

Learn2Aggregate: Supervised Generation of Chvátal-Gomory Cuts Using Graph Neural Networks

Arnaud Deza^{*1}, Elias B. Khalil¹, Zhenan Fan², Zirui Zhou², Yong Zhang²

¹Department of Mechanical & Industrial Engineering, University of Toronto

²Huawei Technologies Canada

Abstract

We present *Learn2Aggregate*, a machine learning (ML) framework for optimizing the generation of Chvátal-Gomory (CG) cuts in mixed-integer linear programming (MILP). The framework trains a graph neural network to classify useful constraints for aggregation in CG cut generation. The ML-driven CG separator selectively focuses on a small set of impactful constraints, improving runtimes without compromising the strength of the generated cuts. Key to our approach is the formulation of a constraint classification task which favours a sparse aggregation of constraints. This, in conjunction with a careful constraint labeling scheme and a hybrid of deep learning and feature engineering, results in enhanced CG cut generation across five diverse MILP benchmarks. On the largest test sets, our method closes roughly *twice* as much of the integrality gap as the standard CG method while running 40% faster. This performance improvement is due to eliminating 75% of the constraints prior to aggregation.

Code — <https://github.com/khalil-research/ML4Cuts>

Extended version — <https://arxiv.org/abs/2409.06559>

1 Introduction

MILP is an established mathematical optimization framework used in many industrial applications such as operation room scheduling, supply chain, and transportation. The integrality requirement makes MILP NP-hard in general and thus highly challenging to solve exactly in practice. Cutting planes (or cuts) play a pivotal role in strengthening the linear programming (LP) relaxation of MILPs, leading to tighter bounds on the optimal integer value. Cuts can be separated heuristically or by optimizing an appropriate measure of cut quality. For instance, Gomory Mixed-Integer (GMI) cuts are read from the simplex tableau of the LP relaxation; a simple, fast heuristic which may nonetheless generate weak cuts.

In contrast, optimization-based separation, such as that of Chvátal-Gomory (CG) cuts, produces strong cuts but requires solving challenging auxiliary MILPs. A CG cut is a weighted aggregation of a subset of the original constraints followed by a rounding-down of the resulting coefficients.

However, optimizing over the extremely large set of possible aggregation weights for all constraints becomes computationally prohibitive for large MILPs. Screening rules to eliminate constraints for the special case of $\{0, 1/2\}$ -CG cuts have been proposed (Koster, Zymolka, and Kutschka 2009), but they have not been explored for general CG cuts.

This paper bridges the gap between heuristic and optimization-based cut generation by using machine learning (ML) to accelerate the optimization of aggregation-based cuts. We address the computational bottleneck of cut optimization by leveraging a model trained to screen constraints, eliminating unnecessary ones from consideration for aggregation. This allows us to solve a *reduced* separation problem over a much smaller subset of constraints and can be seen as a heuristic version of an exact optimization formulation for cut separation. Our method involves a care-




Figure 1: A binary two-variable problem. Constraints are solid and their associated normal vectors (excluding the bound constraints) are drawn, along with the maximization objective vector. Point “LP Opt.” is the solution of the LP relaxation whereas “ILP Opt.” is the desired integer optimum. The dashed cut, $x_1 + x_2 \leq 1$ in purple, is a strong cut which we would like to derive by aggregating the (normal vectors of the) dashed constraints; this cut tightens the LP such that its optimum is “ILP Opt.”. The normal vector of the desired cut is in the cone of the blue and orange constraints and thus can be derived by aggregating them with appropriate weights (e.g., multiplying both of them with 0.9, adding them up, then rounding down the resulting vector; this is a CG cut). The green constraint is not useful and can be excluded from consideration when aggregating.

^{*}These authors contributed equally.

ful combination of the following empirical observations and technical ideas:

1. On five families of MILP problems (facility location, packing, stable set, p-median, set covering), we observe that CG cuts generated via optimization (Fischetti and Lodi 2005) typically aggregate a tiny fraction of all constraints (Table 1). If those constraints could be identified *in advance of cut generation*, the optimization would likely be much faster. Figure 1 illustrates the idea.
2. CG optimization can produce many cuts of varying quality. We argue for learning only from non-dominated cuts, i.e., cuts that are binding at the LP optimum after they are added. This results in better labels for training a constraint classifier.
3. Our constraint classifier is a hybrid of feature engineering — representing constraints and variables using metrics established in the mathematical optimization literature — and graph neural networks (GNNs). This design choice makes training the GNN far easier than using a pure deep learning approach.

We evaluate the ML-based separator against heuristic CG cuts and exact separation on the five families of problems in terms of final integrality gap closure (IGC, i.e., how close the LP relaxation bound is to the optimal integer value) and total cut generation time. We find that:

1. On average, we remove 75% of the constraints, leading to mean IGC improvements of 2%, 23%, and 93% at time reductions of 57%, 64%, and 41% for small, medium, and large instances, respectively. These are substantial improvements that validate our motivating observation 1. above as well as the GNN design described in 3.
2. We show that learning only from non-dominated cuts (2. above) is strictly better than the more straightforward approach of learning from all CG cuts.
3. A GNN trained on small instances of a given problem performs well on medium and large instances. This can save substantial data collection and training time in some applications.

2 Background

Integer Linear Programming (ILP)

An ILP is an optimization problem involving n discrete variables $\mathbf{x} \in \mathbb{Z}^n$ that aims to minimize a linear objective function $\mathbf{c}^\top \mathbf{x}$ over m linear constraints $\mathbf{A}\mathbf{x} \leq \mathbf{b}$. Assuming the variables are nonnegative, an ILP reads as:

$$z^{IP} = \min\{\mathbf{c}^\top \mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \mathbb{Z}_+^n\} \quad (1)$$

This formulation induces two key polyhedra: the continuous polyhedron P and its integer hull P_I , defined as:

$$P := \{\mathbf{x} \in \mathbb{R}_+^n : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}, \quad (2)$$

$$P_I := \text{conv}\{\mathbf{x} \in \mathbb{Z}_+^n : \mathbf{A}\mathbf{x} \leq \mathbf{b}\} = \text{conv}(P \cap \mathbb{Z}_+^n). \quad (3)$$

The polyhedron P represents the feasible region of the LP relaxation, while P_I is the convex hull of feasible integer points. Solving ILPs is challenging given their NP-hard nature, but modern solvers effectively tackle large instances

using a combination of exact and heuristic methods. Solvers mainly rely on *Branch and Bound* (B&B) (Land and Doig 1960), a tree search algorithm that iteratively solves LP relaxations, tightens bounds by branching on fractional variables, prunes suboptimal branches, and uses cuts to further constrain P and approximate P_I . Here, we focus on how to get strong cuts quickly.

Cutting Planes

Cutting planes are valid linear inequalities to problem (1) of the form $\boldsymbol{\alpha}^\top \mathbf{x} \leq \alpha_0$, $\boldsymbol{\alpha} \in \mathbb{R}^n$, $\alpha_0 \in \mathbb{R}$. They are “valid” in the sense that adding them to polyhedron P is guaranteed not to cut off any integer solutions in P_I . Additionally, one seeks cuts that separate the current LP solution, \mathbf{x}_{LP}^* , from P_I . Although adding more cuts can help achieve tighter relaxations in principle, a clear trade-off exists: as more cuts are added, the size of the LP relaxation grows resulting in an increased cost in LP solving at the nodes of the B&B tree (Achterberg 2007). Adding too few cuts, however, may lead to a large number of nodes in the search tree as more branching is required.

A typical “cutting loop” consists of a number of rounds in which a new LP optimum is separated by adding one or more cuts until an integer LP optimum is found (indicating convergence to an optimal solution) or other termination criteria are met. Despite theoretical finite convergence results for the cutting plane method using Gomory cuts, numerical issues will often prevent convergence to an optimal solution in practice. It is thus typical to cut in rounds until the LP has been tightened sufficiently.

A standard metric to evaluate the cutting plane method is the *integrality gap* (IG). Let $z^t \in \mathbb{R}$ be the objective value of the LP after t rounds of cutting and let the bound difference be $g^t := z^{IP} - z^t \geq 0$. The integrality gap closure (IGC) (Tang, Agrawal, and Faenza 2020) is measured as :

$$IGC^t := 100 \cdot \frac{g^0 - g^t}{g^0} = 100 \cdot \frac{z^t - z^0}{z^{IP} - z^0} \in [0, 100] \quad (4)$$

IGC represents the percentage of the integrality gap that has been closed after t separation rounds.

General Valid Inequalities: Chvátal-Gomory Cuts

This work focuses on the widely used CG cuts (Chvátal 1973) which are generated by aggregating a subset of the original constraints of P followed by rounding. CG cuts are defined as the following type of valid inequality for P_I :

$$\boldsymbol{\alpha}^\top \mathbf{x} \leq \alpha_0, \text{ with } \boldsymbol{\alpha} = \lfloor \mathbf{u}^\top \mathbf{A} \rfloor \in \mathbb{Z}^n, \alpha_0 = \lfloor \mathbf{u}^\top \mathbf{b} \rfloor \in \mathbb{Z}. \quad (5)$$

The entries of \mathbf{u} are the nonnegative CG aggregation coefficients, one for each constraint. $\boldsymbol{\alpha}$ are the integer cut coefficients and α_0 is the right-hand side coefficient of the cut. In line with prior research on CG cuts (Fischetti and Lodi 2005), we focus on rank-1 cuts, i.e, cuts that only depend on the constraints defining P .

Chvátal-Gomory Cut Generation: CG-MIP

It is possible to generate a CG cut heuristically by reading the aggregation coefficients from the simplex tableau of

the LP relaxation. However, Fischetti and Lodi (2005) have shown that stronger cuts with higher IGC can be generated if one *optimizes* for an appropriate objective function. Concretely, to separate a point \mathbf{x}^* , Fischetti and Lodi (2005) propose an auxiliary MILP for CG cut separation, the CG-MIP in (6). Here, $J(\mathbf{x}^*) = \{j \in \{1, \dots, n\} : x_j^* > 0\}$ denotes the support of \mathbf{x}^* and δ is a small user-defined parameter (we use $\delta = 0.01$).

$$\begin{aligned}
& \max \sum_{j \in J(\mathbf{x}^*)} \alpha_j x_j^* - \alpha_0 & (6) \\
& \text{subject to } f_0 = \mathbf{u}^\top \mathbf{b} - \alpha_0, \\
& f_j = \mathbf{u}^\top \mathbf{A}_j - \alpha_j, & \forall j \in J(\mathbf{x}^*) \\
& 0 \leq f_j \leq 1 - \delta, & \forall j \in J(\mathbf{x}^*) \cup \{0\} \\
& 0 \leq u_i \leq 1 - \delta, & \forall i \in \{1, \dots, m\} \\
& \alpha_j \in \mathbb{Z}, & \forall j \in J(\mathbf{x}^*) \cup \{0\} \\
& \sum_{j \in J(\mathbf{x}^*)} \alpha_j x_j^* - \alpha_0 \geq 0.01.
\end{aligned}$$

As in all MILP-based separators, CG-MIP’s objective maximizes cut violation which must be at least 0.01. CG-MIP has $\mathcal{O}(n)$ integer and $\mathcal{O}(n + m)$ continuous variables, making its size linear in the original ILP’s. The constraints implement (5): the aggregation of the inequalities using the variable weights α_j and their rounding-down. CG-MIP can return a set of cuts using the off-the-shelf capability of modern MILP solvers to collect multiple feasible solutions (the solution pool). While effective, CG-MIP’s reliance on solving a MILP makes it computationally expensive, especially in the B&B tree, leading to its default deactivation in modern solvers (e.g., in SCIP (Bestuzheva et al. 2021)). Note that CG-MIP applies only to pure ILP but can be extended to MILP by the projected Chvátal-Gomory cuts of Bonami et al. (2008). More details of this simple extension, which we use for the facility location problem, are left to Appendix A.

3 Related Work

ML for MILP Bengio, Lodi, and Prouvost (2021) survey recent ML approaches to automating some decisions in MILP solvers. Notable examples include the learning of computationally challenging variable selection rules for B&B (Khalil et al. 2016; Gasse et al. 2019), learning to schedule heuristics (Chmiela et al. 2021), or estimating variable values (Nair et al. 2021; Khalil, Morris, and Lodi 2022).

Representing MILPs for ML In recent years, graph neural networks (GNNs) have emerged as a popular architecture for several ML applications for MILP (Cappart et al. 2023). GNNs can handle sparse MILP instances and are permutation-invariant, making them well-suited for representing MILP instances. The GNN operates on the so-called *variable-constraint graph* (VCG) of a MILP, first introduced by Gasse et al. (2019). The VCG has n variable nodes and m constraint nodes corresponding to the decision variables and constraints of (1). Edges between a variable node j and constraint node k represent the presence of variable x_j in constraint k whenever the weight of the edge, A_{jk} , is nonzero.

ML for cutting planes tasks Recent studies demonstrate success in integrating ML into cutting plane sub-routines such as cut selection (Tang, Agrawal, and Faenza 2020; Paulus et al. 2022; Turner et al. 2022), cut addition (Berthold, Francobaldi, and Hendel 2022), and cut removal (Puigdemont et al. 2024). For a comprehensive survey on this topic, we refer the reader to Deza and Khalil (2023).

These approaches are orthogonal to ours: they rely on existing (heuristic) cut generators and seek only to select from those generators’ cuts. The gap closed by a cut selection strategy is inherently upper-bounded by that of the whole set of cuts being considered. For example, while Tang, Agrawal, and Faenza (2020) successfully applied ML to select Gomory cuts from the simplex tableau, these weak, unoptimized cuts limited their experiments to very small instances. In contrast, we focus on larger instances where exact separation is challenging but produces stronger cuts. As a point of comparison, our smallest packing instances have 100 variables/constraints whereas the largest considered by Tang, Agrawal, and Faenza (2020) have 60; our largest have 500.

ML for cutting plane generation Despite this success, little attention has been paid to ML for cut *generation*, with only two papers to our knowledge. Chételat and Lodi (2023) frame cut generation as a continuous optimization problem over weights parameterizing families of valid inequalities (GMI cuts) that are optimized via gradient descent to maximize the dual bound of the LP relaxation with the generated cuts. Empirical results demonstrate improved dual bounds over classical GMI cuts read from the simplex tableau, although at a higher computational cost. This method applies on a per-instance basis and is not designed for the typical ML setting of learning over a distribution of instances.

In a similar vein, Dragotto et al. (2023) train a recurrent neural network to dynamically adjust the split-cut separator parameters (parametric LP) from Balas and Saxena (2008). This framework shows good in-distribution generalization and effectively reduces integrality gaps, albeit at a significant computational overhead. This is due to the use of `cvxpylayers` (Agrawal et al. 2019), which does not leverage sparse linear algebra, requiring recomputation of objects during both forward/backward passes. This restricts their experiments to small-scale MILP instances and very few cuts generated, making cut generation significantly slower than solving the parametric LP of Balas and Saxena (2008).

Both approaches highlight the potential of ML-driven cut generation but come with increased computational complexity, making them significantly slower than traditional methods, a bottleneck which we seek to overcome here.

4 Methodology

We propose *Learn2Aggregate*, a binary classification framework to identify useful constraints for aggregation, thereby reducing the computational complexity of an exact separator for a family of aggregation-based cuts, here CG cuts.

ML-Driven Constraint Aggregation

At each iteration t of the cutting plane method, the state S_t of the algorithm is represented by the tuple $\{\mathbf{A}, \mathbf{b}, \mathbf{x}_{LP}^t, U\}$

where \mathbf{x}_{LP}^t denotes the current (fractional) LP relaxation solution and the set U represents the k CG cuts generated thus far as characterized by their aggregation coefficients $U = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\}$. We consider CG-MIP to be our ‘‘oracle’’ for good cuts and use it in this cutting loop to generate the set U . The binary classification target y_i of a constraint i is defined based on its participation in the formation of any CG cut. Specifically, y_i is given by:

$$y_i = \mathbb{I}\{\exists j \in \{1, 2, \dots, k\} \text{ s.t. } u_{ji} > 0\}, \quad (7)$$

where u_{ji} is the i -th coefficient of the j -th aggregation vector \mathbf{u}_j and \mathbb{I} is the indicator function. In words, the binary label y_i is 1 iff the constraint i has a nonzero coefficient in one or more of the aggregation vectors U . Table 1 shows that for the representative set of MILPs we consider, the sparsity of the aggregation vectors \mathbf{u} of the cuts produced by CG-MIP is high, motivating our constraint classification approach.

Instance-Level Constraint Classification

Constraint classification can be done at the constraint or instance level. Constraint-level prediction involves a model that predicts whether a constraint should be utilized based on a feature vector that represents it. Instance-level prediction involves training a GNN that operates on the VCG representation of a MILP where the model simultaneously classifies all constraints. The instance-level approach has a better chance at capturing the complex interdependencies between variables and constraints than a local constraint-level one. Due to space limitations, we only present results for instance-level prediction.

We model the CG cut separation problem at iteration t using the VCG representation $G = \{\mathcal{C}, \mathcal{V}, \mathcal{E}\}$ of the MILP instance, where \mathcal{C} and \mathcal{V} denote sets of constraints and variables, respectively, and \mathcal{E} represents the edges between them. The constraint and variable feature matrices are represented by $\mathbf{C} \in \mathbb{R}^{m \times 53}$ and $\mathbf{V} \in \mathbb{R}^{n \times 18}$, respectively. Those features capture the relationship between the fractional LP solution being separated, \mathbf{x}_{LP}^t , and the constraints. Matrix $\mathbf{E} \in \mathbb{R}^{(m \cdot n) \times 1}$ contains edge features. A complete description of all features is included in Appendix C.

Integrating Efficient Cut Selection into Generation

Our goal is to learn to select constraints that lead to cuts maximizing dual-bound improvement (i.e., the IGC defined in (4)) rather than maximizing violation, a simple linear surrogate of cut quality. At round t , CG-MIP generates k cuts to separate the current fractional solution \mathbf{x}_{LP}^t . Which of these cuts is most useful in tightening the LP? We can find out by adding all k cuts to the current relaxation and solving the new LP to obtain the relaxed solution \mathbf{x}_{LP}^{t+1} . We argue that cuts that are tight to \mathbf{x}_{LP}^{t+1} , i.e., for which $\alpha^\top \mathbf{x}_{LP}^{t+1} = \alpha_0$, dominate the other cuts. As such, classification labels y_i for round t are generated based only on these tight cuts, rather than all k cuts. Only the deepest cuts are used to label the constraints, resulting in a refined signal for training.

Training and Neural Network Architecture

We train the graph convolutional neural network proposed in (Gasse et al. 2019), denoted as π , with a slight modification




Figure 2: Schematic of the GNN architecture with constraint features \mathbf{C} , variable features \mathbf{V} , and edge features \mathbf{E} .

in order to output logits z_i for each constraint i . The GNN model architecture is depicted in Fig. 2. The model initially embeds all input feature vectors \mathbf{c}_i and $\mathbf{v}_j, \forall i \in \mathcal{C}, j \in \mathcal{V}$ to a common embedding size h through a 2-layer perceptron. Then, the model performs one graph convolution involving two interleaved half-convolutions: (1) \mathcal{V} -side convolution from constraints to variables,

$$\mathbf{v}_j \leftarrow \phi_{\mathcal{V}} \left(\mathbf{v}_j, \sum_{i:(i,j) \in \mathcal{E}} \psi_{\mathcal{V}}(\mathbf{c}_i, \mathbf{v}_j, e_{i,j}) \right), \forall j \in \mathcal{V},$$

and (2) \mathcal{C} -side convolution from variables to constraints,

$$\mathbf{c}_i \leftarrow \phi_{\mathcal{C}} \left(\mathbf{c}_i, \sum_{j:(i,j) \in \mathcal{E}} \psi_{\mathcal{C}}(\mathbf{c}_i, \mathbf{v}_j, e_{i,j}) \right), \forall i \in \mathcal{C}.$$

In these equations, $\phi_{\mathcal{C}}, \phi_{\mathcal{V}}, \psi_{\mathcal{C}}$, and $\psi_{\mathcal{V}}$ are 2-layer perceptrons with ReLU activation functions. To obtain π , we apply a final 2-layer perceptron on the constraint nodes \mathbf{c}_i to produce logits z_i , which are then transformed into probability scores \hat{y}_i using a sigmoid function: $\hat{y}_i = \sigma(z_i) = \frac{1}{1+e^{-z_i}}$. The model is trained to minimize a weighted binary cross-entropy loss, accounting for the significant class imbalance where class-0 labels vastly outnumber class-1 labels. The loss for predictions $\hat{\mathbf{y}}$ of m target values \mathbf{y} is given by:

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) = -\frac{1}{m} \sum_{i=1}^m w_1 y_i \log(\hat{y}_i) + w_0 (1 - y_i) \log(1 - \hat{y}_i),$$

where w_1 and w_0 are weights assigned to class-1 and class-0, respectively, to counterbalance the skewed label distribution (see ‘Sparsity’ in Table 1). A tunable classification threshold of $\tau \in (0, 1)$ is used to classify a constraint: $\tilde{y}_i = \mathbb{I}\{\hat{y}_i \geq \tau\}$.

5 Experimental Setup

We now describe our experimental setup. More details can be found in the extended version Deza et al. (2024).

Datasets We experiment with five families of MILPs used in the research area of ML for MILP. They span a diverse range of combinatorial structures as can be seen in Table 1. For each family of problems, three instance sizes are considered. For each family and size, a dataset of 200 instances is generated. We use MIPLearn (Xavier et al. 2020) for the generation of PM and SS instances and Ecole (Prouvost et al.

Problem	Size	#Vars / Cons	Sparsity (%)
Stable Set (SS)	S	50 / 203	85.82
	M	75 / 535	90.43
	L	100 / 1264	93.14
P-Median (PM)	S	110 / 32	70.59
	M	240 / 47	58.14
	L	420 / 62	59.52
Set Cover (SC)	S	500 / 125	81.93
	M	500 / 250	92.74
	L	1000 / 250	98.03
Binary Packing (BP)	S	100 / 100	93.06
	M	300 / 300	96.82
	L	500 / 500	97.94
Facility Location (FL)	S	2550 / 2601	97.44
	M	5100 / 5151	98.91
	L	10100 / 10201	99.36

Table 1: MILP benchmarks and their sizes: S, M, and L refer to small, medium, and large respectively.

2020) for the generation of FL and SC instances. Details regarding problem formulation and instance generation is also included in Appendix B. To collect training data (on a random subset of the 200 instances), we deploy the full CG-MIP separator in a pure cutting plane method to collect features/labels as well as baseline performance for the full separator. Table 1 summarizes the instance size statistics as well as their average sparsities as measured by the fraction of constraints in an instance that receive a label of zero based on the procedure described in the last paragraph of Section 4.

Baselines We use two baselines to evaluate the proposed ML-driven separator: (1) the full exact CG-MIP separator and (2) heuristic GMI cuts from the simplex tableau. The former can produce strong cuts at a high computational cost whereas the latter produces weak cuts extremely quickly.

Cutting plane setting We evaluate our approach in a cutting plane method where CG-MIP is used to generate CG cuts at every round. As done by Fischetti and Lodi (2005), GMI cuts read from the simplex tableau are added to the root LP relaxation prior to any CG cut generation. We restrict the maximum number of cut generation rounds to 100 and terminate early if IGC stagnates over 7 rounds to avoid diminishing returns. Additionally, motivated by the CG-MIP implementation in SCIP, we use a complemented MIR cut generation heuristic based on CG-MIP’s solution pool.

CG-MIP parameters CG-MIP can run unnecessarily long without appropriate safeguards. CG-MIP’s execution is stopped if an integer (optimal) solution has been found or if one of the following conditions are met: (1) A 15-second time limit is hit, (2) a feasible solution limit of 5000 is hit, or (3) an incumbent (i.e., improving) solution limit of 1000. The conservative time limit of 15 seconds is used to return cuts fast if any are found. Otherwise, the time limit is doubled and the solution limit is set to 1, i.e., the solver terminates on finding the first cut. The time limit is doubled until a total time limit of 180 seconds is reached, in which case the cutting plane method terminates. Additionally, to ensure CG-MIP focuses on finding feasible solutions rather

than proving optimality, we set SCIP’s internal parameter *emphasis* to feasibility (Appendix D). We use the same parameters when deploying the trained classifier to solve a reduced CG-MIP separator, ensuring a fair comparison.

Training details For each dataset, we use 100, 50, and 50 instances for training, validation, and testing. GNN hyperparameters are tuned using 80 random search trials. After training, the GNN’s classification threshold (between zero and one) is tuned to maximize the F1-score on validation data. Tuning this threshold is crucial for striking a balance between accurately identifying useful constraints and avoiding the risk of over-constraining CG-MIP, which could lead to infeasibility in the cut generation process. Models are trained with Pytorch 1.10 (Paszke et al. 2019) and SCIP 8.0.0 (Bestuzheva et al. 2021) is used as the MIP solver. More details can be found in Tables 8 and 9 of the appendix.

Evaluation metrics To compare the reduced and full CG-MIP separator, we focus on size reduction and end-to-end metrics. The former is measured by the average percentage of constraints excluded from aggregation based on the classifier’s predictions. The latter focus on the final IGC after successive cut generation rounds and cumulative CG-MIP runtime in seconds. For these, we report the mean IGC and time using shifted geometric means with shifts of 0.5 and 5, respectively, as is standard in MILP benchmarking (Achterberg 2007); see Appendix D for a definition. Additionally, we compare both separators by calculating the ratio of the means of the reduced separator for IGC and time, to those of the full separator. We perform a win/loss/tie analysis where we compare the reduced separator against the full separator for each instance. A *win* (or *loss*) in time indicates that the reduced separator is 10% faster (or slower), while a *win* (or *loss*) in IGC signifies a 1% higher (or lower) final IGC. An *absolute win* occurs when time favours the reduced separator whilst beating or tying in IGC, and an *absolute loss* occurs when both metrics incur a loss.

6 Experimental Results

Does ML-driven constraint selection accelerate cut generation and close more gap? We answer this question before analyzing the effect of using only tight cuts for constraint labeling and our models’ size generalization capabilities.

ML closes more gap in less time. Table 2 shows, on unseen test instances, the ratio of the shifted geometric means between the reduced separator and the full separator for the final IGC, total cut generation time (seconds), and number of cutting plane rounds. It also includes the percentage of constraints removed by the reduced separator (% Size Reduction) and the absolute win/loss counts. The results demonstrate that we are able to achieve significant reductions — above 60% size for all problems — in the size of the CG-MIP separator while maintaining comparable or superior final IGC (most IGC ratios exceed 1). For instance, across binary packing datasets, the reduced separator achieves an average IGC improvement of 20% with a 60% reduction in cut generation time, while removing over 94% of constraints. We observe a consistent trend across problems and sizes, on

Dataset	Size	Ratio		Abs. Perf Win/Loss	% Size Reduction
		IGC \uparrow	Time \downarrow		
SS	S	0.96	0.33	84 / 2	63.50
	M	1.28	0.19	80 / 2	67.99
	L	2.59	0.9	48 / 4	72.53
PM	S	1.0	0.3	84 / 0	62.91
	M	1.0	0.56	74 / 4	49.30
	L	1.0	0.65	82 / 2	37.98
SC	S	0.97	0.47	80 / 0	68.33
	M	1.31	0.23	78 / 4	67.14
	L	2.99	0.66	60 / 6	70.05
BP	S	1.01	0.49	84 / 6	90.95
	M	1.14	0.31	80 / 2	95.33
	L	1.45	0.37	76 / 8	96.55
FL	S	1.14	0.54	28 / 4	92.80
	M	1.41	0.51	34 / 6	96.22
	L	1.63	0.39	50 / 2	97.18
Average	S	1.02	0.43	72 / 2	75.70
	M	1.23	0.36	69 / 4	75.20
	L	1.93	0.59	63 / 4	74.86

Table 2: Ratios of mean test IGC and time of the reduced separator to those of the full separator followed by the absolute profile of win/loss counts and CG-MIP size reduction. Ratios are bolded when the reduced separator outperforms the full separator. Higher IGC ratios and lower time ratios are better for our method, respectively.

average removing 75% of constraints while achieving a 2%, 23%, and 93% mean IGC improvement at a 57%, 64%, and 41% mean time improvement for small, medium, and large instances respectively. In the appendix, we include raw values for mean IGC and time in Table 6 as well as test accuracy and recall in Table 5 .

Zooming in on the cutting loops. Figure 3 shows the evolution of the primary evaluation metric, the IGC, over rounds of cutting by both the ML-reduced and full separators. Of note is that the former typically dominates the latter even very early on in the process. For most of these instance sets, our method typically terminates sooner than the full separator; this may explain the running time reductions observed in Table 2. However, for large instances of Set Cover and Stable Set (last two bottom-right sub-plots), ML runs for more rounds. This may be explained by the hardness of these two problems at scale, while noting that they also exhibit the largest improvements in IGC: ratios of 2.59 and 2.99, respectively (Table 2). To complement Fig. 3, the appendix includes Fig. 6, a plot of the distributions of the number of rounds for both methods.

Is there value in optimization-based separation? To motivate the use of optimization-based separation, we compare both full and reduced CG-MIP to GMI cuts read from the simplex tableau in Fig. 4. The vertical axis corresponds to an IGC value and the horizontal axis corresponds to the percentile of test instances for which the final IGC is at most that. For most datasets, the reduced separator (in red) has the largest area under its curve, followed by the full separator (blue) and GMI cuts (yellow). This not only shows the ben-

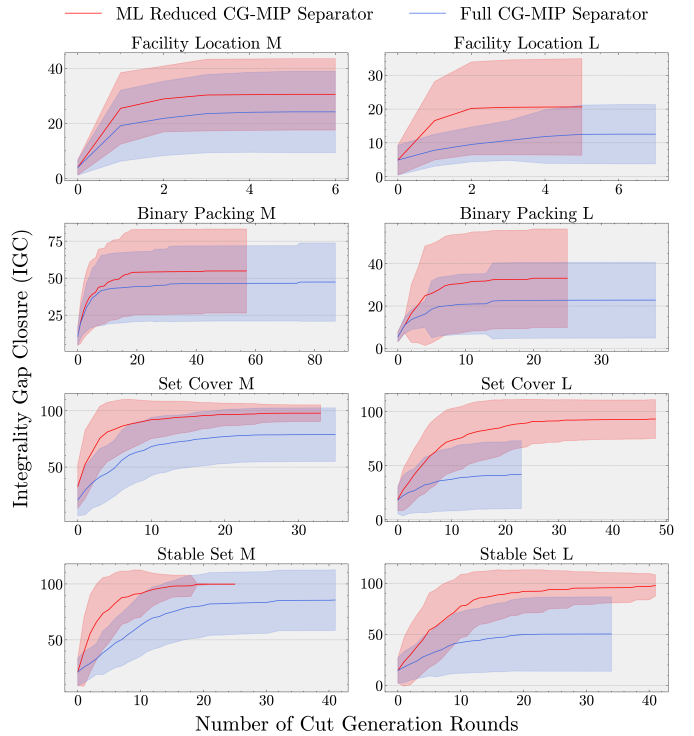


Figure 3: Mean IGC and standard deviation (shaded) v.s. round for medium/large test instances. The reduced and full CG-MIP separators are shown in red and blue, respectively.

efit of the reduced separator over the full separator in terms of final IGC, but also the superiority of optimization-based cut generation over heuristic cut generation. The green curve represents the IGC obtained by adding one round of GMI cuts to the root LP relaxation at “round zero” of the cutting plane method, which all three methods use, and hence lower bounds the other curves.

Is there value in learning only from tight cuts? As described in Section 4, the main variant of our method uses only tight (binding) cuts from CG-MIP’s solution pool for labeling constraints for model training. As an ablation analysis, we compare classification models trained under both the aforementioned setting and the “full” setting in which all cuts are used. Table 3 and Fig. 5 summarize the results. In both settings, the reduced CG-MIP separators dominate the full CG-MIP separator, as expected. However, we observe slightly better performance when learning from tight cuts. For example, in Fig. 5, the green curves (tight cuts) consistently dominate both blue (all cuts) and red (baseline CG-MIP) curves. We believe this is due to judiciously ignoring some less-useful constraints and focusing on fewer, more impactful ones. This is also supported by the larger size reduction of CG-MIP and consequentially lower ratios of cut generation time in Table 3.

Can we learn on small instances and generalize to larger ones? Table 4 shows how models trained on small instances fare on medium and large instances. In most cases,

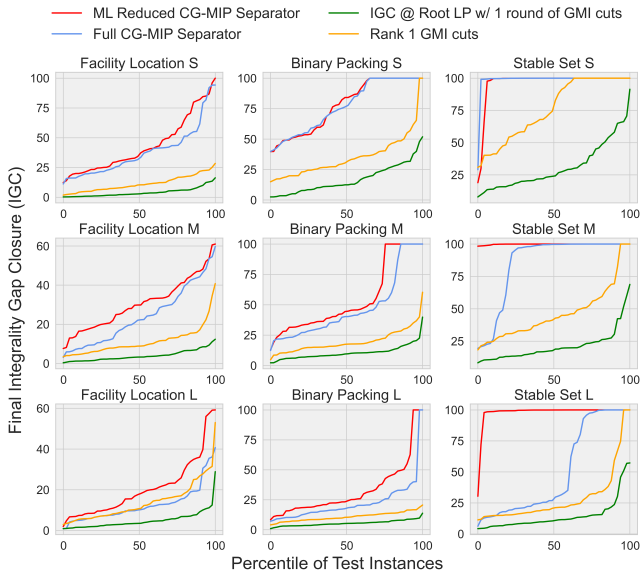


Figure 4: Plot of final test IGC v.s. percentile of instances. The reduced separator is shown in red, the full separator in blue, GMI cuts in yellow, and 1 round of GMI cuts in green. A larger area under a curve is preferred. Due to space limits, plots for SC and PM datasets are in Appendix Fig. 7.

Dataset	Size	Cut Pool	Ratio		% Size Reduction
			IGC \uparrow	Time \downarrow	
FL	M	All	1.54	0.84	96.03
		Tight	1.76	0.76	96.42
	L	All	1.95	0.62	96.59
		Tight	2.04	0.56	97.43
SS	M	All	1.21	0.46	61.29
		Tight	1.3	0.41	68.84
	L	All	1.42	0.79	65.12
		Tight	2.52	1.39	71.17

Table 3: Ratio of mean test final IGC and time as cut pool being considered for classification labels varies.

we observe a substantial size reduction of CG-MIP while maintaining or improving the final IGC in significantly less time. Figure 8 in the appendix illustrates the logarithmic total time comparison via boxplots, revealing that the reduced separator generally achieves a lower total time distribution in most datasets, highlighting the computational efficiency of the reduced CG-MIP separator.

7 Discussion and Conclusion

This paper presents a framework for integrating ML into an established MILP framework for optimization-based separation of aggregation-based cuts. We reduce the computational burden of exact separators by identifying a small important subset of constraints to consider for aggregation. Although we focus on CG cuts, our framework can be applied to other families of aggregation-based cuts. The framework relies on

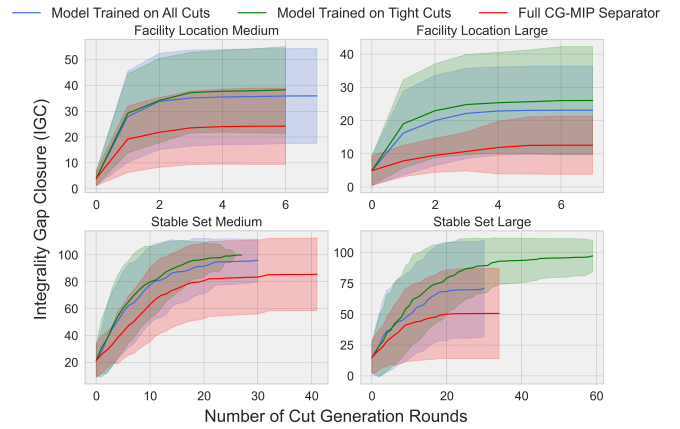


Figure 5: Plot of mean test IGC v.s. cut generation round for full and reduced separators with two labeling strategies.

Dataset	Size Test	Ratio		Abs. Perf Win / Loss	% Size Reduction
		IGC \uparrow	Time \downarrow		
SS	M	1.28	0.36	74 / 0	73.04
	L	2.65	1.3	34 / 4	77.57
PM	M	1.0	0.68	68 / 0	44.16
	L	1.0	0.81	60 / 4	31.74
SC	M	1.03	0.16	76 / 0	77.61
	L	1.14	0.1	30 / 4	87.42
BP	M	1.04	0.36	66 / 4	97.92
	L	1.18	0.48	60 / 10	98.00
FL	M	1.44	0.7	34 / 2	94.45
	L	1.92	0.51	54 / 0	98.03

Table 4: End-to-end metrics of reduced separator using GNN trained on small instances evaluated on larger instances.

the use of graph neural networks trained on a supervised binary classification task to identify constraints that lead to the generation of strong CG cuts. Extensive computational experiments demonstrate that the trained models can effectively reduce the size of the separation problem with little effect on cut quality and smaller runtimes across various MILPs. Models generalize well to unseen data and scale effectively to larger problem sizes.

Although our approach demonstrates faster solution times and improved final IGC, several limitations and future directions remain. First, our evaluation focuses on a pure cutting plane method, not integrated within a MILP solver, which poses challenges due to stricter time constraints for cut generation. Second, the use of supervised learning for classification could be enhanced by exploring reinforcement or imitation learning, which may better capture the iterative nature of the cutting plane method. Finally, to further reduce CG-MIP’s size, we can learn to predict upper bounds on u through regression rather than binary classification. This could further reduce CG-MIP’s search space, as initial findings indicate that many coefficient values are often very small nonnegative numbers.

References

- Achterberg, T. 2007. Constraint integer programming. PhD thesis. Technical report, TU Berlin.
- Agrawal, A.; Amos, B.; Barratt, S.; Boyd, S.; Diamond, S.; and Kolter, Z. 2019. Differentiable Convex Optimization Layers. In *Advances in Neural Information Processing Systems*.
- Balas, E.; and Saxena, A. 2008. Optimizing over the split closure. *Mathematical Programming*, 113(2): 219–240.
- 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.
- Berthold, T.; Francobaldi, M.; and Hendel, G. 2022. Learning to use local cuts. *arXiv preprint arXiv:2206.11618*.
- Bestuzheva, K.; Besançon, M.; Chen, W.-K.; Chmiela, A.; Donkiewicz, T.; van Doornmalen, J.; Eifler, L.; Gaul, O.; Gamrath, G.; Gleixner, A.; et al. 2021. The SCIP optimization suite 8.0. *arXiv preprint arXiv:2112.08872*.
- Bonami, P.; Cornuéjols, G.; Dash, S.; Fischetti, M.; and Lodi, A. 2008. Projected Chvátal–Gomory cuts for mixed integer linear programs. *Mathematical Programming*, 113(2): 241–257.
- Cappart, Q.; Chételat, D.; Khalil, E. B.; Lodi, A.; Morris, C.; and Veličković, P. 2023. Combinatorial optimization and reasoning with graph neural networks. *Journal of Machine Learning Research*, 24(130): 1–61.
- Chételat, D.; and Lodi, A. 2023. Continuous cutting plane algorithms in integer programming. *Operations Research Letters*, 51(4): 439–445.
- Chmiela, A.; Khalil, E.; Gleixner, A.; Lodi, A.; and Pokutta, S. 2021. Learning to schedule heuristics in branch and bound. *Advances in Neural Information Processing Systems*, 34: 24235–24246.
- Chvátal, V. 1973. Edmonds Polytopes and a Hierarchy of Combinatorial Problems. *Discrete Mathematics*, 4: 305–337.
- Deza, A.; and Khalil, E. B. 2023. Machine Learning for Cutting Planes in Integer Programming: A Survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-2023*. International Joint Conferences on Artificial Intelligence Organization.
- Deza, A.; Khalil, E. B.; Fan, Z.; Zhou, Z.; and Zhang, Y. 2024. Learn2Aggregate: Supervised Generation of Chvátal–Gomory Cuts Using Graph Neural Networks. *arXiv preprint arXiv:2409.06559*.
- Dragotto, G.; Clarke, S.; Fisac, J. F.; and Stellato, B. 2023. Differentiable Cutting-plane Layers for Mixed-integer Linear Optimization. *arXiv:2311.03350*.
- Fey, M.; and Lenssen, J. E. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Fischetti, M.; and Lodi, A. 2005. Optimizing over the first Chvátal closure. *Mathematical Programming*, 110: 3–20.
- Gasse, M.; Chételat, D.; Ferroni, N.; Charlin, L.; and Lodi, A. 2019. Exact combinatorial optimization with graph convolutional neural networks. *Advances in neural information processing systems*, 32.
- Khalil, E.; Le Bodic, P.; Song, L.; Nemhauser, G.; and Dilkina, B. 2016. Learning to branch in mixed integer programming. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Khalil, E. B.; Morris, C.; and Lodi, A. 2022. Mip-gnn: A data-driven framework for guiding combinatorial solvers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 10219–10227.
- Koster, A. M.; Zymolka, A.; and Kutschka, M. 2009. Algorithms to separate Chvátal–Gomory cuts. *Algorithmica*, 55(2): 375–391.
- Land, A.; and Doig, A. 1960. An Automatic Method of Solving Discrete Programming Problems. *Econometrica*, 28(3): 497–520.
- Nair, V.; Bartunov, S.; Gimeno, F.; von Glehn, I.; Lichocki, P.; Lobov, I.; O’Donoghue, B.; Sonnerat, N.; Tjandraatmadja, C.; Wang, P.; Addanki, R.; Hapuarachchi, T.; Keck, T.; Keeling, J.; Kohli, P.; Ktena, I.; Li, Y.; Vinyals, O.; and Zwols, Y. 2021. Solving Mixed Integer Programs Using Neural Networks. *arXiv:2012.13349*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv:1912.01703*.
- Paulus, M. B.; Zarpellon, G.; Krause, A.; Charlin, L.; and Maddison, C. 2022. Learning to cut by looking ahead: Cutting plane selection via imitation learning. In *International conference on machine learning*, 17584–17600. PMLR.
- Prouvost, A.; Dumouchelle, J.; Scavuzzo, L.; Gasse, M.; Chételat, D.; and Lodi, A. 2020. Ecole: A Gym-like Library for Machine Learning in Combinatorial Optimization Solvers. *arXiv:2011.06069*.
- Puigdemont, P.; Skoulakis, S.; Chrysos, G.; and Cevher, V. 2024. Learning to Remove Cuts in Integer Linear Programming. *arXiv:2406.18781*.
- Tang, Y.; Agrawal, S.; and Faenza, Y. 2020. Reinforcement learning for integer programming: Learning to cut. In *International conference on machine learning*, 9367–9376. PMLR.
- Turner, M.; Koch, T.; Serrano, F.; and Winkler, M. 2022. Adaptive Cut Selection in Mixed-Integer Linear Programming. *arXiv preprint arXiv:2202.10962*.
- Xavier, A. S.; Qiu, F.; Gu, X.; Becu, B.; and Dey, S. S. 2020. MIPLearn: an extensible framework for learning-enhanced optimization. URL <https://anl-ceeesa.github.io/MIPLearn>.