



MARK TURNER¹, TIMO BERTHOLD², MATHIEU BESANÇON³,
THORSTEN KOCH⁴

Cutting Plane Selection with Analytic Centers and Multiregression

¹  [0000-0001-7270-1496](https://orcid.org/0000-0001-7270-1496)

²  [0000-0002-6320-8154](https://orcid.org/0000-0002-6320-8154)

³  [0000-0002-6284-3033](https://orcid.org/0000-0002-6284-3033)

⁴  [0000-0002-1967-0077](https://orcid.org/0000-0002-1967-0077)

Zuse Institute Berlin
Takustr. 7
14195 Berlin
Germany

Telephone: +49 30 84185-0
Telefax: +49 30 84185-125

E-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Cutting Plane Selection with Analytic Centers and Multiregression

Mark Turner^{1,2}, Timo Berthold^{1,3},
Mathieu Besançon², and Thorsten Koch^{1,2}

¹ Institute of Mathematics, Technische Universität Berlin, Germany

² Zuse Institute Berlin, Germany

³ Fair Isaac Deutschland GmbH, Berlin, Germany

{turner, besancon, koch}@zib.de, timoberthold@fico.com

Abstract. Cutting planes are a crucial component of state-of-the-art mixed-integer programming solvers, with the choice of which subset of cuts to add being vital for solver performance. We propose new distance-based measures to qualify the value of a cut by quantifying the extent to which it separates relevant parts of the relaxed feasible set. For this purpose, we use the analytic centers of the relaxation polytope or of its optimal face, as well as alternative optimal solutions of the linear programming relaxation. We assess the impact of the choice of distance measure on root node performance and throughout the whole branch-and-bound tree, comparing our measures against those prevalent in the literature. Finally, by a multi-output regression, we predict the relative performance of each measure, using static features readily available before the separation process. Our results indicate that analytic center-based methods help to significantly reduce the number of branch-and-bound nodes needed to explore the search space and that our multiregression approach can further improve on any individual method.

1 Introduction

Branch-and-cut is the algorithm at the core of most Mixed-Integer Programming (MIP) solvers. A key component of branch-and-cut is the resolution of Linear Programming (LP) relaxations of the original problem over partitions of the variable domains. *Cutting planes* – or cuts – tighten those relaxations around integer-feasible points. Given a MIP:

$$\operatorname{argmin}_{\mathbf{x}} \{ \mathbf{c}^\top \mathbf{x} \mid \mathbf{A} \mathbf{x} \leq \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{|\mathcal{J}|} \times \mathbb{R}^{n-|\mathcal{J}|} \} \quad (\text{P})$$

a cut is an inequality $\alpha^\top \mathbf{x} \leq \beta$ that is violated by at-least one solution of the LP relaxation but that does not increase the optimal value of the problem when added i.e. it is valid for (P). Thereby, the inequality added as a constraint to (P) tightens the relaxation, potentially increasing the relaxation's optimal value. The use of cutting planes is one of the crucial aspects to solving MIPs efficiently [3]. A well-designed cutting plane separation procedure often helps to reduce the branch-and-bound tree size while accelerating the overall solving process.

In MIP solvers, two key algorithms related to cuts are their generation and their selection. Cut generation is the problem of computing a set of cuts that tighten the relaxation and separate the current continuous relaxation solution from the feasible MIP solutions. Modern MIP solvers implement various general-purpose and specialised cutting plane generation algorithms. Since the generation of cuts is, in general, far less expensive than solving the LP relaxation, many cuts are generated from the same relaxation. The cut selection algorithm takes the set of all cut candidates generated so far and selects a subset that is actually added to the LP relaxation. This two-step process of generation and selection constitutes a single *separation round*.

At the root node, the MIP solver interleaves separation rounds with solving the enhanced LP relaxation until the branch-and-bound search is started. At other search tree nodes, the solver often only performs a limited cut loop, if at all. We focus in this paper on globally-valid cuts, i.e. cuts that are valid for the original problem, as opposed to *local cuts*, which are generated with additional local bounds at a node.

Cut selection is a classical trade-off problem: too little cutting leads to large enumeration trees; too much cutting to a small node throughput and numerical instability. However, carefully selected cuts can help

improve both the dual and the primal bound simultaneously by bringing the relaxation closer to the convex hull of feasible solutions. Since proximity to the convex hull is inherently hard to measure, cut selection methods often try to approximate it by various cheap measures, e.g., efficacy.

Efficacy, or *cutoff distance*, is used in commercial MIP solvers as one of the main criteria for whether to add a cut or not [1,9]⁴. Efficacy measures the shortest distance between the LP optimal solution \mathbf{x}^{LP} and the cut hyperplane $\alpha^\top \mathbf{x} \leq \beta$. The function eff that maps a cut, and the LP solution \mathbf{x}^{LP} to the efficacy is defined as:

$$\text{eff}(\alpha, \beta, \mathbf{x}^{LP}) := \frac{\alpha^\top \mathbf{x}^{LP} - \beta}{\|\alpha\|}$$

Introduced in [16], *directed cutoff distance* is the signed distance between the LP solution \mathbf{x}^{LP} and the cut hyperplane in the direction of a primal solution, $\hat{\mathbf{x}}$. The measure has the property that the directed projection of \mathbf{x}^{LP} onto the cut hyperplane is inside of the feasible region, and by using the best primal solution available, aims to cut in the direction of the optimal solution. We define the directed cutoff distance, with the function dcd as follows:

$$\text{dcd}(\alpha, \beta, \mathbf{x}^{LP}, \hat{\mathbf{x}}) := \frac{\alpha^\top \mathbf{x}^{LP} - \beta}{|\alpha^\top \mathbf{y}|}, \text{ where } \mathbf{y} = \frac{\hat{\mathbf{x}} - \mathbf{x}^{LP}}{\|\hat{\mathbf{x}} - \mathbf{x}^{LP}\|}$$

The last measure we consider, although not based on a distance, is *expected improvement*, see [31], which corresponds to the difference in objective between \mathbf{x}^{LP} and its orthogonal projection onto the hyperplane of a cut. We denote the measure as exp-improv , and define it as:

$$\text{exp-improv}(\alpha, \beta, \mathbf{c}, \mathbf{x}^{LP}) := \|\mathbf{c}\| \cdot \frac{\alpha^\top \mathbf{c}}{\|\alpha\| \|\mathbf{c}\|} \cdot \text{eff}(\alpha, \beta, \mathbf{x}^{LP})$$

In this paper, we propose new distance-based measures for the quality of cuts. These measures are designed to retain soundness properties in common cases that hinder the applicability, and to reduce the size of the search space, without focusing on runtime improvement. To this end, we perform extensive computational experiments to analyse the effects of our newly introduced measures and those that exist in the literature. Unlike previous beliefs, the choice of distance measure significantly impacts solver performance both in time and number of nodes, with different measures performing better on different groups of instances. Motivated by this observation, we design a multi-output regression model, which predicts the relative performance of each measure using static features readily available before the separation process. Finally, in Appendix A, we complement our investigation of cut selection with an analysis of the impact of dense cutting planes on solver performance. The scope of this paper is to establish a model that aims at reducing search space size rather than runtime.

2 Related Work

There is a prevailing sentiment in the MIP community, supported by a set of older computational studies, see [1,4,31], that inexpensive heuristics are sufficient for cut selection. Specifically, these studies suggest that cheap ranking metrics, predominantly efficacy, are sufficiently effective when combined with filtering mechanisms that ensure no two overly parallel cuts are added. The studies [1,31] argue that a weighted sum of different metrics is most effective for ranking cuts as opposed to any single metric.

Recently, the research focus for cuts in MIP has been on using deep learning to either calculate scores directly from a set of measures or to predict parameter values in scoring functions, see [7,27,19,6,29,23]. In [7] and [23] a neural network is trained to predict the objective value improvement of cuts when added. In [27],

⁴ Confirmed as a main criterion in FICO Xpress 8.14

a neural network is trained using evolutionary strategies to select Gomory cuts. A neural network is also trained in [19], this time using multiple instance learning, to map a set of aggregate cut features to a scoring function. In [6], the cut selection parameter space is shown to be partitionable into a set of regions, such that all parameter choices within a region select the same set of cuts. Finally, for cut scores based on weighted criteria, [29] provides an illustrative example of worst-case scenario for parameter grid search in the cut selection parameter space, and phrases learning cut selection parameters as a reinforcement learning problem.

Closest to our work are papers that introduced cut selection measures other than efficacy, which were still based on a notion of measuring distances. *Directed cutoff distance* was introduced in SCIP 6.0 [16], other measures such as *rotated cutoff distance* and *distance with bounds* were explored in [31], and *depth* was introduced in [25]. Measures based on non-distance arguments are also prevalent in cut selection, albeit often as smaller weighted complements to a distance measure, see *objective parallelism* and *integer support* in [1,31], and enumeration of lattice points in [20].

In the presented distance measures of Section 1, an LP-optimal solution is used as a reference point. Note, however, that this optimum is not necessarily unique, with the set of minimisers frequently being a higher-dimensional face of the LP-feasible region due to dual degeneracy [15]. The optimal solution used in these degenerate cases inevitably has a large impact on cut generation [33], and therefore cut selection. This has been noted in previous research, with work such as [13] using multiple LP solutions from different LP random seeds to generate different sets of cuts. Additionally, the patent [2], proposes using a second LP optimal solution at the same cutting round to filter cuts derived from the original LP solution. They provide an example, where a second LP solution that prioritises integrality is found, which then filters all cuts that do not separate it. Dual degeneracy is one major aspect motivating some of the newly-proposed measures in our work.

Finally, analytic centers have been used in other aspects of MIP solving, namely, presolving [8], cut generation [14], branching [8], and heuristics [5,22], motivating the measures introduced in this paper.

3 Contributions and methodology

The contribution of the present paper is threefold. First, we introduce new distance-derived cut quality measures, the most important of which utilise analytic centers, and analyse properties of interest for these measures in cases of dual degeneracy and infeasible projections. Secondly, we present an extensive set of computational experiments on the effectiveness of our new measures and those commonly found in practice, showing that the choice of cut selection measure does have a strong influence on root-node and tree-wide performance. Thirdly, we introduce a multi-output regression model that predicts a ranking of distance measures per instance from a set of root node features.

3.1 Analytic Center-Based Methods

We propose two new methods for measuring cut quality: *analytic efficacy*, and *analytic directed cutoff distance*. They are based on the analytic center of the polytope and of the optimal face, respectively. For a given bounded MIP formulation (P), the analytic center is unique and in the relative interior of the feasible region.

When a constraint $a^T x \leq b$ is tight for any feasible solution or in the presence of equality constraints, the analytic center is not well-defined due to a log-barrier term being $+\infty$. In practice, algorithms relax all log-barriers with a fixed slack constant as long as constraints are imposed on solutions.

While the analytic center is invariant under affine transformations of Problem (P), it can change with reformulations. E.g., the analytic center can be shifted by the presence of redundant inequalities. For our work, we assume that the formulation has already been presolved by the MIP solver. Presolving includes tightening both variable bounds and constraints, and removing both redundant constraints and variables. It thereby limits the extent of such edge cases.

Analytic efficacy As opposed to using the \mathbf{x}^{LP} extreme point returned by the MIP solver for eff calculations, we propose a new measure that uses the analytic center of the optimal face of the LP, \mathbf{x}^F . We define \mathbf{x}^F by the following using the notation from Problem (P), where \mathbf{A}_i , \mathbf{b}_i , \mathbf{u}_i , and \mathbf{l}_i are the i -th row or i -th entry of their respective matrix or vector:

$$\mathbf{x}^F := \underset{\mathbf{x} : \mathbf{c}^\top \mathbf{x} = \mathbf{c}^\top \mathbf{x}^{LP}}{\operatorname{argmin}} \left\{ - \sum_{i=1}^m \log(\mathbf{b}_i - \mathbf{A}_i \mathbf{x}) - \left(\sum_{i=1}^n \log(\mathbf{x}_i - \mathbf{l}_i) + \log(\mathbf{u}_i - \mathbf{x}_i) \right) \right\}$$

In practice, we find \mathbf{x}^F using the barrier (or interior point) algorithm on the LP relaxation without crossover, see [8]. This algorithm is available in all modern MIP solvers, and is often run concurrently to the simplex at the root node, making our cut selection algorithms of practical interest for MIP solving.

The obtained center \mathbf{x}^F is different from \mathbf{x}^{LP} only if the current problem presents dual degeneracy. The purpose of evaluating cuts with respect to how much they separate the center of the optimal face is to favour those which cut a greater part or potentially all of the optimal face, thereby more likely favouring an improvement in the dual bound. Compare Figures 1a and 1c for the intuition behind analytic efficacy.

Analytic directed cutoff distance This measure is inspired by directed cutoff distance, and uses the analytic center of the feasible region, \mathbf{x}^C , as opposed to the best incumbent solution $\hat{\mathbf{x}}$. Using the same notation as in Section 3.1, we define the analytic center of the feasible region, \mathbf{x}^C , as:

$$\mathbf{x}^C := \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ - \sum_{i=1}^m \log(\mathbf{b}_i - \mathbf{A}_i \mathbf{x}) - \left(\sum_{i=1}^n \log(\mathbf{x}_i - \mathbf{l}_i) + \log(\mathbf{u}_i - \mathbf{x}_i) \right) \right\}$$

\mathbf{x}^C is often interpreted as the “central-most” point of the polytope and can be computed efficiently by dropping the objective function $\mathbf{c}^\top \mathbf{x}$ and solving the resulting LP relaxation using the barrier algorithm without crossover. The motivation is that the incumbent is not necessarily representative of the feasible set of the MIP since it can be any point of the feasible set generated by heuristics. Furthermore, it introduces an additional source of variability in the cut selection since the best primal solution is prone to updates, particularly during root node cutting when many primal heuristics are employed. By contrast, \mathbf{x}^C is unique and deterministically determined for the LP relaxation. Compare Figures 1b and 1d for the intuition behind analytic directed cutoff distance.

For computational efficiency, we further introduce *approximate analytic directed cutoff distance*, which re-uses the analytic center, \mathbf{x}^C , from the previous separation round provided it is still LP-feasible. This is motivated by the intuition that the analytic center, as the “central-most” point, is rarely separated and remains close to the new analytic center after cuts have been added.

3.2 Multiple LP solutions

We also introduce two distance measures, which mitigate cases of dual degeneracy and do not rely on analytic centers, but rather rely on multiple LP-optimal vertices. Let \mathcal{X}^{LP} be a set of LP optimal solutions. We propose the measures *average efficacy*, denoted avgeff , and *minimum efficacy*, denoted mineff , which respectively take the average and minimum efficacy over all LP solutions in \mathcal{X}^{LP} .

$$\begin{aligned} \operatorname{avgeff}(\boldsymbol{\alpha}, \beta, \mathcal{X}^{LP}) &:= \sum_{\mathbf{x}^{LP} \in \mathcal{X}^{LP}} \frac{\operatorname{eff}(\boldsymbol{\alpha}, \beta, \mathbf{x}^{LP})}{|\mathcal{X}^{LP}|} \\ \operatorname{mineff}(\boldsymbol{\alpha}, \beta, \mathcal{X}^{LP}) &:= \min\{\operatorname{eff}(\boldsymbol{\alpha}, \beta, \mathbf{x}^{LP}) \mid \mathbf{x}^{LP} \in \mathcal{X}^{LP}\} \end{aligned}$$

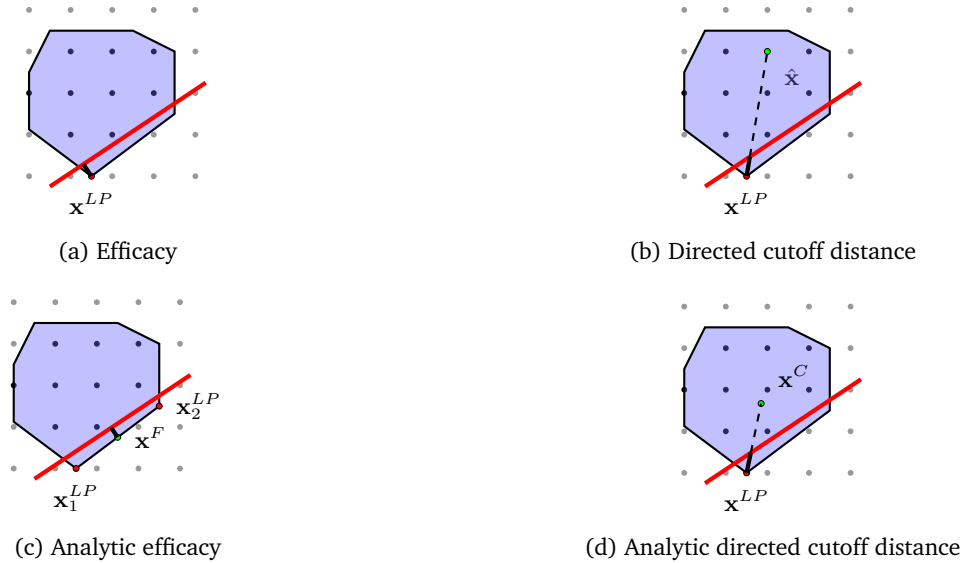


Fig. 1: A visualisation of distance measures. Note that in Figure 1c we see that there are two alternative optimal vertices.

3.3 Properties and limitations of the distance measures

In this section, we highlight situations in which the standard interpretations of some distance measures are limited; we introduce dominance consistency, a soundness property for distance measures, and establish cases under which it holds and under which it does not. Examples of dominance consistency also offer valuable insight into the geometry of different measures.

Efficacy, similarly expected improvement, uses the orthogonal projection of \mathbf{x}^{LP} onto the cut to measure distance. Unlike measures such as directed cutoff distance, the projection point might not be LP feasible, potentially making efficacy non-representative of the strength of the cut. A larger efficacy does not necessarily correspond to a larger part of the polyhedron being cut off nor to a better improvement in dual bound, see Figure 2a. Note that minimum and analytic efficacy would not assign the dotted cut a positive score. Efficacy would prefer the dotted cut even though it does not improve the dual bound while the dashed cut does. Figure 2b visualises a basic example of when the orthogonal projection used for efficacy is LP-infeasible and thus not necessarily a good proxy for cut quality. Analytic efficacy, minimum efficacy, and average efficacy help overcome some limitations of efficacy, namely the dependence on the vertex returned by the LP solve. They are, however, equivalent when the current relaxation is not dual-degenerate since they then compute the distance to the cut using the same, unique optimal vertex.

Directed cutoff distance heavily depends on the incumbent solution, which is typically obtained from a primal heuristic.⁵ This primal solution may be sub-optimal and near a corner of the feasible region, biasing cuts in an unfavourable direction. Additionally, the primal solution may be LP-infeasible for local relaxations of the branch-and-bound tree, reducing the applicability of directed cutoff distance.

Ultimately, one would like a distance measure to be a surrogate for solver efficiency, in our case: the size of the search space in terms of branch and bound nodes. Such a measure is impossible to quantify however, with solution fractionality being insufficient, and the closest analogue, ‘strong cutting’, which measures the dual bound improvement of an added cut, see [23], being computationally intractable. The holy grail of cut

⁵ At the root node, in particular, the incumbent – if existing – will always come from a heuristic, otherwise there would be no more cut rounds.



(a) The dashed cut removes the entire optimal face, and thereby is likely preferable. Efficacy however scores the cut higher (black lines). Analytic, minimum, and average efficacy prefer the dashed cut. (b) The projection from \mathbf{x}^{LP} onto the dashed cut is LP-infeasible, limiting its usefulness as a measure. The dotted cut is dominated but would still be selected by efficacy.

Fig. 2: Two cases showing the limitation of efficacy, in the presence of dual degeneracy 2a, and infeasible projection 2b. The blue polytope represents the LP feasible region, the black dots integer solutions, the red lines are proposed cuts, and \mathbf{x}^{LP} represents a LP optimal solution. In both cases, the dotted cut has a higher efficacy even though it is likely not preferable.

selection is to either identify a computationally tractable measure that is a reasonably good proxy for solver performance, or to learn to adaptively select a suitable measure based on the input instance.

In the following, we will formalise these considerations. Therefore, we first recall the basic concepts of cut dominance and feasible rays. A cut $(A) = (\alpha_A, \beta_A)$ dominates another cut $(B) = (\alpha_B, \beta_B)$ if all points of the polytope cut by (A) are cut by (B) and there exists a point cut by (A) not cut by (B) . We highlight that this definition of dominance is more general than, e.g., [32, Definition 9.2.1] since it only requires dominance to hold in the polytope and not in the whole space or the positive orthant. We define a *feasible ray* as a ray \mathbf{r} starting from an LP-feasible point \mathbf{x} for which there exists $\lambda > 0$ such that $\mathbf{x} + \lambda \mathbf{r}$ is LP-feasible. Such a ray always exists by convexity if the polyhedron is not a single point.

Definition 1 (Dominance consistency). Given a MIP (P) and a relaxation point to cut off \mathbf{x} , a distance measure noted $d(\mathbf{x}, \alpha_X, \beta_X)$ is dominance-consistent w.r.t. a set of cuts iff for any cut (A) and (B) in the set, $d(\mathbf{x}, \alpha_A, \beta_A) > d(\mathbf{x}, \alpha_B, \beta_B)$ implies that (A) is not dominated by (B) .

Note that with this definition, for a given MIP and LP solution \mathbf{x} , one set of cuts (e.g., from one separation round) can imply dominance consistency for a measure while another set of cuts (e.g., from another separation round) might not. We will see that for some measures, dominance consistency applies for all sets of cuts, all relaxation points, and all MIPs. Dominance consistency is a desirable property for a cut selection measure since a fully dominated cut will systematically be inferior to the dominating cut from the cut strength perspective.⁶

Proposition 1 (Consistency of Euclidean distance measures). All measures consisting of the Euclidean distance of a given point \mathbf{x} to the cut hyperplane are dominance-consistent with respect to any set of cuts if, for any two cuts in the set, the cut with the smallest distance measure cuts off \mathbf{x} and the projection of \mathbf{x} onto its hyperplane is LP-feasible.

Proof. Let (A) and (B) be the two cuts of the set. The proof directly applies to more cuts by induction. We assume w.l.o.g. that $d(\mathbf{x}, \alpha_A, \beta_A) > d(\mathbf{x}, \alpha_B, \beta_B)$. The set of cuts with distance measure $d(\mathbf{x}, \cdot, \cdot) = d(\mathbf{x}, \alpha_A, \beta_A)$ is a subset of the hyperplanes tangent to the sphere of radius $d(\mathbf{x}, \alpha_A, \beta_A)$ centered at \mathbf{x} . The cut (B) does not separate the projection of \mathbf{x} onto its own hyperplane, which is by assumption LP-feasible. The projected point lies on the sphere of radius $d(\mathbf{x}, \alpha_B, \beta_B)$ centered at \mathbf{x} , and is contained in the open

⁶ Note that we disregard other cut properties such as density, numerical stability, or orthogonality here.

ball of radius $d(\mathbf{x}, \alpha_A, \beta_A)$ centered at \mathbf{x} . As a tangent hyperplane to the sphere of radius $d(\mathbf{x}, \alpha_A, \beta_A)$, (A) therefore cuts off the projected point by at-least $d(\mathbf{x}, \alpha_A, \beta_A) - d(\mathbf{x}, \alpha_B, \beta_B)$ and can therefore not be dominated. \square

Proposition 1 directly applies to efficacy and analytic efficacy, with the key restriction that the projection of the relaxation point onto the cut must be LP-feasible, which excludes a majority of real instances. Furthermore, since eff is a linear function of \mathbf{x}^{LP} , the property also applies to avgeff as a distance measure, with the point to project \mathbf{x} being the average of the multiple LP solutions. Finally, we can construct counter-examples where dominance-consistency does not hold for measures based on the Euclidean projection of a point onto the cut hyperplane in cases where the projection is not LP-feasible, as shown in Figure 2b.

Proposition 2 (Consistency of the minimum efficacy). *Given the set of LP solutions \mathcal{X}^{LP} , we define the active solutions for a cut (α, β) as the subset $\arg \min_{\mathbf{x} \in \mathcal{X}^{LP}} \text{eff}(\alpha, \beta, \mathbf{x})$. mineff is dominance-consistent with respect to a set of cuts if for any two cuts (A) and (B) such that (B) has a strictly lower mineff , there exists an active solution \mathbf{x}_0 separated by (B) such that its projection onto the hyperplane of (B) is LP-feasible.*

Proof. Similarly to Proposition 1, the cut (B) forms a tangent hyperplane to the sphere centered at \mathbf{x}_0 and of radius equal to the score of (B). That point is not separated by (B) itself but (A) has to separate it by at-least the difference in score between (A) and (B). \square

Proposition 3 (Consistency of directed distance measures). *All measures based on the distance of a point \mathbf{x} to the cut in the direction of an LP-feasible point $\hat{\mathbf{x}}$ are dominance-consistent w.r.t any set of cuts which all separate \mathbf{x} .*

Proof. All points on the segment $[\mathbf{x}, \hat{\mathbf{x}}]$ are LP-feasible. Let (A) and (B) be two cuts of the set such that (A) has a strictly greater measure value. Since the measure corresponds to the length of the segment cut off by cuts, some points are separated by (A) only, (B) cannot dominate (A). \square

Proposition 3 notably applies to directed cutoff distance and (approximate) analytic directed cutoff distance. We note that dominance-consistency does not extend to exp-improv , even with feasible LP projections, with a counterexample visualised in Figure 3.




Fig. 3: Two cuts with LP-feasible orthogonal projections of \mathbf{x}^{LP} . The dashed cut is dominated, but has a better exp-improv score.

3.4 Multi-Output Regression

Machine Learning for MIP has mainly focused on classification tasks, e.g. should an algorithm be run with option A or B, or run at all. In our case, the output of interest is the (relative) advantage of distance measures in terms of some performance criterion. For some instances, different measures may result in (near-)identical performance, e.g. if the same, “obvious” subset of cuts is selected by all methods. Classifying these (near-)ties with one distance measure or choosing an arbitrary threshold to classify a measure as well-performing could prevent the model from fitting well on the important data points where the selection method is significant. We therefore pose our learning task as a multi-output regression that predicts the

relative performance of each method. We aim at an interpretable model to predict the preferred measure, with the goal of outperforming individual distance measures.

The feature space used as input to our model consists of: *dual degeneracy* (fraction of non-basic variables with zero-reduced cost), *primal degeneracy* (fraction of basic variables at their respective bounds), *solution fractionality* (fraction of integer variables with fractional LP values), *thinness* (fraction of equality constraints), and *density* (fraction of non-zero entries in constraint matrix). All features are obtainable at the root node before the separation process begins, are relevant to the separation process, and are easy to retrieve.

4 Experiments

We perform experiments on the MIPLIB 2017 collection set ⁷ [17], which we will now simply refer to as MIPLIB. We define a run as an instance random-seed pair for which we enforce exactly 50 separation rounds at the root node, with a maximum of 10 cuts to be added per round. We use default separators, but increase the amount of cuts that can be generated. Additionally, in order to reduce the variability of the solving process, restarts are disabled, no cuts are allowed to be added after the root node, and the best available MIPLIB solution is provided. All other aspects of the solver are untouched, with our experiments only replacing the cut scoring function in SCIP, not the selection algorithm itself, see [29]. The primary motivation is to assess the performance of different distance-based cut measures, and to determine for which instance characteristics a distance measure is effective. All results are obtained by averaging results over the SCIP random seeds {1, 2, 3}, and instances are filtered subject to Table 1, with 162 instances remaining. Three LP solutions are used for `mineff` and `avgeff`.

Criteria	% of instances removed
Tags: <i>feasibility, numerics, infeasible, no solution</i>	4.5%, 17.5%, 2.8%, 0.9%
Unbounded objective, MIPLIB solution unavailable	0.9%, 2.6%
Root optimal (any measure)	13.3%
Root node with separation rounds longer than 600s (any measure)	13.4%
No optional cuts generated (all measures)	2.3%
Numerical issues (any measure)	1.2%
Failed to prove optimality in branch and cut within 7200s (all measures)	21.7%

Table 1: Criteria for which we removed instances from the MIPLIB collection and percentages of instances affected by each criterion

For all experiments, SCIP 8.0.2 [9] is used, with PySCIPOpt [21] as the API, and Xpress 8.14 [12] as the LP solver. All experiments are run on a cluster equipped with Intel Xeon Gold 5122 CPUs with 3.60GHz and 96GB main memory. The code used for all experiments is available and open-source⁸. The structure of this section is as follows. In Subsection 4.1, we present results of our distance-based cut measures on root node restricted runs. In Subsection 4.2 we present results of our distance measures generalised to branch and cut. Finally, in Subsection 4.3 we present the performance of our support vector regression model on selecting distance measures.

4.1 Root Node Results

Table 2 provides a summary of all cut selection measures we evaluated. We compare head-to-head results on the primal-dual difference after 50 separation rounds. We say that a scoring measure has outperformed

⁷ MIPLIB 2017 – The Mixed Integer Programming Library <https://miplib.zib.de/>.

⁸ <https://github.com/Opt-Mucca/Analytic-Center-Cut-Selection>

Function	Measure	Description
<code>eff</code>	Efficacy	See Section 1
<code>dcd</code>	Directed cutoff distance	See Section 1
<code>a-eff</code>	Analytic Efficacy	See Subsection 3.1
<code>a-dcd</code>	Analytic directed cutoff distance	See Subsection 3.1
<code>app-a-dcd</code>	Approximate analytic directed cutoff distance	See Subsection 3.1
<code>avgeff</code>	Average efficacy	See Subsection 3.2
<code>mineff</code>	Minimum efficacy	See Subsection 3.2
<code>exp-improv</code>	Expected Improvement	See Section 1

Table 2: Summary of all distance measures

another for an instance, if it is at-least as good over all random seeds, and strictly better for at-least one. Curiously, we observe a clear hierarchy of distance-based cut measures for root-restricted dual bound performance over MIPLIB. That is, $a-dcd \geq app-a-dcd \geq a-eff \geq mineff \geq avgeff \geq dcd \geq eff \geq exp-improv$.

	a-dcd	app-a-dcd	a-eff	mineff	avgeff	eff	dcd	exp-improv
a-dcd	-	0.29/0.16	0.32/0.3	0.29/0.29	0.32/0.26	0.38/0.22	0.4/0.21	0.47/0.22
app-a-dcd	0.16/0.29	-	0.3/0.3	0.3/0.27	0.3/0.27	0.38/0.22	0.37/0.23	0.43/0.25
a-eff	0.3/0.32	0.3/0.3	-	0.25/0.23	0.22/0.2	0.31/0.14	0.37/0.26	0.41/0.23
mineff	0.29/0.29	0.27/0.3	0.23/0.25	-	0.14/0.13	0.28/0.12	0.32/0.26	0.36/0.24
avgeff	0.26/0.32	0.27/0.3	0.2/0.22	0.13/0.14	-	0.27/0.12	0.31/0.25	0.38/0.24
eff	0.22/0.38	0.22/0.38	0.14/0.31	0.12/0.28	0.12/0.27	-	0.26/0.25	0.32/0.25
dcd	0.21/0.4	0.23/0.37	0.26/0.37	0.26/0.32	0.25/0.31	0.25/0.26	-	0.36/0.26
exp-improv	0.22/0.47	0.25/0.43	0.23/0.41	0.24/0.36	0.24/0.38	0.25/0.32	0.26/0.36	-

Table 3: Entry coordinate (i, j) is a tuple of win / loss percentage over all instances for dual bound improvement of measure i over measure j .

A head-to-head comparison, while helpful for ranking measures, contains limited information on the distribution of performance. For this reason, we also visualise the primal-dual difference results using boxenplots in Figure 4, see [18] for a description. Figure 4 shows a comparison of each method against the so-called *virtual best*. Therefore, we divide, instance by instance, the average gap (over all seeds) by the best average gap among the eight methods. We observe similar results to Table 3 in that a-dcd, app-a-dcd, and a-eff are superior to other methods in terms of dual bound improvement. We conclude that using an analytic center for cut selection is beneficial for closing the primal-dual gap during root node cutting.

4.2 Branch and Bound Generalisation

While root node performance can be used as a surrogate for solver performance, there is no guarantee that results generalise to the entire solving process. We therefore extend our experiments to the branch-and-bound tree with a time limit of two hours. The head-to-head results of each distance measure for the number of nodes and solve time are displayed in Table 4. For the number of nodes, note that we removed all instances where a measure timed out. We observe in the number of nodes comparison that a-dcd remains the superior method, however the ordering of methods is now less clear. Most interesting is the drop in performance of app-a-dcd compared to the root node results, which suggests that the analytic center from previous separation rounds is often not a good direction for distance measures. This is supported by the fact that 26.5% of the analytic centers from previous rounds are LP infeasible. In the solve time comparison, we see that the ‘cheaper’ measures, `eff`, `dcd`, `exp-improv`, which require no additional LP solver calls, are

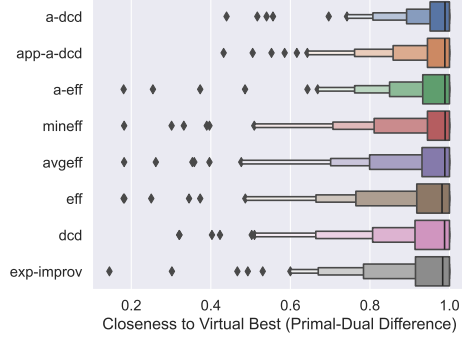


Fig. 4: Boxplots of distance measures root-node performance. Each instance compared to the virtual best and averaged over random seeds.

	a-dcd	app-a-dcd	a-eff	mineff	avgeff	eff	dcd	exp-improv
a-dcd	-	0.17/0.16	0.29/0.17	0.25/0.17	0.23/0.19	0.27/0.17	0.23/0.17	0.28/0.15
app-a-dcd	0.16/0.17	-	0.2/0.22	0.23/0.2	0.19/0.27	0.22/0.19	0.21/0.22	0.23/0.18
a-eff	0.17/0.29	0.22/0.2	-	0.15/0.15	0.2/0.2	0.17/0.17	0.2/0.19	0.18/0.17
mineff	0.17/0.25	0.2/0.23	0.15/0.15	-	0.11/0.17	0.17/0.1	0.17/0.18	0.18/0.15
avgeff	0.19/0.23	0.27/0.19	0.2/0.2	0.17/0.11	-	0.17/0.14	0.21/0.19	0.23/0.22
eff	0.17/0.27	0.19/0.22	0.17/0.17	0.1/0.17	0.14/0.17	-	0.16/0.2	0.19/0.2
dcd	0.17/0.23	0.22/0.21	0.19/0.2	0.18/0.17	0.19/0.21	0.2/0.16	-	0.19/0.17
exp-improv	0.15/0.28	0.18/0.23	0.17/0.18	0.15/0.18	0.22/0.23	0.2/0.19	0.17/0.19	-

(a) Number of nodes

	a-dcd	app-a-dcd	a-eff	mineff	avgeff	eff	dcd	exp-improv
a-dcd	-	0.37/0.14	0.33/0.16	0.28/0.15	0.28/0.17	0.22/0.3	0.16/0.3	0.19/0.29
app-a-dcd	0.14/0.37	-	0.23/0.2	0.25/0.23	0.24/0.23	0.17/0.35	0.15/0.3	0.16/0.36
a-eff	0.16/0.33	0.2/0.23	-	0.23/0.23	0.26/0.29	0.14/0.41	0.15/0.35	0.1/0.37
mineff	0.15/0.28	0.23/0.25	0.23/0.23	-	0.2/0.2	0.12/0.4	0.11/0.33	0.12/0.29
avgeff	0.17/0.28	0.23/0.24	0.29/0.26	0.2/0.2	-	0.1/0.37	0.12/0.34	0.12/0.32
eff	0.3/0.22	0.35/0.17	0.41/0.14	0.4/0.12	0.37/0.1	-	0.18/0.22	0.2/0.2
dcd	0.3/0.16	0.3/0.15	0.35/0.15	0.33/0.11	0.34/0.12	0.22/0.18	-	0.21/0.2
exp-improv	0.29/0.19	0.36/0.16	0.37/0.1	0.29/0.12	0.32/0.12	0.2/0.2	0.2/0.21	-

(b) Time

Table 4: Entry coordinate (i, j) is a tuple of win / loss percentage over all instances for measure i over measure j .

superior over the more ‘expensive’ measures. This suggests that while our introduced methods, especially a-dcd, can reduce the number of nodes and have better root node performance, the total solve time is not similarly improved. We note that while dcd is superior in Table 4b, we believe that our experimental design is overly favourable since we start with an optimal solution.

Similarly to Subsection 4.1, we visualise an instance-wide comparison to the virtual best of all measures for number of nodes and solve time in Figure 5. The results confirm our conclusion from Table 4a that a-dcd is the best performing measure, and that eff is the worst one w.r.t. number of nodes. All other measures, however, have similar distributions, making stronger conclusions difficult. We note that 90.4% of cuts have infeasible projections when scoring by eff, confirming the practical geometric limitation of efficacy presented in Section 3.3. For solve time, we observe the improved performance of ‘cheaper’ methods through their relatively high median values. We also observe that a-dcd has the smallest performance variability of all measures, while the standard eff has the largest performance variability. This implies that using an analytic center for cut selection can help to reduce performance variability, an interesting observation by its own.

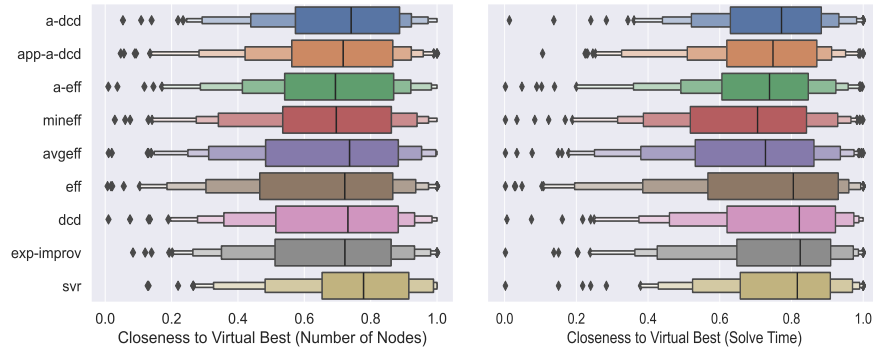


Fig. 5: Boxenplots of measures' tree performance. Nodes (left), time (right).

4.3 Regression Model Results

We have thus far observed that our newly introduced measures, especially a-dcd, have superior root node dual bound performance than traditional measures, and often result in smaller branch-and-bound trees. No single measure is however dominant, as seen in Tables 4a and 4b, with no single measure ever having less than 10% of instances as wins in the head-to-head contest. This motivates the need for an adaptive method, which decides on a distance measure at the start of the solving process that will best perform on the instance.

We use support vector regression (SVR) with a cubic kernel function, see [26], implemented in `scikit-learn` [24] with default parameters. We train on instance-seed pairs, with the virtual best number of nodes for each pair divided by the number of nodes under the distance measure as output. Our model was trained using 5-fold cross-validation, with 10% of pairs retained for validation. We were able to achieve comparable performance with regression forests and alternative kernels, likewise with default parameter sets. The final model was selected due to its ease of interpretation and potential embedding in a MIP solver.

We observe in Figure 5 that our trained model clearly outperforms any individual distance measure w.r.t. number of nodes. Further, when considering the shifted geometric mean of the number of nodes, it is 12% smaller than that of the best overall performing distance measure. This strong result does not generalise to solve time, however. The distribution looks similar to a-dcd, albeit with a better median. The shifted geometric mean of our model w.r.t. solve time is 8% larger than the single best-performing distance measure. Note that there are situations where available memory is a limiting factor – e.g., super-computing – which makes node savings important. Finally, we visualise the decision boundaries of the trained model over the two first principal components from a PCA of the original features, maintaining 71% of the explained variance. We determine decision boundaries with the largest regression value over all distances and visualise the result in Figure 6, with the component equations printed below.

5 Conclusion

In this paper, we reassessed the question of cut selection through the lens of distance measures. Motivated by geometric properties of polyhedra encountered in MIPs, we defined measures based on analytic centers and multiple LP solutions. We showed their performance, and more importantly, that the relative performance of distance measures can be learned for new instances with an interpretable and implementable model. We found that the introduced measures help to reduce root node gap, size of the branch and bound tree and performance variability. The focus of our work is on the improved evaluation of individual cuts; future directions will build upon these measures to enhance the whole separation process, incorporating combinations of cut measures, and generalising the promising node reductions to improved runtime.

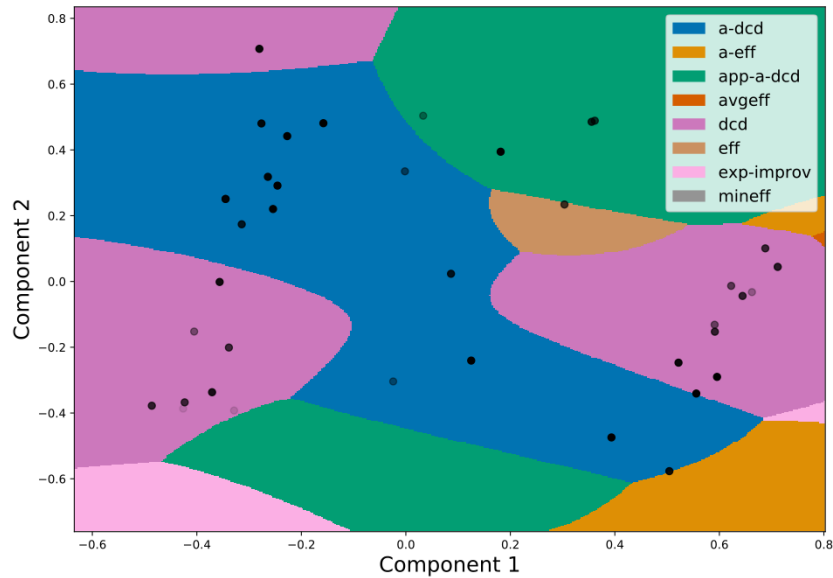


Fig. 6: Decision regions in transformed feature space. Dots are validation instances, with their opacity the relative performance of the predicted measure.

Component 1: $0.947 \text{ dual_deg} - 0.205 \text{ primal_deg} + 0.22 \text{ frac} - 0.089 \text{ thin} - 0.063 \text{ density}$

Component 2: $-0.27 \text{ dual_deg} - 0.733 \text{ primal_deg} + 0.467 \text{ frac} - 0.256 \text{ thin} + 0.326 \text{ density}$

Acknowledgements

The work for this article has been conducted in the Research Campus MODAL funded by the German Federal Ministry of Education and Research (BMBF) (fund numbers 05M14ZAM, 05M20ZBM). The described research activities are funded by the Federal Ministry for Economic Affairs and Energy within the project UNSEEN (ID: 03EI1004-C).

References

1. Achterberg, T.: Constraint integer programming. Ph.D. thesis, TU Berlin (2007)
2. Achterberg, T.: LP relaxation modification and cut selection in a MIP solver. US Patent US8463729B2 (2013), <https://patents.google.com/patent/US8463729B2/en>
3. Achterberg, T., Wunderling, R.: Mixed integer programming: Analyzing 12 years of progress. In: Facets of combinatorial optimization, pp. 449–481. Springer (2013)
4. Andreello, G., Caprara, A., Fischetti, M.: Embedding $\{0, 1/2\}$ -cuts in a branch-and-cut framework: A computational study. *INFORMS Journal on Computing* **19**(2), 229–238 (2007)
5. Baena, D., Castro, J.: Using the analytic center in the feasibility pump. *Operations Research Letters* **39**(5), 310–317 (2011). <https://doi.org/10.1016/j.orl.2011.07.005>, <https://www.sciencedirect.com/science/article/pii/S0167637711000824>
6. Balcan, M.F.F., Prasad, S., Sandholm, T., Vitercik, E.: Sample complexity of tree search configuration: Cutting planes and beyond. *Advances in Neural Information Processing Systems* **34** (2021)
7. Baltean-Lugojan, R., Bonami, P., Misener, R., Tramontani, A.: Scoring positive semidefinite cutting planes for quadratic optimization via trained neural networks. *optimization-online preprint 2018/11/6943* (2019)
8. Berthold, T., Perregaard, M., Mészáros, C.: Four good reasons to use an interior point solver within a MIP solver. In: *Operations Research Proceedings 2017*, pp. 159–164. Springer (2018)
9. Bestuzheva, K., Besançon, M., Chen, W.K., Chmiela, A., Donkiewicz, T., van Doornmalen, J., Eifler, L., Gaul, O., Gamrath, G., Gleixner, A., Gottwald, L., Graczyk, C., Halbig, K., Hoen, A., Hojny, C., van der Hulst, R., Koch, T.,

- Lübbecke, M., Maher, S.J., Matter, F., Mühmer, E., Müller, B., Pfetsch, M.E., Rehfeldt, D., Schlein, S., Schlösser, F., Serrano, F., Shinano, Y., Sofranac, B., Turner, M., Vigerske, S., Wegscheider, F., Wellner, P., Weninger, D., Witzig, J.: The SCIP Optimization Suite 8.0 (2021)
10. Chmiela, A., Muñoz, G., Serrano, F.: On the implementation and strengthening of intersection cuts for qcqps. *Mathematical Programming* pp. 1–38 (2022)
 11. Dey, S.S., Molinaro, M.: Theoretical challenges towards cutting-plane selection. *Mathematical Programming* **170**(1), 237–266 (2018)
 12. FICO Xpress Optimization. <https://www.fico.com/en/products/fico-xpress-optimization>, accessed: 2022-11-10
 13. Fischetti, M., Lodi, A., Monaci, M., Salvagnin, D., Tramontani, A.: Improving branch-and-cut performance by random sampling. *Mathematical Programming Computation* **8**(1), 113–132 (2016)
 14. Fischetti, M., Salvagnin, D.: Yoyo search: a bisection cutting-plane method (2009)
 15. Gamrath, G., Berthold, T., Salvagnin, D.: An exploratory computational analysis of dual degeneracy in mixed-integer programming. *EURO Journal on Computational Optimization* **8**(3-4), 241–261 (2020)
 16. Gleixner, A., Bastubbe, M., Eifler, L., Gally, T., Gamrath, G., Gottwald, R.L., Hendel, G., Hojny, C., Koch, T., Lübbecke, M., Maher, S.J., Miltenberger, M., Müller, B., Pfetsch, M., Puchert, C., Rehfeldt, D., Schlösser, F., Schubert, C., Serrano, F., Shinano, Y., Viernickel, J.M., Walter, M., Wegscheider, F., Witt, J.T., Witzig, J.: The SCIP Optimization Suite 6.0. Tech. Rep. 18-26, ZIB, Takustr. 7, 14195 Berlin (2018)
 17. Gleixner, A., Hendel, G., Gamrath, G., Achterberg, T., Bastubbe, M., Berthold, T., Christophel, P., Jarck, K., Koch, T., Linderoth, J., et al.: MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation* pp. 1–48 (2021)
 18. Hofmann, H., Kafadar, K., Wickham, H.: Letter-value plots: Boxplots for large data. Tech. rep., had.co.nz (2011)
 19. Huang, Z., Wang, K., Liu, F., Zhen, H.l., Zhang, W., Yuan, M., Hao, J., Yu, Y., Wang, J.: Learning to select cuts for efficient mixed-integer programming. arXiv preprint arXiv:2105.13645 (2021)
 20. Lodi, A., Pesant, G., Rousseau, L.M.: On counting lattice points and chvátal-gomory cutting planes. In: *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. pp. 131–136. Springer (2011)
 21. Maher, S., Miltenberger, M., Pedroso, J.P., Rehfeldt, D., Schwarz, R., Serrano, F.: PySCIPOpt: Mathematical programming in python with the SCIP optimization suite. In: *International Congress on Mathematical Software*. pp. 301–307. Springer (2016)
 22. Naoum-Sawaya, J.: Recursive central rounding for mixed integer programs. *Computers & operations research* **43**, 191–200 (2014)
 23. Paulus, M.B., Zarpellon, G., Krause, A., Charlin, L., Maddison, C.: Learning to cut by looking ahead: Cutting plane selection via imitation learning. In: *International Conference on Machine Learning*. pp. 17584–17600. PMLR (2022)
 24. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
 25. Poirrier, L., Yu, J.: On the depth of cutting planes. arXiv preprint arXiv:1903.05304 (2019)
 26. Smola, A.J., Schölkopf, B.: A tutorial on support vector regression. *Statistics and computing* **14**(3), 199–222 (2004)
 27. Tang, Y., Agrawal, S., Faenza, Y.: Reinforcement learning for integer programming: Learning to cut. In: *International Conference on Machine Learning*. pp. 9367–9376. PMLR (2020)
 28. Turner, M., Berthold, T., Besançon, M., Koch, T.: Cutting Plane Selection with Analytic Centers and Multiregression - Density Filtering Plots (Dec 2022). <https://doi.org/10.5281/zenodo.7433671>, <https://doi.org/10.5281/zenodo.7433671>
 29. Turner, M., Koch, T., Serrano, F., Winkler, M.: Adaptive Cut Selection in Mixed-Integer Linear Programming. arXiv preprint arXiv:2202.10962 (2022)
 30. Walter, M.: Sparsity of lift-and-project cutting planes. In: *Operations Research Proceedings 2012*, pp. 9–14. Springer (2014)
 31. Wesselmann, F., Stuhl, U.: Implementing cutting plane management and selection techniques. Tech. rep., Technical report, University of Paderborn (2012)
 32. Wolsey, L.A.: *Integer programming*. John Wiley & Sons (2020)
 33. Zanette, A., Fischetti, M., Balas, E.: Can pure cutting plane algorithms work? In: *International Conference on Integer Programming and Combinatorial Optimization*. pp. 416–434. Springer (2008)

A Cut density filtering

In this section, we focus on another key criterion to evaluate and select cuts: cut *density* or its complement *sparsity*, i.e. the number of non-zero coefficients of the cut. In complement to the evaluation of distance measures, we therefore assess the effect of filtering cuts based on their density, with the goal of shedding some light on a second aspect of cuts beyond the extent to which they cut off parts of the relaxation.

A.1 Motivation

The density of cuts is known to increase with the rank at which they are generated [30], where the rank of a cut is the number of separation rounds that were used to obtain the cut; such higher-rank cuts are known to be numerically unstable [11]. Dense cuts also typically slow down individual LP solves [30,10], largely due to LU factorisations being more difficult on dense matrices.

Density is a more nuanced signal than distance measures for whether to favour a cut. Unlike the latter, one cannot consider that higher sparsity levels are systematically better. Indeed, even though sparser cuts are numerically preferable, they can, all other things equal, be weaker by design, since one can always construct weaker sparse cuts out of dense ones.

We denote the density by `dens`, and relative density by `reldens`, and define them as:

$$\begin{aligned} \text{dens}(\boldsymbol{\alpha}) &:= |\{i \in \{1, \dots, n\} \mid \alpha_i \neq 0\}| \\ \text{reldens}(\boldsymbol{\alpha}) &:= \frac{\text{dens}(\boldsymbol{\alpha})}{n} \end{aligned}$$

To the best of our knowledge, the literature is relatively sparse on the effects of density on the performance and stability of MIP solvers. Specifically for stability, consider the following cut of dimension $N \in \mathbb{N}$, where each variable x_i is binary:

$$\sum_{i \in N} x_i \geq 1$$

When N is large enough, the individual variables can take value $\frac{1}{N} \approx 0$ in the continuous relaxation, thereby satisfying integrality tolerances, and the cut itself.

A.2 Computational experiments

In order to assess the impact of dense cuts on the solving process, we introduce filters that remove any cut over a given relative density threshold at the start of the cut selection algorithm, and before scoring. We restrict ourselves to scoring cuts exclusively with efficacy for this experiment, and use relative density thresholds 5%, 10%, 20%, 40%, and 80%, with the scoring functions denoted as `eff-05`, `eff-10`, `eff-20`, `eff-40`, and `eff-80`.

We display the results of our density experiments at the root node in Table 5a. We see a common and expected trend of larger primal-dual differences at the end of the root node when a more restrictive density filtering is used; removing the denser cuts weakens the root-node relaxation. A corresponding decrease in solve time also follows when the filtering becomes more restrictive, although we note that this decrease is more associated with the like decrease in cuts added and amount of separation rounds, as opposed to longer LP solves. For `eff-05`, we observe that the bound performance is better, albeit still worse than `eff`, in the $[0.8, 1]$ instance subset compared to that of all other instance subsets. This is due to a few instance outliers that have substantially worse performance under `eff-05` and maximum relative densities of 0.5 under `eff`. The detailed plots for this analysis are available on the archive [28].

Time-specific results of our density experiments when expanded to branch-and-bound are displayed in Table 5b. We observe that solve time decreases for instances that would not add cuts of relative density greater than 0.4 when using `eff-10` and `eff-20`, and decreases across all instance sets when using `eff-80`. For the most aggressive filtering `eff-05`, however, solve time increases across all instance subsets and fewer instances are solved overall. Aside from `eff-05`, and `eff-10` for the instance subset $[0.4, 1]$, all filtering methods result in at least as many instances being solved to optimality as under `eff`. An increase in LP iteration throughput can be observed across all filtering methods and instance subsets, supporting the claim that dense cuts slow down LP solves.

Finally, we display node-specific results of our density experiments when expanded to branch-and-bound in Table 5c. We observe a slight increase of nodes needed to prove optimality when using density-based filtering, although this is partially driven by outliers as evidenced in [28], highlighting the issues of performance variability for cut filtering. The exception of this node increase is `eff-10` and subsets containing instances with maximum relative densities less than 0.2 for which density filtering improves performance. We also observe for all filtering methods and instance subsets, an increase in node throughput and LP iterations per node, empirically demonstrating that dense cuts slow down LP solves.

max-density	instances	eff-05				eff-10				eff-20				eff-40				eff-80			
		time	bound	cuts	round	time	bound	cuts	round	time	bound	cuts	round	time	bound	cuts	round	time	bound	cuts	round
[0, 1]	162	0.69	1.13	0.53	0.6	0.75	1.04	0.62	0.67	0.81	1.03	0.69	0.72	0.88	1.01	0.77	0.79	0.97	1.01	0.89	0.9
[0.05, 1]	134	0.62	1.16	0.46	0.54	0.7	1.05	0.56	0.61	0.76	1.03	0.63	0.67	0.85	1.02	0.73	0.75	0.96	1.01	0.87	0.88
[0.1, 1]	112	0.56	1.2	0.4	0.48	0.63	1.07	0.5	0.56	0.71	1.04	0.58	0.62	0.82	1.02	0.68	0.71	0.95	1.01	0.85	0.86
[0.2, 1]	83	0.47	1.24	0.32	0.4	0.54	1.08	0.41	0.47	0.62	1.05	0.47	0.52	0.76	1.03	0.59	0.63	0.93	1.01	0.8	0.82
[0.4, 1]	57	0.39	1.31	0.25	0.33	0.46	1.09	0.33	0.38	0.53	1.06	0.37	0.42	0.65	1.04	0.46	0.51	0.9	1.01	0.72	0.74
[0.8, 1]	23	0.44	1.12	0.16	0.24	0.5	1.09	0.19	0.26	0.55	1.08	0.22	0.29	0.59	1.07	0.26	0.33	0.75	1.04	0.42	0.48

(a) Root node results. Time (s) is that to the end of all separation round. Bound is the primal-dual difference after the separation rounds. Cuts is the number of cuts added after all separation rounds. Rounds are the number of rounds.

max-density	instances	eff-05				eff-10				eff-20				eff-40				eff-80			
		time	it	it/s	Δ -sol	time	it	it/s	Δ -sol	time	it	it/s	Δ -sol	time	it	it/s	Δ -sol	time	it	it/s	Δ -sol
[0, 1]	162	1.15	1.16	1.04	-20	0.97	0.99	1.05	+0	0.99	1.01	1.03	+2	1.01	1.01	1.02	+2	0.99	0.99	1.01	+3
[0.05, 1]	134	1.19	1.19	1.04	-19	0.96	0.99	1.06	+0	0.99	1.02	1.04	+2	1.01	1.02	1.02	+2	0.99	0.99	1.01	+3
[0.1, 1]	112	1.24	1.26	1.05	-18	0.95	0.99	1.08	+0	0.99	1.02	1.05	+2	1.01	1.02	1.03	+2	0.99	0.99	1.01	+3
[0.2, 1]	83	1.39	1.42	1.06	-18	0.98	1.04	1.11	+0	0.98	1.02	1.07	+2	1.01	1.02	1.04	+2	0.98	0.98	1.02	+3
[0.4, 1]	57	1.63	1.62	1.05	-14	1.03	1.09	1.12	-1	0.99	1.03	1.07	+0	1.02	1.04	1.06	+2	0.98	0.97	1.02	+3
[0.8, 1]	23	1.11	1.14	1.11	+0	1.15	1.19	1.11	+0	1.17	1.17	1.06	+0	1.1	1.13	1.1	+3	0.95	0.93	1.05	+3

(b) Tree results for runtime (s) to optimality or 7200s time limit. It is the number of LP iterations, and it/s the iterations per second. Δ -sol is the difference in number of instance-seed pairs solved compared to `eff`.

max-density	instances	eff-05			eff-10			eff-20			eff-40			eff-80		
		nodes	nodes/s	it/nodes	nodes	nodes/s	it/nodes	nodes	nodes/s	it/nodes	nodes	nodes/s	it/nodes	nodes	nodes/s	it/nodes
[0, 1]	134	1.04	1.17	0.92	0.97	1.1	0.97	1.02	1.07	0.98	1.03	1.04	1.0	1.0	1.01	1.0
[0.05, 1]	110	1.06	1.21	0.9	0.96	1.12	0.96	1.03	1.09	0.98	1.04	1.05	0.99	1.0	1.02	1.0
[0.1, 1]	93	1.11	1.28	0.86	0.95	1.15	0.96	1.03	1.1	0.97	1.05	1.06	0.99	1.0	1.02	0.99
[0.2, 1]	67	1.09	1.31	0.87	1.0	1.2	0.95	1.04	1.15	0.95	1.07	1.08	0.99	1.0	1.03	0.99
[0.4, 1]	47	1.08	1.32	0.88	1.01	1.22	0.96	1.03	1.16	0.97	1.1	1.12	0.99	0.99	1.04	0.99
[0.8, 1]	21	1.1	1.38	0.88	1.1	1.31	0.91	1.08	1.25	0.95	1.08	1.22	0.94	0.99	1.1	0.97

(c) Tree results. Nodes is the number of nodes to prove optimality. Nodes/s is the node throughput per second, and it/nodes is the number of LP iterations per node.

Table 5: Summary of density results. Rows labelled $[d, 1]$ consider instances where the default runs with efficacy added a cut of relative density of at-least d . Entries are the instance-wise shifted geometric mean normalised by the shifted geometric mean using `eff`. The shifts are: (time, 1), (bound, 1), (cuts, 10), (round, 1), (it, 100), (it/s, 10), (nodes, 10), (nodes/s, 1), (it/n, 10).