduction

In recent years, there has been a substantial rise on the use of learning techniques to improve the performance of branch-and-bound schemes. Most of the focus so far has been set on mixed integer linear programming (MILP) problems and, in particular, on learning to choose the branching variable (see [Lodi and Zarpellon \(2017\)](#) and [Bengio et al. \(2021\)](#) and references therein). Since branch-and-bound schemes are also at the core of most global optimization algorithms for solving non-linear and non-convex optimization problems, it is natural to wonder to what extent the insights gained for (discrete) MILP problems can be successfully transferred to the (continuous) non-linear setting. To this end, in this paper we study variable selection for spatial branching in non-linear programming (NLP) problems.

We develop our analysis for polynomial optimization (PO) problems, a subclass of NLP problems known to be \mathcal{NP} -hard in the general case. PO has a wide range of practical

applications in the context of control, energy and water networks, process engineering, facility location, economics and equilibrium, and finance. PO generalizes several special cases that have been thoroughly studied in optimization, including mixed binary linear optimization, convex/non-convex quadratic optimization, and complementarity problems.

A key factor of the efficiency of state-of-the-art global solvers in non-linear optimization is the design of the branch-and-bound algorithm and, in particular, the branching rules. Inspired by the recent literature on “learning to branch”, we explore ways to harness learning-based approaches to improve the performance of the branch-and-bound search for PO and, more specifically, we focus on variable selection in spatial branching. Among the various approaches that have been proposed for solving problems in PO, here we build upon the branch-and-bound scheme embedded in the Reformulation-Linearization Technique (RLT) introduced by [Sherali and Tuncbilek \(1992\)](#). Although we have framed our contribution in the context of the RLT technique for PO, the insights are extensible to any branch-and-bound technique designed to find global optima for any class of NLP problems.

2. Contribution to the Literature

Machine learning (ML) and the neighbouring field of statistical learning have both been effectively leveraged in several areas to develop algorithms that learn from observing the performance of different tasks. While mathematical optimization intrinsically lies at the core of learning methods, recent years have seen a rise of research in the opposite direction: learning and predicting the best optimization approaches in different settings. In particular, a very active area of research concerns the use of ML techniques in the algorithmic design of optimization solvers, independently of the application domain in which such algorithms are used. A substantial part of the associated literature, to which this paper belongs, focuses on improving the performance of branch-and-bound schemes.

Because of the wide ranging applications and the ease of use of state-of-the-art solvers, mixed integer linear programming represents an excellent testing ground for the aforementioned research. Further, since these problems are solved using branch-and-bound techniques, the richness of the data generated in the process naturally lends itself to learning. As already mentioned, variable selection in branch-and-bound schemes is probably the topic that has attracted the most attention so far and its outstanding importance for the performance of such schemes has long been recognized.

The integration of learning in the algorithmic design of the branching decisions can involve different paradigms and approaches, and occur on various levels. Strong branching has proven to be the variable selection rule that leads to smallest trees (Linderoth and Savelsbergh 1999), but the associated computational costs make it’s overall performance not competitive with more sophisticated rules such as reliability branching (Achterberg et al. 2005). This motivates the research in Khalil et al. (2016) and Alvarez et al. (2017), where the authors devise learning approaches to approximate strong branching without the associated computational overhead, and test their approaches on the MIPLIB instances. Another natural learning approach related to branching is taken in Liberto et al. (2016) where the authors, inspired by the portfolio approaches to algorithm selection (Gomes and Selman 2001, Leyton-Brown et al. 2003), learn a clustering-based classifier to switch between branching rules and again test their approach on MIPLIB instances. Because of the richness of the different learning paradigms, novel approaches in this area are constantly appearing such as the use of neural networks in Gasse et al. (2019) and Cappart et al. (2021). They approximate strong branching by encoding the branching policies into a graph convolutional neural network which allows exploiting the natural bipartite graph representation of MILP problems, thereby reducing the amount of manual feature engineering. Then, they test their approach on four benchmarks of specific combinatorial optimization problems.

Although some work has been done on using learning for NLP, this area is still not as mature as the one for MILP. For instance, for non-convex quadratic programming problems with box constraints, Baltean-Lugojan et al. (2019) try to learn effective linear outer-approximations of semidefinite constraints. To generate the corresponding cuts, a neural network is used to select the most promising submatrices without the computational burden of solving semidefinite programming problems. For convex quadratic problems, Bonami et al. (2018) use ML to decide between applying branch and bound on the quadratic problem and applying it on an equivalent MILP reformulation.

Despite the extensive work done on learning to branch in MILP problems, it is an issue not yet well understood and a very active area of research. The understanding of variable selection is even more limited in spatial branching for NLPs, where there has not been much research in the topic (an exception is Belotti et al. (2009)), much less in the learning context.

Our research aims to bridge this gap by proposing a learning framework for spatial branching. We start by noticing that even relatively similar branching rules exhibit a lot of variability in their relative performance to one another across instances of different test sets. This is

not something specific to this setting since, as argued in [Liberto et al. \(2016\)](#) and references therein, work in portfolio algorithms has shown that there is often no single approach that performs best on all instances. Thus, it is important to know to what extent one can learn which rule will perform best on a given instance and this is the main objective of this paper. We follow an algorithm selection approach in the spirit of the one used in [Liberto et al. \(2016\)](#) for MILP problems, but our learning takes place in a regression framework instead of a clustering one. In the present work, learning is done offline and the branching rule for a new instance is selected based on its underlying features. Given the promising results we obtain, the proposed approach can be seen as a first step towards more sophisticated (possibly online) learning schemes for variable selection in NLP problems. Given that the learning takes place offline, our approach entails no computational overhead once the solving process starts, so there is no node-efficiency *vs.* time-efficiency trade-off. We develop our analysis for PO problems by building upon the global solver RAPOSa ([González-Rodríguez et al. 2020](#)). We perform our numerical experiments using a comprehensive set of instances that builds upon well established benchmarks including randomly generated instances from [Dalkiran and Sherali \(2016\)](#), MINLPLib instances ([Bussieck et al. 2003](#)), and QPLIB instances ([Furini et al. 2018](#)), which illustrate that our approach can be successfully applied to general PO problems without focusing on a particular class of problems or structure.

The contribution of this paper is threefold: (1) developing a framework for learning to select the most efficient branching rule in non-linear optimization problems, (2) defining graph-based features and branching rules that significantly contribute to the outcome of the learning process, and (3) introducing a new performance measure which, by combining information related to both running time and optimality gap, allows for a smooth comparison in sets of instances of varying structure and difficulty.

The remainder of this paper is organized as follows. In [Section 3](#), we present the core elements on which our learning framework stands: i) the RLT-based spatial branch-and-bound scheme for PO and ii) the set of branching rules for algorithm selection. In [Section 4](#), we describe the learning methodology, motivate and define the novel key performance indicator (KPI), and describe the features used for learning. In [Section 5](#), we show the performance of the learning approach. Finally, conclusions and future research directions are drawn in [Section 6](#).

3. Motivation and Problem Description

As we have discussed, the main goal of this paper is to provide a first study of the impact that learning may have in the context of spatial branching. In this section we describe the polynomial optimization setting in which we develop our analysis, some graphs and graph-related features for polynomial optimization problems, and the branching rules that constitute our portfolio for algorithm selection.

3.1. Polynomial Optimization and RLT Technique

The Reformulation-Linearization Technique was originally developed in [Sherali and Tuncbilek \(1992\)](#). It was designed to find global optima in polynomial optimization problems of the following form:

$$\begin{aligned}
 & \text{minimize} && p_0(x) \\
 & \text{subject to} && p_r(x) \geq \alpha_r; \quad r = 1; 2; \dots; R_1 \\
 & && p_r(x) = \alpha_r; \quad r = R_1 + 1; \dots; R \\
 & && x \in \mathcal{R}^n,
 \end{aligned} \tag{PO}$$

where $\mathcal{R} = \{x_1, \dots, x_n\}$ denotes the set of variables, each $p_r(x)$ is a polynomial of degree $r \leq N$ and $\alpha_r = \{x \in \mathcal{R}^n : 0 \leq x_j \leq u_j < 1; \forall j \in \mathcal{N}\} \subset \mathcal{R}^n$ is a hyperrectangle containing the feasible region. Then, $R = \max\{r \in \{0, \dots, R_0\} : p_r\}$ is the degree of the problem and $\mathcal{M}(N; \mathcal{R})$ represents all possible monomials of degree

RLT is based on the construction of a linear relaxation of the polynomial problem, which is then solved using a branch-and-bound scheme. This is accomplished by replacing each monomial of the problem with a corresponding RLT variable. For example, associated to a monomial of the form $x_1 x_2 x_4$ one would define the RLT variable X_{124} . More generally, RLT variables are defined as

$$X_J = \prod_{j \in J} x_j; \tag{1}$$

where J is a multiset containing the information about the multiplicity of each variable in the underlying monomial. Then, at each node of the branch-and-bound tree, one would solve the corresponding linear relaxation. In order to get a tighter relaxation and ensure convergence, new constraints, called bound-factor constraints, are also added. At each node, branching is performed on one of the variables that present some violation of the RLT-defining identities (1). Therefore, one of the crucial steps in any implementation of the RLT is the branching variable selection, which we discuss in Section 3.3.

3.2. Graphs Associated to Polynomial Optimization Problems

The relevance of graph representations of general optimization problems has long been recognized and, as discussed in [Bienstock and Muñoz \(2018\)](#) and [Faenza et al. \(2022\)](#), polynomial optimization problems are not an exception. Our interest on these graphs is twofold: first, as a means to define new branching rules and, second, to define features that can be taken as input by the learning process.

To each problem of the form PO , one can associate different graphs that capture some characteristics of its structure. We introduce two different graphs, both based on the concept of intersection graphs as defined in [Fulkerson and Gross \(1965\)](#). The first one is the variables intersection graph VIG , analogous to the one used in [Bienstock and Muñoz \(2018\)](#). It is built by assigning a vertex to each variable and connecting two vertices if the corresponding variables appear together in some monomial of PO . The second one is the constraints-monomials intersection graph CMIG , where we have one vertex for each monomial, one vertex for each constraint, and one for the objective function. We then connect one vertex associated to a monomial with one associated to a constraint (or the objective function) if the monomial appears in the constraint (or objective function).

Once the graphs are defined, we are interested in parameters that summarize some properties of their underlying structure. Below we describe several standard parameters that we use in different parts of our analysis (for further details refer, for instance, to [Latora et al. \(2017\)](#)).

- ^ Density or edge density . Given a graph G , the edge density of G is defined as the number of edges in G divided by the number of edges in the complete graph with the same number of nodes.

- ^ Treewidth . Given a graph G , the treewidth is the smallest width of any tree-decomposition of G ([Robertson and Seymour 1984](#)). It gives a measure of how close a graph is to a tree. Trees have treewidth 1.

- ^ Modularity . Given a division of a graph G into communities, the modularity of the division is a measure of the strength of that division (i.e., graphs with high modularity have dense connections between the nodes within communities but sparse connections between nodes in different communities). The modularity of G is the maximum modularity over all the possible divisions into communities of the graph. Since the exact computation of this parameter is a hard computational problem, we use the greedy algorithm in [Clauset \(2004\)](#) to compute an approximate value.

^ Transitivity . The transitivity of a graph G is the ratio between the number of triangles (subgraphs that are complete graphs of order 3) and the number of triads \mathcal{T} .

^ Eigenvector centrality or eigencentality . For each node in a graph, its centrality score is the corresponding value in the first eigenvector of the graph adjacency matrix.

We now move to the definition of different branching rules, two of which heavily rely on the eigencentralities of the nodes associated with the variables of the problem.

3.3. Branching Rules

State-of-the-art solvers for non-linear optimization build upon two main themes to define their variable selection rules: i) adaptations of MILP methods such as strong branching and, most notably, reliability branching (Achterberg et al. 2005), as described in Belotti et al. (2009) and ii) contributions of the variables to the violations of non-linear terms using variants of the so-called violation transfer (Tawarmalani and Sahinidis 2004, Belotti et al. 2009). In the specific context of RLT, past literature has mainly focused on violations of the RLT-defining identities. In the seminal paper by Sherali and Tuncbilek (1992), starting from an optimal solution $(X; x)$ at a node in the branch-and-bound tree, the violation of each variable in PO is computed according to the following expression:

$$j = \max_{t \in \{1, \dots, n\}} \max_{\{g \in \mathcal{J}^2(N; t) : |J| = t\}} f_j X_{J|f|g} - x_j X_{J|g}. \quad (2)$$

Then, the variable that maximizes this value is chosen for branching. More recent contributions such as Sherali et al. (2012a,b) and González-Rodríguez et al. (2020) rely on variations of (2) where the maxima are replaced with sums. In particular, González-Rodríguez et al. (2020) define various branching rules based on the formula:

$$j = \sum_{t \in \{1, \dots, n\}} \sum_{\{g \in \mathcal{J}^2(N; t) : |J| = t\}} w(j; J) j X_{J|f|g} - x_j X_{J|g} \quad (3)$$

where $w(j; J)$ are the weights associated with the violations. In order to define our portfolio of branching rules we complement the original formulation in (2) with various rules defined by using different weights inside the sum, which can be seen in Table¹.

Figure 1 shows the results of a preliminary computational experiment to assess the performance of the different branching rules on PO problems. More precisely, we used three

¹ The reason behind choosing only one rule based on maxima is that previous analysis have shown that maxima-based rules are normally outperformed by the sum-based counterparts (González-Rodríguez et al. 2020).

Table 1 Definition of the branching rules.

Branching rule	Definition
Maximum (max)	Eq. (2).
Sum (sum)	Eq. (3), with $w(j; J) = 1$ for all j and J .
Dual (dual)	Eq. (3), with $w(j; J)$ defined as the sum of the absolute values of the shadow prices of the problem constraints containing x_j .
Range (range)	Eq. (3), with $w(j; J) = \frac{\min\{u_j - x_j, x_j - l_j\}}{u_j - l_j} g$, where u_j and l_j are the upper and lower bound of x_j at the root node, i.e. PO, and u_j and l_j are the bounds at the current node.
Eigencentality VIG (eig-VI)	Eq. (3), with $w(j; J)$ defined as the eigencentality of x_j 's node in VIG.
Eigencentality CMIG (eig-CMI)	Eq. (3), with $w(j; J)$ defined as the eigencentality of x_j 's node in CMIG (0 if x_j never appears alone in a monomial).

different sets of instances (refer to Section 5.1 for details): DS (Dalkiran and Sherali 2016), MINLPLib (Bussieck et al. 2003), and QPLIB (Furini et al. 2018). For our comparison of the branching rules we use a key performance indicator (KPI) whose values are normalized to $[0; 1]$ (details are given in Section 4.2). The closer the normalized KPI is to 1 for a given branching rule, the better the performance of that rule. Figure 1a shows, for the problems in each of the three sets of instances, the boxplots of this normalized KPI using each of the described branching rules. As expected, no branching rule outperforms the rest in the different data sets. While the sum rule performs reasonably well in the DS instances, it does not for the ones in MINLPLib. In QPLIB, the eig-CMI rule clearly outperforms all other rules, but it would be a poor choice for many of the MINLP instances. These observations already hint at the promising potential of the integration of learning techniques in the branching rule selection process. In Figure 1b we rank, for each problem, the branching rules from 1 (best) to 6 (worst), according to the normalized KPI. As we observe from the learning process, it is important not only to choose the best possible branching rule, but also to ensure that, in case of mistake, the selected rule performs reasonably well. For instance, for those MINLPLib instances where the max rule is not the best one, its performance is on average significantly worse than other methods. In those cases it would be preferable to choose, for instance, the range rule, even when it is not the best possible choice.

(a) Boxplots of the normalized KPI for each of the branching rules.

(b) For each instance branching rules were ranked from 1 (best) to 6 (worst), according to the KPI. Stacked bar graphs represent percentage of problems in each rank position per branching rule. Numbers inside bars indicate the average value of the normalized KPI for the problems in each rank position.

Figure 1 Comparative analysis of the performance of the different branching rules.

4. Proposed Methodology

4.1. Overview of the Learning Framework

Our learning framework is related to the approach for switching heuristics in [Liberto et al. \(2016\)](#), where the authors use a cluster-based method to learn to choose from a portfolio of branching rules in MILP problems. In this paper, we develop a different, regression-based, learning technique. The main elements of our setting, some of which are novel with respect to past literature, are the following:

- ^ Spatial branching . We separate from previous literature by studying branching variable selection in non-linear problems.
 - ^ Algorithm portfolio . The goal is learning to choose from the branching rules defined in Section 3.3.
 - ^ Data collection . Learning is performed on a rich set of very diverse instances (see Section 5.1). All instances are solved with all branching rules, gathering all the relevant data on their performance. Thus, all the learning takes place online, with no computational overhead when solving new instances (beyond the computation of the features of the given instance).
 - ^ Graph-based branching rules and features . A novel part of our approach is the use of graphs to define both branching rules and features (see Tables 1 and 2, respectively). Remarkably, both of them turn out to play an important role in the results.
 - ^ Learning technique . It consists of running individual regressions for each branching rule, with the goal of predicting its relative performance with respect to the remaining ones. Then, given a new instance, the rule with the best predicted relative performance is chosen.
 - ^ KPI. We present a novel KPI, which we refer to as pace, and that is discussed in detail in Section 4.2. The pace allows to jointly study the performance in instances solved by all methods (where running time is normally used) and instances not solved by at least one method (where optimality gap is normally used).
 - ^ Learning on static features. Learning uses instance-specific features, with no dynamic information being gathered on a node by node basis along the branch-and-bound tree. Thus, the resulting learned branching rule is static, in the sense that it selects the same branching rule in each and every node of the tree associated to a given instance. This is different from most of the literature in MILP problems where dynamic features are often included (see, for instance, [Liberto et al. \(2016\)](#), [Khalil et al. \(2016\)](#), and [Alvarez et al. \(2017\)](#)).
- It is worth noting that, despite the variability in performance illustrated in Figure 1, all the rules in our portfolio revolve around the same baseline computation (violations of RLT-defining identities). This common core probably limits to what extent they can complement each other and, as a consequence, may result in learned rules with less potential to significantly outperform the original ones. Similarly, the choice to perform learning only with static features, may result in learning a less effective branching rule. Because of the

two aforementioned points, our analysis provides some sort of lower bound on the potential for learning in spatial branching. One of our main contributions is to show that, even under these limitations, there is much potential to improve the overall performance of the branch-and-bound algorithm by conducting online learning on the branching rule.

4.2. A Novel Key Performance Indicator

Carrying out a fair comparison of the performance of different algorithms on instances of varying difficulty often involves some challenges. Ideally, one would like to have a KPI that allows to compare the performance on all instances together, but the most widely used KPIs, running time and optimality gap, fail to do so. To illustrate this, let's define the following subsets of the set of instances:

- ^ A = Instances solved by all algorithms within the time limit,
- ^ B = Instances solved by some but not all algorithms within the time limit, and
- ^ C = Instances not solved by any algorithms within the time limit.

The running time is a good KPI in A and uninformative in C, whereas the optimality gap is a good KPI in C and uninformative in A. Further, neither the time nor the gap allows to differentiate between algorithms for all instances in B, since the running time does not distinguish between algorithms that have reached the time limit (but with different gaps) and the optimality gap does not distinguish between algorithms that have solved the given instance (but with different running time). Because of this, it is customary to split the performance analysis separating the problems in B and C or A and C [B. For instance, in [Alvarez et al. \(2017\)](#) they separated the problems that were solved by all methods from the problems that were not solved by at least one of the compared methods. An important limitation of such approaches is that the sets A, B and C depend on the algorithms to be compared, so the results may substantially change if a new algorithm is added to the study or an existing one is removed. In [Gupta et al. \(2020\)](#), when discussing limitations of their work stemming from the NP-hard nature of MILP solving, the authors acknowledge that One can consider the primal-dual bound gap after a time limit as an evaluation metric for the bigger instances, but this is misaligned with the solving time objective.

We propose a novel approach to define KPIs that allows to jointly compare all instances while being equivalent to the running time in A, mimicking the optimality gap in C, and being able to completely differentiate between algorithms in B. Thus, this new family of KPIs overcomes all the issues mentioned above. First, recall the standard definition of the

optimality gap at the end of the algorithm, OG^{end} , as a function of the lower and upper bounds (LB and UB, respectively):

$$OG^{\text{end}} = \frac{UB^{\text{end}} - LB^{\text{end}}}{UB^{\text{end}} + \epsilon};$$

where ϵ is a small constant needed to avoid divisions by zero. We can now define a KPI based on the pace at which the optimality gap is closed during the execution:

$$\text{pace}^{\text{OG}} = \frac{\text{time}}{OG^{\text{end}} - OG^{\text{root}} + \epsilon};$$

Note that, in our setting, OG^{root} is common for all branching rules. So define pace^{OG} represents the time needed to improve the optimality gap in one unit. Suppose now that we are comparing two algorithms according to their $\text{pace}_1^{\text{OG}}$ and $\text{pace}_2^{\text{OG}}$:

- ^ In set A, OG^{end} is the same for all rules. Hence $\text{pace}_1^{\text{OG}} = \text{pace}_2^{\text{OG}} = \text{time}_1 = \text{time}_2$.
- ^ In set C, the running time is the same for all rules. Hence $\text{pace}_1^{\text{OG}} = \text{pace}_2^{\text{OG}}$ is of the form $(OG_2^{\text{end}} + K) / (OG_1^{\text{end}} + K)$, where K is a constant term.
- ^ In set B, pace^{OG} may recognize the difference between two rules that have solved an instance if they have different running times and, similarly, between two rules that have not solved an instance if they have different gaps. Further, as desired, rules that have solved an instance will have a better pace than those that have not.

Given the above properties pace^{OG} stands out as natural candidate for a streamlined comparison of branching rules, jointly for all instances. Unfortunately, although branch-and-bound schemes always compute a lower bound at the root node, an initial upper bound is often not available. Yet, since the main goal of the branching rules is often to deliver a faster increase of the lower bounds, it is natural to define a variant of pace^{OG} that is based on the pace at which the lower bound increases:

$$\text{pace}^{\text{LB}} = \frac{\text{time}}{LB^{\text{end}} - LB^{\text{root}} + \epsilon}; \quad (4)$$

Note that pace^{LB} is always well defined and measures the amount of time needed to improve the lower bound in one unit. This new KPI, pace^{LB} , will guide both the learning process described in Section 4.4 and the illustration of the results in Section 5.

4.3. Features

A key element of our approach is the identification and selection of input variables (features) that are able to explain and predict the behaviour of the aforementioned KPI for each branching rule. Following some of the ideas in [Liberto et al. \(2016\)](#), [Khalil et al. \(2016\)](#), and [Alvarez et al. \(2017\)](#) for MILP and our knowledge of PO problems, we have considered 34 features. They are summarized in Table 2. Features are categorized depending on whether they represent characteristics of the variables, constraints, monomials, coefficients, or other attributes of the polynomial optimization problems. Additionally, we have considered some novel features related to the graph representations of PO, VIG and CMIG, as described in Section 3.2. Many of the chosen features consist of counts or statistical measures (such as average, median, and variance) of the characteristics of the optimization problem. These are static features, as they only depend on the formulation of the original problem. They represent global information of the problem instance, without any node-specific information that changes throughout the branch-and-bound tree.

Table 2 List of features used by the different learning techniques.

Variables	No. of variables, variance of the density of the variables ¹ Average/median/variance of the ranges of the variables ² Average/variance of the no. of appearances of each variable ³ Pct. of variables not present in any monomial with degree greater than one Pct. of variables not present in any monomial with degree greater than two
Constraints	No. of constraints, Pct. of equality/linear/quadratic constraints
Monomials	No. of monomials Pct. of linear/quadratic monomials, Pct. of linear/quadratic RLT variables Average pct. of monomials in each constraint and in the objective function ⁴
Coefficients	Average/variance of the coefficients
Other	Degree and density of PO ⁵ No. of variables divided by no. of constraints/degree No. of RLT variables/monomials divided by no. of constraints
Graphs	Density, modularity, treewidth, and transitivity of VIG and CMIG ⁶

¹ The density of a variable is the number of different monomials in which it appears divided by the number of different monomials in the problem.

² The range of a variable is the difference between its upper and its lower bound.

³ For each variable we compute in how many constraints (and objective function) that variable is present.

⁴ For each constraint (and objective function) we compute the percentage of monomials present in it with respect to the total number of monomials of PO.

⁵ The density of PO is its number of different monomials divided by total number of monomials a problem with the same degree and number of variables might have.

⁶ We do not include transitivity for CMIG since, given the structure of the graph, its value is always 0.

4.4. Quantile Regression Learning Techniques

In supervised learning, data consists of input output pairs. Given an optimization problem of the form PO we use, as input, the information from the features vector. Regarding the output, we compute for each of the branching rules described in Section 3.3 the so-called $pace^{LB}$, that is, the ratio between the running time and the improvement in the lower bound. It is convenient to normalize the values of $pace^{LB}$ to $[0; 1]$ by dividing the best pace among all rules, i.e., the smallest, by the pace of each rule. We refer to this normalized pace as $npace^{LB}$. Thus, the closer this value is to one for a given branching rule, the better the performance of that rule. Note that the $pace^{LB}$ of a given branching rule is independent of the other rules, but its $npace^{LB}$ is not.

While our ultimate goal is the selection of the best branching rule, our variables of interest are quantitative outputs. Therefore, we address the problem as a regression problem. The task is to learn, for each rule, a real-valued function that models and predicts $pace^{LB}$. Then, given an instance, the rule that corresponds to the highest predicted normalized pace is selected.

The primary goal of classical regression models is to estimate the conditional mean of a response variable $y \in \mathbb{R}$, given the features vector $x \in \mathbb{R}^d$. Let y denote the $npace^{LB}$ of a given rule. In the preliminary computational experiment in Figure 1a we observe an asymmetric behavior of y (negative skewness), as well as the presence of outliers. This makes conventional regression models based on the conditional mean unsuitable. In this context, quantile regression is presented as a more appropriate methodology, since it does not make any assumptions about the distribution of the response variable and is more robust to the presence of outliers. Quantile regression aims to estimate $Q_\alpha(x)$, the α -conditional quantile of y , with $\alpha \in (0; 1)$. More precisely, $Q_\alpha(x) = \inf \{y : F(y|x) \geq \alpha\}$; being $F(y|x)$ the conditional cumulative distribution function of y . Thus, the focus is on the conditional distribution of the response variable rather than just on its conditional mean, $\mu(x) = E(y|x)$. In our context it might be relevant to find out, for example, that for certain instances a given rule exceeds, with high probability, a given value of y . We refer to [Koenker \(2005\)](#) and [Koenker et al. \(2017\)](#) for a detailed overview of quantile regression.

There are several quantile regression methods in the statistics and ML literature. We focus on nonparametric methods, which allow for more flexible specifications of the relation between the response and the explanatory variables than their parametric counterparts. We present below a quick overview of the three main methods used in our analysis.

Quantile generalized additive models. Generalized additive models assume that the conditional mean of the response has an additive structure such as $\mu(x) = \sum_{j=1}^m f_j(x)$, where the m additive terms represent predictors. The f_j 's are nonparametric functions which can be modeled using an expansion of basis functions; spline basis functions in our case. Hence, we get a flexible and interpretable statistical method, capable of characterizing non-linear regression effects. This leaves behind the traditional linear model that fails to represent the real effects of the inputs, which are normally non-linear. For detailed information of generalized additive models, we refer to [Hastie et al. \(2009\)](#). Quantile generalized additive models, see [Fasiolo et al. \(2021\)](#), extend the previous idea and assume that the conditional quantile of y is given as a linear combination of unknown smooth functions of the features vector. Thus, it uses a very flexible technique that combines the ideas of quantile regression and generalized additive models.

Stochastic gradient boosting for quantile regression. Boosting is a supervised learning technique designed for classification problems, but also applicable to regression. The basic idea of this ensemble method is to create a highly accurate prediction rule as a weighted combination of weak learners predictions. In this work we focus on the stochastic gradient boosting algorithm, introduced in [Friedman \(2002\)](#). The stochastic gradient boosting algorithm is a modification of the gradient boosting algorithm by [Friedman \(2001\)](#), that improves the performance by incorporating randomness in the procedure. More specifically, a subsample of the training data is randomly selected at each iteration, and used to fit the base learner (tree). We use stochastic gradient boosting for quantile regression, which generalizes the described method to the context of quantile regression by using an asymmetrically weighted absolute loss function ([Kriegler and Berk 2010](#)).

Quantile regression forests. Random forests are supervised learning algorithms that can be used both in classification and regression problems. They were introduced in [Breiman \(2001\)](#) as ensemble methods that grow several individual decision trees and aggregate them to make a single prediction. For regression, they give an approximation of the conditional mean of a response variable. Quantile regression forests, introduced in [Meinshausen \(2006\)](#), are a generalization of random forests that compute an estimation of the conditional distribution of the response variable by taking into account all the observations in every leaf of every tree and not just their average. One of the advantages of random forests, as well as quantile

random forests, is that we can use the complete dataset, without randomly splitting the data into training and test set, to evaluate the model using the Out-Of-Bag predictions.

5. Computational Results

5.1. Experimental Setup

We use three different sets of instances to test the performance of the learning approach on a variety of instances with different characteristics. The first one is taken from [Dalkiran and Sherali \(2016\)](#) and consists of 180 instances of randomly generated polynomial optimization problems of different degrees, number of variables, and density. The second dataset comes from the well known benchmark MINLPLib ([Bussieck et al. 2003](#)), a library of Mixed-Integer non-linear Programming problems. We have selected from MINLPLib those instances that are PO problems with box-constrained and continuous variables, resulting in a total of 168 instances. The third dataset comes from another well known benchmark, QPLIB ([Furini et al. 2018](#)), a library of quadratic programming instances, for which we made a selection analogous to the one made for MINLPLib, resulting in a total of 63 instances. Hereafter we refer to the first dataset as DS, to the second one as MINLPLib, and to the third one as QPLIB. Moreover, we have disregarded instances solved at the root node, since the branching rule plays no role in them. We also discarded unconstrained problems, as some of the features have the number of constraints in the denominator. Our final dataset consists of a total of 354 instances.

All the executions reported in this paper have been performed on the supercomputer Finisterrae II, at Galicia Supercomputing Centre (CESGA). Specifically, we used computational nodes powered with 2 deca-core Intel Haswell 2680v3 CPUs with 128GB of RAM and 1TB of hard drive.

The branching rules defined in Section 3.3 have been coded in RAPOSa ([González-Rodríguez et al. 2020](#)), a state-of-the-art implementation of the RLT outlined in Section 3.1. The analyses are performed using a configuration of RAPOSa with a minimal number of enhancements (calls to an NLP local solver and J-sets ([Dalkiran and Sherali 2013](#)); refer to [González-Rodríguez et al. \(2020\)](#) for details). Each instance of our final dataset was run with six different configurations of RAPOSa, based on the six different branching rules defined in Section 3.3. The time limit of each execution was set to one hour.

We follow the standard learning procedure. We randomly split the complete set of instances into two disjoint sets: the training set (70%) and the test set (30%). With the objective

of obtaining a better performance, we gather the instances into families related to the groupings defined in the corresponding libraries, where those belonging to the same group share similar characteristics. The within-family proportion of instances is maintained through the splitting process. Moreover, in order to gain robustness, we construct 10 partitions of the dataset into training and test data and report aggregate results over all the partitions.

As discussed in Section 4.4, when performing quantile regression one has to specify the conditional quantile to be estimated. Due to the negative asymmetry present in the KPI (see Figure 1), we have opted for values below the median ($\tau = 0.5$), where more differentiated behaviors can be observed between the branching rules and where it is more relevant to make a good selection. Although the different values evaluated lead to similar results, the best ones were obtained with $\tau = 0.3$, which is therefore the value used for the numerical analysis.

5.2. Numerical Results and Analysis

This section presents the numerical results that allow to evaluate the performance of the obtained ML-based branching rules. The objective is to show the impact of the ML approach when embedded within RAPOSa and also to provide further evidence of its effectiveness across various sets of instances. As a first step we compare the three ML-based rules obtained with the learning methodologies described in Section 4.4, that is, quantile generalized additive models (Q-GAM), stochastic gradient boosting for quantile regression (SGB-QR), and quantile regression forests (Q-RF). As RAPOSa includes several enhancements such as warm start and bound tightening (González-Rodríguez et al. 2020), we provide results of the ML-based branching without any of these enhancements (Baseline) and including them (WS and WS + BT).

Table 3 presents the results for the pace^{LB} ,² where the reported number corresponds to the average, across the 10 test sets, of the geometric mean pace^{LB} in the instances of each test set. The best branching rule is the original branching rule where the aforementioned average is smaller, that is, the one performing best. For the Baseline configuration, the original is the best branching rule while for the other two setups, the dual is the best branching rule. The optimal rule corresponds with choosing, for each instance, the original rule that performs best. We can see that the three learning approaches Q-GAM, SGB-QR, and Q-RF have similar performance. Since Q-RF performs slightly better, hereafter we use it as the reference

² We take $\epsilon = 0.001$ in Equation (4).

Table 3 ML-based rules' performance with respect to pace^{LB} (average over test sets).

	Baseline	WS	WS + BT
Best branching rule (across all instances)	1.875	1.534	0.819
ML-based branching rule:			
Q-GAM	1.439	1.181	0.782
SGB-QR	1.439	1.174	0.763
Q-RF	1.415	1.153	0.745
Optimal rule (instance by instance)	1.187	0.956	0.653
Improvement after learning (with Q-RF)	24.5%	24.8%	9.0%
Optimal improvement (upper bound for learning)	36.7%	37.6%	20.2%

ML-based rule. For the Baseline configuration, compared to selecting the best branching rule, Q-RF improves the pace by 24.5%. Importantly, this is not far from the improvement obtained by the optimal rule, 36.7%, which is an upper bound for our algorithm selection approach. For the WS + BT, the improvement with Q-RF is 9% and the optimal one is 20.2%. This is due to the fact that bound tightening techniques already improve the bounds significantly, reducing the search space and, therefore, reducing as well the room for improvement using different branching rules. Interestingly, the WS enhancement improves the KPI from 1.875 to 1.534, which is 18.2%. This value is lower than using Q-RF, which implies that given the choice between applying warm starting or a machine-learning branching approach, the latter may be better (and requiring no computational overhead). Table 4 shows that, when using out-of-bag predictions for the Q-RF methodology, we get similar results (the reported values correspond to the geometric means of pace^{LB} across all the instances of the dataset and the best original rule was always the dual).

Table 4 ML-based rules' performance with respect to pace^{LB} (Out-of-bag).

	Baseline	WS	WS + BT
Best branching rule (across all instances)	1.736	1.399	0.733
Out-of-bag Q-RF	1.191	1.050	0.678
Optimal rule (instance by instance)	0.983	0.810	0.601
Improvement after learning (with Q-RF)	31.4%	24.9%	7.5%
Optimal improvement (upper bound for learning)	43.4%	42.1%	18.0%

In order to get additional insights from the above results, we now explore in depth the performance of the ML-based branching rule Q-RF in the Baseline setup, building upon the

out-of-bag predictions:³ In Figure 2, the complete set of instances for the three libraries is used (total of 354 instances). The boxplots show that npace^{LB} is significantly closer to 1 for Q-RF than for any of the original rules. For the bar chart, it also validates that Q-RF is outperforming the other branching rules as it is ranked number 1 (best) in terms of the pace more often (38.98% compared to range which is 36.72%). The numbers inside the bars indicate average value of npace^{LB} for the problems in each rank position. Thus, even when Q-RF is ranked 6th (worst) its average normalized pace value is 0.7 which is also better than the corresponding value for all the other six branching rules (ranging from 0.43-0.69). In particular, the plots in Figure 2 show that, although the learned criterion chooses the best rule in less than 50% of the instances, it consistently chooses rules whose performance is very close to that of the best rule. This last observation also allows to understand why the improvement after learning with Q-RF, reported in Tables 3 and 4, is close to the optimal improvement.

Figure 2 Performance in Baseline configuration. Left, npace^{LB} boxplots of all the branching rules. Right, branching rules ranked from 1 (best) to 6 (worst), according to npace^{LB} .

In Figure 3 we separate the results in Figure 2 for the three sets of instances, DS, MINLPLib, and QPLIB. The number of instances within each library is indicated in parenthesis. There are some differences across datasets. For instance, the box plots show little variation in npace^{LB} across the branching rules for DS set compared to the other two sets.

³ We have also run the corresponding analysis for WS and WS+BT configurations, obtaining qualitatively similar results.

The maxrule is the worst one, while the other rules are close in terms of normalized pace. For MINLPLib, range performs slightly better than Q-RF as the number of times that range ranked first is 45.73% compared to 43.41% for Q-RF. For QPLIB, both Q-RF and eig-CMI strongly dominate in terms of performance, ranking first 52.72% and 56.36% of the instances, respectively. Note that for this set, the range rule, which performed very well for MINLPLib, is the worst branching rule for QPLIB. In particular, it is ranked 6th in around 20% of the instances with an average pace^{LB} of 0.32 in them, while the other rules, when ranked 6th, have an average pace^{LB} ranging from 0.69 to 0.96. Therefore, the ML-based rule seems to understand which are the features within and across datasets that drive the performance of the original rules.

Next, we further explore to what extent the ML-based rule manages to learn the patterns behind optimal rule selection. Figure 4 presents a comparison of the frequencies with which the original rules are chosen by the ML-based rule and the optimal rule. The behavior of the ML-based branching rule strongly resembles that of the theoretical optimal branching rule in all three datasets. For DS the dominant rule is range, with dual and eig-CMI also playing an important role. For MINLPLib, range dominates again and, surprisingly, maxrule is chosen quite often and so is again dual. For QPLIB, eig-CMI is chosen for almost 50% of the instances for both the ML-based rule and the optimal rule. Interestingly, in this last set of instances, the range rule is rarely selected as the criteria for branching. Overall, these findings are consistent with the results presented in Figure 3. Recall that our online learning is performed jointly on all sets of instances, so the results in Figure 4, showing that the ML-based rule correctly adapts to the instances in DS, MINLPLib, and QPLib, are a sign of the quality of the learning process and of its potential to learn from highly heterogeneous instances.

One advantage of the learning techniques used in our analysis is the interpretability. In particular, (quantile) random forests allow to assign scores to the different features and, hence, understand the ones that are the most important for the associated predictions. Since our Q-RF builds upon six independent regressions, one for each of the original branching rules, we have six different scores for each explanatory feature. Figure 5 presents, for each of the six regressions, the feature importance scores of the top 15 features overall. A first observation is that the novel graph-related features are important, as three of them appear

(a) Boxplots of the normalized pace using the described branching rules.

(b) Branching rules ranked from 1 (best) to 6 (worst), according to the normalized pace.

Figure 3 Performance in Baseline configuration. Comparison across the sets of instances.

in the top 8: CMIG.mod, VIG.mod, and CMIG.Dens. Further, CMIG.mod, the modularity on CMIG, is in second position. Different features have different importance score for learning different branching rules. In particular, graph-related features are important for the corresponding graph-based rules: CMIG.mod and CMIG.Dens are particularly important for eig-CMI and VIG.mod score is very high for eig-VI. Additionally, the most important features for learning range and dual branching rules are those related to the ranges of variables.

To conclude, we compare the performance of the different branching rules by showing the corresponding performance profiles (Dolan and Moré 2002). In Figure 6 we present the per-

Figure 4 Stacked bar graphs represent percentage of times that each branching rule is selected by Q-RF and the optimal branching rule.

Figure 5 Stacked bar graphs represent, for each branching rule, variable importance score of the features obtained using Q-RF.

formance profiles for the three sets of instances together and individually. Consistently with all the results we have already reported so far, the ML-based rule Q-RF performs noticeably better than all the original branching rules. From the DS library plot, Q-RF performs slightly better than 4 of the 6 branching rules (while `max` and `eig-VI` are significantly worse). For the MINLPLib library, Q-RF is clearly the best and range is in second position. For the QPLIB, Q-RF is slightly better than `eig-CMI` and dominates the other 5 branching rules. Therefore, despite having that the original rules performing best vary significantly for the different sets

of instances, Q-RF manages to slightly outperform the best rules in each individual set and solidly outperforms all the original rules overall.

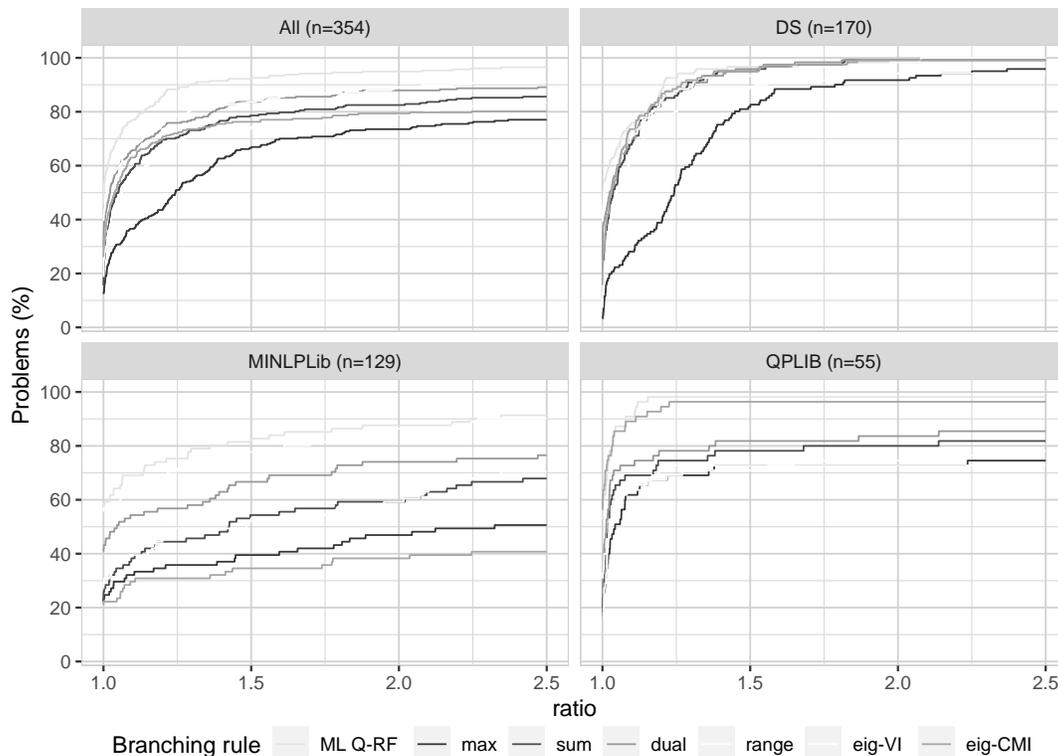


Figure 6 Performance profiles using the pace KPI for different sets of instances.

6. Conclusions and Future Research

We have presented a framework for learning rules for spatial branching and shown in different numerical experiments that the resulting gains in performance can be significant (up to 30%). These experiments have been performed jointly on a rich set of diverse instances taken from standard NLP libraries and from past literature. Interestingly, we have also shown that the interpretability of our framework allows to understand the importance of the different explanatory features for the resulting ML-based rule. This kind of understanding, which may shed light for the feature-design of future contributions, has been overlooked by most of the past literature on “learning to branch”.

As we briefly discussed in Section 4.1, the choice to perform learning only with static features limits the potential of our learning setting. Arguably, neglecting node by node information and limiting the scope of the algorithm selection to be instance specific, instead of node specific, may result in learning a less effective branching rule. Yet, given that this

paper is the first contribution on “learning to branch” in the non-linear setting, we believe that the static learning setting provides a natural starting point. First, since it involves no computational overhead for the computation of node-specific features, the differences in performance of the learned rule with respect to the original ones directly reflects the net effect of the learning process. Second, learning static rules allows for a clean definition of the “optimal” rule for each instance (this would not be entirely straightforward with dynamic rules, where the selection of the best “myopic” rule in each node might not lead to the best performance on the given instance). This is relevant since, once we pin down what can be achieved with perfect learning, one can easily assess the quality of the learning behind the rules resulting from our approach. Third, static learning provides a benchmark on which to assess the performance of more advanced dynamic techniques and, in particular, to evaluate to what extent the online computational overhead of node-specific features is compensated by the reduction on the tree size.

Given the above discussion, a prominent direction for future research is the incorporation of dynamic features and study the performance of the resulting rules taking the approach developed in this paper as a benchmark. Additionally, future research should aim at having richer portfolios of branching rules, incorporating for instance rules that build upon reliability branching or violation transfer (see, for instance, [Tawarmalani and Sahinidis \(2004\)](#), [Achterberg et al. \(2005\)](#), [Belotti et al. \(2009\)](#)).

Acknowledgments

This research has been funded by FEDER and the Spanish Ministry of Science and Technology through projects MTM2014-60191-JIN and MTM2017-87197-C3. Brais González-Rodríguez and Ignacio Gómez-Casares acknowledge support from the Spanish Ministry of Education through FPU grants 17/02643 and 20/01555, respectively. Beatriz Pateiro-López acknowledges support from Grant PID2020-116587GB-I00 funded by MCIN/AEI/10.13039/501100011033 and ED431C 2021/24 funded by Consellería de Cultura, Educación e Universidade. Bissan Ghaddar’s research is supported by Natural Sciences and Engineering Research Council of Canada Discovery Grant 2017-04185 and by the David G. Burgoyne Faculty Fellowship.

References

- Achterberg, T., Koch, T., and Martin, A. (2005). Branching rules revisited. *Operations Research Letters*, 33(1):42–54.
- Alvarez, A. M., Louveaux, Q., and Wehenkel, L. (2017). A machine learning-based approximation of strong branching. *INFORMS Journal on Computing*, 29(1):185–195.

- Baltean-Lugojan, R., Bonami, P., Misener, R., and Tramontani, A. (2019). Scoring positive semidefinite cutting planes for quadratic optimization via trained neural networks. Technical report, Optimization-online 6943.
- Belotti, P., Lee, J., Liberti, L., Margot, F., and Wächter, A. (2009). Branching and bounds tightening techniques for non-convex MINLP. *Optimization Methods and Software*, 24(4-5):597–634.
- Bengio, Y., Lodi, A., and Prouvost, A. (2021). Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421.
- Bienstock, D. and Muñoz, G. (2018). LP formulations for polynomial optimization problems. *SIAM Journal on Optimization*, 28(2):1121–1150.
- Bonami, P., Lodi, A., and Zarpellon, G. (2018). Learning a classification of mixed-integer quadratic programming problems. In *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 595–604. Springer.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.
- Bussieck, M. R., Drud, A. S., and Meeraus, A. (2003). MINLPLib—a collection of test models for mixed-integer nonlinear programming. *INFORMS Journal on Computing*, 15:114–119.
- Cappart, Q., Chételat, D., Khalil, E., Lodi, A., Morris, C., and Veli kovi , P. (2021). Combinatorial optimization and reasoning with graph neural networks. Technical report, arXiv:2102.09544.
- Clauset, A. (2004). Finding community structure in very large networks. *Physical Review E*, 70(6).
- Dalkiran, E. and Serali, H. D. (2013). Theoretical filtering of rlt bound-factor constraints for solving polynomial programming problems to global optimality. *Journal of Global Optimization*, 57(4):1147–1172.
- Dalkiran, E. and Serali, H. D. (2016). RLT-POS: Reformulation-linearization technique-based optimization software for solving polynomial programming problems. *Mathematical Programming Computation*, 8:337–375.
- Dolan, E. D. and Moré, J. J. (2002). Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91:201–213.
- Faenza, Y., Muñoz, G., and Pokutta, S. (2022). New limits of treewidth-based tractability in optimization. *Mathematical Programming*, 191:559–594.
- Fasiolo, M., Wood, S. N., Zaran, M., Nedellec, R., and Goude, Y. (2021). Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 116(535):1402–1412.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232.
- Friedman, J. H. (2002). Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378. Nonlinear Methods and Data Mining.

- Fulkerson, D. and Gross, O. (1965). Incidence matrices and interval graphs. *Pacific Journal of Mathematics*, 15(3):835–855.
- Furini, F., Traversi, E., Belotti, P., Frangioni, A., Gleixner, A., Gould, N., Liberti, L., Lodi, A., Misener, R., Mittelmann, H., Sahinidis, N., Vigerske, S., and Wiegele, A. (2018). QPLIB: a library of quadratic programming instances. *Mathematical Programming Computation*, 1:237–265.
- Gasse, M., Chételat, D., Ferroni, N., Charlin, L., and Lodi, A. (2019). Exact combinatorial optimization with graph convolutional neural networks. Technical report, arXiv:1906.01629.
- Gomes, C. P. and Selman, B. (2001). Algorithm portfolios. *Artificial Intelligence*, 126(1):43–62. Tradeo s under Bounded Resources.
- González-Rodríguez, B., Ossorio-Castillo, J., González-Díaz, J., González-Rueda, Á. M., Penas, D. R., and Rodríguez-Martínez, D. (2020). Computational advances in polynomial optimization: Raposa, a freely available global solver. Technical report, Optimization-online 7942.
- Gupta, P., Gasse, M., Khalil, E. B., Mudigonda, P. K., Lodi, A., and Bengio, Y. (2020). Hybrid models for learning to branch. In *NeurIPS*.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer, New York, NY, 2 edition.
- Khalil, E. B., Le Bodic, P., Song, L., Nemhauser, G., and Dilkina, B. (2016). Learning to branch in mixed integer programming. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*.
- Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs. Cambridge University Press.
- Koenker, R., Chernozhukov, V., He, X., and Peng, L., editors (2017). *Handbook of Quantile Regression*. CRC Press.
- Kriegler, B. and Berk, R. (2010). Small area estimation of the homeless in Los Angeles: An application of cost-sensitive stochastic gradient boosting. *The Annals of Applied Statistics*, 4(3):1234–1255.
- Latora, V., Nicosia, V., and Russo, G. (2017). *Complex networks: principles, methods and applications*. Cambridge University Press.
- Leyton-Brown, K., Nudelman, E., Andrew, G., McFadden, J., and Shoham, Y. (2003). A portfolio approach to algorithm select. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI'03*, pages 1542–1543, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Liberto, G. D., Kadioglu, S., Leo, K., and Malitsky, Y. (2016). Dash: Dynamic approach for switching heuristics. *European Journal of Operational Research*, 248(3):943–953.
- Linderoth, J. T. and Savelsbergh, M. W. P. (1999). A computational study of search strategies for mixed integer programming. *INFORMS Journal on Computing*, 11(2):173–187.
- Lodi, A. and Zarpellon, G. (2017). On learning and branching: a survey. *Top*, 25(2):207–236.
- Meinshausen, N. (2006). Quantile regression forests. *Journal of Machine Learning Research*, 7:983–999.

-
- Robertson, N. and Seymour, P. D. (1984). Graph minors. III. planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64.
- Sherali, H. D., Dalkiran, E., and Desai, J. (2012a). Enhancing rlt-based relaxations for polynomial programming problems via a new class of v -semidefinite cuts. *Computational Optimization and Applications*, 52(2):483–506.
- Sherali, H. D., Dalkiran, E., and Liberti, L. (2012b). Reduced rlt representations for nonconvex polynomial programming problems. *Journal of Global Optimization*, 52:447–469.
- Sherali, H. D. and Tuncbilek, C. H. (1992). A global optimization algorithm for polynomial programming problems using a reformulation-linearization technique. *Journal of Global Optimization*, 2(1):101–112.
- Tawarmalani, M. and Sahinidis, N. V. (2004). Global optimization of mixed-integer nonlinear programs: A theoretical and computational study. *Mathematical Programming*, 99(3):563–591.