



MARK TURNER¹, TIMO BERTHOLD², MATHIEU BESANÇON³

A Context-Aware Cutting Plane Selection Algorithm for Mixed-Integer Programming

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)

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

A CONTEXT-AWARE CUTTING PLANE SELECTION ALGORITHM FOR MIXED-INTEGER PROGRAMMING

A PREPRINT

 **Mark Turner***[†]
turner@zib.de

 **Timo Berthold***[‡]
timoberthold@fico.com

 **Mathieu Besançon**[†]
besancon@zib.de

July 17, 2023

ABSTRACT

The current cut selection algorithm used in mixed-integer programming solvers has remained largely unchanged since its creation. In this paper, we propose a set of new cut scoring measures, cut filtering techniques, and stopping criteria, extending the current state-of-the-art algorithm and obtaining a 5% performance improvement for SCIP over the MIPLIB 2017 benchmark set.

1 Introduction and Related Work

We focus on this paper on the selection of cutting planes, an essential component of the branch-and-cut framework, which is the main algorithmic paradigm to solve Mixed-Integer Linear Programs (MILP) classically defined as:

$$\underset{\mathbf{x}}{\operatorname{argmin}}\{\mathbf{c}^\top \mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{|\mathcal{J}|} \times \mathbb{R}^{n-|\mathcal{J}|}\} \quad (1)$$

A cut, parameterised by $(\alpha, \beta) \in \mathbb{R}^{n+1}$, 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 (1). It is used to tighten the Linear Programming (LP) relaxation of (1). Cuts are generated in rounds, where we refer to the process of computing cuts, applying a subset of the computed cuts, and resolving the LP as a *separation round*. In general, cuts are cheap to compute, and more are computed than actually applied to the LP. The process of deciding which cuts to add to the LP is called *cut selection*, and is the focus of this paper, in which we introduce a new cut selection algorithm.

Cut selection was considered largely “solved”, following computational results on a variety of MILP solvers [1, 2, 3], in that cheap heuristic selection rules were sufficient for good performance. Recently this conclusion has been challenged, with largely machine learning-driven research attempting to extract further performance, see [4] for an overview. Such research, however, is often limited in how it can be deployed in a MILP solver due to the complexity of deploying black-box predictors and their brittle generalisation. In this paper we introduce a new cut selection algorithm, which revises the three major aspects of cut selection, namely cut scoring measures, cut filtering techniques, and stopping criteria. We determine default parameter values of our algorithm by leveraging methodology [5, 6, 7] and software [8] from the algorithm configuration community. Our algorithm and the accompanying default values will be integrated in the next release of SCIP [9].

*Institute of Mathematics, Technische Universität Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany

[†]Zuse Institute Berlin, Department of Mathematical Optimization, Takustr. 7, 14195 Berlin

[‡]Fair Isaac Germany GmbH, Takustr. 7, 14195 Berlin, Germany

*Corresponding author.

2 New Cut Selection Techniques

Our new cut selection algorithm expands and generalises the existing one from SCIP [9] in three aspects, namely, cut scoring, cut filtering, and stopping criteria. For all three aspects, we incorporate additional information from the cut and from the current context in which the separation takes place.

2.1 Cut Scoring

The most studied aspect of cut selection is the scoring of cuts, which subsumes their ranking. This is both due to the large variety of scoring measures that exist, see [3, 10], and the common assumption of a fixed selection rule outside of scoring [11, 12, 13]. We introduce a new set of cut scoring measures that incorporate information from other non-cutting plane aspects of the MILP solving process.

2.1.1 Pseudo-cost Scoring

Pseudo-costs [14], the dominant decision-maker in state-of-the-art MILP branching rules [15], estimate scores for branching candidates based on the historical objective value improvement observed when branching on the candidates. It has recently been shown in [16] that cut scoring measures can complement pseudo-cost-based branching rules for improved solver performance. We address the reverse direction here, and use pseudo-costs to augment cut scores, where we let $\psi(x_i)$ be the pseudo-cost of variable x_i . We define the pseudo-cost scoring of a cut, $(\alpha, \beta) \in \mathbb{R}^{n+1}$, as the expected objective change as predicted by the pseudo-costs and express it as follows:

$$\text{psc}(\alpha, \beta, \mathbf{x}^{LP}) := \sum_{i, \alpha_i \neq 0} \psi(x_i) \left(\left| x_i^{LP} - \alpha_i \frac{\alpha^\top \mathbf{x}^{LP} - \beta}{\|\alpha\|^2} \right| \right) \quad (2)$$

We normalise the result by the maximum pseudo-cost score of a cut in the separation round, so all pseudo-cost scores are in the range $[0, 1]$.

2.1.2 Lock Scoring

Down-locks and up-locks, introduced in [1], can be interpreted for a variable x_i , as the number of constraints that “block” the shifting of x_i in the direction of the variable’s lower or upper bound. Let ζ_i^- be the number of down-locks and ζ_i^+ be the number of up-locks for variable x_i . We define the lock scoring of a cut, $(\alpha, \beta) \in \mathbb{R}^{n+1}$, as:

$$\text{loc}(\alpha) := \frac{1}{n} \sum_{i, \alpha_i \neq 0} (\zeta_i^- + \zeta_i^+) \quad (3)$$

As with the pseudo-cost measure, we normalise the result by the maximum among all cuts in the separation round so the scores are in the range $[0, 1]$. Additionally, as it is initially unclear if we want to reward a cut featuring variables with many locks, i.e. further restrict heavily restricted variables, we also introduce the complement of this measure, i.e. promote cuts that restrict variables with few restrictions.

2.1.3 Sparsity Scoring

Sparsity refers to the fraction of zero coefficients in a cut. In general, dense cuts slow down LP solves [17]. We therefore introduce a score that promotes sparse cuts, which, given a cut $(\alpha, \beta) \in \mathbb{R}^{n+1}$, and the parameters $\text{maxsps} \in \mathbb{R}_+$, and $\text{endsps} \in [0, 1]$, is defined as:

$$\text{sps}(\alpha) := \max\left\{ \text{maxsps} - \frac{\text{maxsps}}{\text{endsps}} \cdot \|\alpha\|_0, 0 \right\} \quad (4)$$

2.2 Cut Filtering

The standard approach in the literature for cut filtering is the parallelism-based approach [1, 2, 3]. The approach selects the highest-scoring cut, removes any remaining cut from consideration that is too parallel to the selected cut, and then repeats until some stopping criteria is met or no more non-filtered cuts are left to be selected. Such a procedure was shown to be absolutely necessary for good performance in [1, 2], however that is only w.r.t. trivial selection approaches. We therefore present and discuss two other approaches for cut filtering.

2.2.1 Density Filtering

Density-based filtering, unlike parallelism-based filtering, is performed at the start of the selection process. It immediately removes any generated cut above a given density threshold. Preliminary results for such an approach already exist, see [10, 18], and suggest that 40% density, i.e. a cut with at most 40% of its entries being non-zero, is a potentially desirable setting.

2.2.2 Parallelism-Based Penalties

It is possible to simply impose a penalty on any cut parallel to previously-added ones as opposed to outright removing it from consideration. In our algorithm, this involves reducing the score of all remaining cuts considered parallel to a selected cut. If the score falls below some threshold, which we set to 0, then the cut is removed from consideration.

2.3 Stopping Criteria

In general for computational studies, and MILP solvers by extension, there is a hard-coded limit on the number of cuts that can be added per round. Typically, this limit is not reached, with the standard parallelism filtering algorithm removing enough cuts by itself as detailed in Subsection 2.2. As we are employing alternate filtering algorithms, however, we believe it is necessary to have an instance-dependent limit. We therefore place a limit on the number of *non-zeros* added to the LP at each separation round, where the limit is a multiple of the number of columns in the LP.

3 Experiments and Computational Results

We conduct the following two experiments: First we select an appropriate training set for determining the best algorithm configuration of our new cut selector. Second we deploy the best found configuration on a test set to determine the actual MILP solver performance improvement.

Our training set is a combination of instances from the MIPLIB 2017 collection set⁴ [19], strIPlib⁵ [20], and SNDlib-MIPs [21]. We select the training set by removing any instance from consideration that takes less than 5 seconds or longer than 120 seconds to solve, takes longer than 10 seconds to presolve, or solves in less than 50 or more than 20000 nodes. This selection procedure is repeated with an optimal primal solution preloaded. From the remaining instances, we create a mapping to a feature space, where the features are the proportion of each constraint and variable type, the average row density, the maximum row density, and the objective vector density. Using this mapping, we select 40 instances that are maximally diverse by maximising the distance between any two selected instances in the feature space. Our test set is the MIPLIB 2017 benchmark set.

For all experiments, SCIP 8.0.3 [9] is used, with random seeds {1, 2, 3, 4, 5}, PySCIPOpt [22] as the API, and Xpress 9.0.2 [23] as the LP solver restricted to a single thread. For training, we use in non-exclusive mode a cluster equipped with Intel Xeon Gold 6342 CPUs running at 2.80GHz. For testing we run in exclusive mode on a cluster equipped with Intel Xeon Gold 5122 CPUs running at 3.60GHz, where each run is restricted to 48GB memory and a 2-hour time limit. The code used for all experiments is available and open-source⁶, and will be integrated in the next release of SCIP.

In addition to the new cut scoring measures introduced in Subsection 2.1, we use standard scoring measures from the literature. Namely, we use efficacy, *eff*, expected improvement, *exp*, objective parallelism, *obp*, and integer support, *isp* [3]. We normalise efficacy and expected improvement as in [13], ensuring values in the range [0, 1]. Notably, we disable directed cutoff distance due to it being unfairly advantaged by preloading the optimal solution in training. All cut scoring measures are included in a weighted sum rule, where the multipliers are hyperparameters.

3.1 Training

We use SMAC [8] to determine the parameter values of our new cut selection algorithm. We run 5 instances of SMAC on our training set using the random seeds 1-5 and a limit of 300 complete passes of SMAC over our training set. Our training metric is the ratio of shifted geometric means of solve time compared to SCIP default with a shift of one second.

Seed	1	2	3	4	5
Time (s)	0.874	0.909	0.853	0.917	0.891
Nodes	0.847	0.951	0.862	0.920	0.910

Table 1: Ratio of shifted geometric means (1s for solve time and 10 for nodes) for our approach vs SCIP default on instance-seed pairs of training set.

⁴MIPLIB 2017 – The Mixed Integer Programming Library <https://miplib.zib.de/>.

⁵<https://striplib.or.rwth-aachen.de/login/>

⁶<https://github.com/Opt-Mucca/Context-Aware-Cut-Selection>

The runs generated largely different parameter choices, with all choices resulting in an 8-15% improvement w.r.t. solve time over SCIP default on the training set. The generated parameter choices differed mostly in how to score a cut, where we could find no common parameter choices. For filtering and penalising cuts, however, there was an overwhelming consensus for the following parameters: dense cuts should be filtered at some 40 – 50% threshold, and parallel cuts should not be filtered, but rather lightly penalised. We additionally note that due to the instance restrictions put on our training set, we do not expect the level of performance to perfectly generalise to larger MILP instances.

3.2 Test

As the cut selection performance space is highly non-linear, aggregating the parameter choices from the various SMAC runs is difficult. We therefore use a combination of the results from SMAC (e.g. density and parallelism filtering), and expert knowledge on the impact of cuts on the overarching solving process.

MIPLIB 2017 benchmark set				
97% instance-seed pairs affected				
Time (s)	Nodes	Δ -Solved	Δ -Dual	Δ -Primal
0.95	0.92	14(668)	1(532)	-11(532)

Table 2: Ratio of shifted geometric means (10s for solve time and 100 for nodes) for our approach vs SCIP default over affected instance-seed pairs of the test set (solved to optimality for both cut selection algorithms). Δ is the difference in wins over the instance-seed pairs for amount solved, and the respective bounds for unsolved instances. Bold entries are when better than SCIP default.

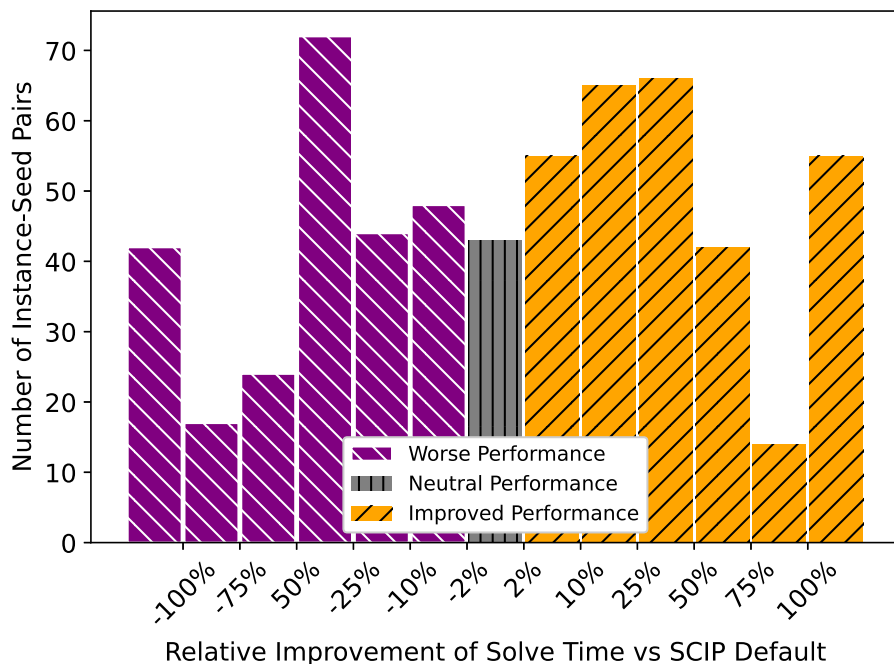


Figure 1: Bar plot of relative improvement of our approach vs SCIP default over affected instance-seed pairs of MIPLIB 2017 benchmark set.

Comparing our new cut selector to SCIP default, we observe in Table 2 an improvement of 5% for solve time and 8% for number of nodes. The distribution of relative time improvement over the instance-seed pairs is visualised in Figure 1, and is verified as statistically significant by a Wilcoxon signed-rank test (with a p-value below 5%).

4 Conclusion

In this paper we introduced a new cut selection algorithm, which combines new context-aware scoring measures, filtering methods, and stopping criteria. We incorporate the algorithm into SCIP, and obtain a 5% improvement in solve time over the MIPLIB 2017 benchmark set.

Acknowledgements

We thank Antonia Chmiela and Michael Winkler for helpful discussions about the experimental design. 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] Tobias Achterberg. *Constraint integer programming*. PhD thesis, TU Berlin, 2007.
- [2] Giuseppe Andreello, Alberto Caprara, and Matteo Fischetti. Embedding $\{0, 1/2\}$ -cuts in a branch-and-cut framework: A computational study. *INFORMS Journal on Computing*, 19(2):229–238, 2007.
- [3] Franz Wesselmann and U Stuhl. Implementing cutting plane management and selection techniques. Technical report, Technical report, University of Paderborn, 2012.
- [4] Arnaud Deza and Elias B Khalil. Machine learning for cutting planes in integer programming: A survey. *arXiv preprint arXiv:2302.09166*, 2023.
- [5] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration (extended version). *Technical Report TR-2010–10, University of British Columbia, Computer Science, Tech. Rep.*, 2010.
- [6] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Automated configuration of mixed integer programming solvers. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems: 7th International Conference, CPAIOR 2010, Bologna, Italy, June 14-18, 2010. Proceedings 7*, pages 186–202. Springer, 2010.
- [7] Lin Xu, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Hydra-MIP: Automated algorithm configuration and selection for mixed integer programming. In *RCRA workshop on experimental evaluation of algorithms for solving problems with combinatorial explosion at the international joint conference on artificial intelligence (IJCAI)*, pages 16–30, 2011.
- [8] Marius Lindauer, Katharina Eggenberger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Ruhkopf, René Sass, and Frank Hutter. SMAC3: A versatile Bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54):1–9, 2022.
- [9] Ksenia Bestuzheva, Mathieu Besançon, Wei-Kun Chen, Antonia Chmiela, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Oliver Gaul, Gerald Gamrath, Ambros Gleixner, Leona Gottwald, Christoph Graczyk, Katrin Halbig, Alexander Hoen, Christopher Hojny, Rolf van der Hulst, Thorsten Koch, Marco Lübbecke, Stephen J. Maher, Frederic Matter, Erik Mühmer, Benjamin Müller, Marc E. Pfetsch, Daniel Rehfeldt, Steffan Schlein, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Boro Sofranac, Mark Turner, Stefan Vigerske, Fabian Wegscheider, Philipp Wellner, Dieter Weninger, and Jakob Witzig. Enabling Research through the SCIP Optimization Suite 8.0. *ACM Transactions on Mathematical Software*, 2023.
- [10] Mark Turner, Timo Berthold, Mathieu Besançon, and Thorsten Koch. Cutting plane selection with analytic centers and multiregression. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 52–68. Springer, 2023.
- [11] Yunhao Tang, Shipra Agrawal, and Yuri Faenza. Reinforcement learning for integer programming: Learning to cut. In *International Conference on Machine Learning*, pages 9367–9376. PMLR, 2020.
- [12] Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik. Improved sample complexity bounds for branch-and-cut. In *28th International Conference on Principles and Practice of Constraint Programming (CP 2022)*, 2022.
- [13] Mark Turner, Thorsten Koch, Felipe Serrano, and Michael Winkler. Adaptive Cut Selection in Mixed-Integer Linear Programming. *arXiv preprint arXiv:2202.10962*, 2022.

- [14] Tobias Achterberg, Thorsten Koch, and Alexander Martin. Branching rules revisited. *Operations Research Letters*, 33(1):42–54, 2005.
- [15] Tobias Achterberg and Timo Berthold. Hybrid branching. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems: 6th International Conference, CPAIOR 2009 Pittsburgh, PA, USA, May 27-31, 2009 Proceedings 6*, pages 309–311. Springer, 2009.
- [16] Mark Turner, Timo Berthold, Mathieu Besançon, and Thorsten Koch. Branching via cutting plane selection: Improving hybrid branching. *arXiv preprint arXiv:2306.06050*, 2023.
- [17] Matthias Walter. Sparsity of lift-and-project cutting planes. In *Operations Research Proceedings 2012*, pages 9–14. Springer, 2014.
- [18] Herman Appelgren. Sparse cutting planes for mixed integer optimization, master’s thesis, 2023.
- [19] Ambros Gleixner, Gregor Hendel, Gerald Gamrath, Tobias Achterberg, Michael Bastubbe, Timo Berthold, Philipp Christophel, Kati Jarck, Thorsten Koch, Jeff Linderoth, et al. MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation*, pages 1–48, 2021.
- [20] Michael Bastubbe, Alexander Helber, Lukas Kirchner, Marco Lübbecke, Niklas Rieken, and Jonas Witt. strIPLib: structured Integer Program Library. Work in progress.
- [21] Mark Turner, Timo Berthold, Mathieu Besançon, and Thorsten Koch. SNDlib-MIPs: A new set of homogeneous MILP instances, June 2023.
- [22] Stephen Maher, Matthias Miltenberger, Joao Pedro Pedroso, Daniel Rehfeldt, Robert Schwarz, and Felipe Serrano. PySCIPOpt: Mathematical programming in python with the SCIP optimization suite. In *International Congress on Mathematical Software*, pages 301–307. Springer, 2016.
- [23] FICO Xpress Optimization. <https://www.fico.com/en/products/fico-xpress-optimization>. Accessed: 01-07-2023.