

# Learning Pseudo-Backdoors for Mixed Integer Programs

Aaron Ferber,<sup>1</sup> Jialin Song,<sup>2</sup> Bistra Dilkina,<sup>1</sup> Yisong Yue<sup>2</sup>

<sup>1</sup> University of Southern California

<sup>2</sup> California Institute of Technology

afferber@usc.edu, jssong@caltech.edu, bdilkina@usc.edu, yyue@caltech.edu

## Abstract

We propose a machine learning approach for quickly solving Mixed Integer Programs (MIP) by learning to prioritize a set of decision variables, which we call pseudo-backdoors, for branching that results in faster solution times. Learning-based approaches have seen success in the area of solving combinatorial optimization problems by being able to flexibly leverage common structures in a given distribution of problems. Our approach takes inspiration from the concept of strong backdoors, which corresponds to a small set of variables such that only branching on these variables yields an optimal integral solution and a proof of optimality. Our notion of pseudo-backdoors corresponds to a small set of variables such that only branching on them leads to faster solve time (which can be solver dependent). A key advantage of pseudo-backdoors over strong backdoors is that they are much amenable to data-driven identification or prediction. Our proposed method learns to estimate the solver performance of a proposed pseudo-backdoor, using a labeled dataset collected on a set of training MIP instances. This model can then be used to identify high-quality pseudo-backdoors on new MIP instances from the same distribution. We evaluate our method on the generalized independent set problems and find that our approach can efficiently identify high-quality pseudo-backdoors. In addition, we compare our learned approach against Gurobi, a state-of-the-art MIP solver, demonstrating that our method can be used to improve solver performance.

## Introduction

Mixed integer programs (MIPs) are widely used mathematical models for combinatorial optimization problems (Conforti et al. 2014) that are generally solved to optimality with a tree search algorithm called branch-and-bound (Land and Doig 2010). Backdoors, initially introduced for SAT (Williams, Gomes, and Selman 2003; Dilkina, Gomes, and Sabharwal 2009) and then generalized to MIPs (Dilkina et al. 2009), are defined as subsets of integer variables such that only branching on them yields an optimal integral solution and a certificate of optimality. While their contribution to SAT has theoretical and practical limitations (Järvisalo and Junttila 2007; Järvisalo and Niemelä 2008; Semenov et al. 2018), solve time speedup in MIPs

has been observed by prioritizing “backdoors” in branching (Fischetti and Monaci 2011). We consider a subset of integer variables a *pseudo-backdoor* if prioritizing those variables leads to faster MIP solve time.

We introduce a data-driven approach to predicting pseudo-backdoors for distributions of MIPs with a *scoring model* to identify high-quality pseudo-backdoors among a sample of candidates, and a subsequent *classification model* to decide whether to use a candidate pseudo-backdoor by setting branching priority accordingly or just use the default solver. We conduct empirical evaluations on the Generalized Independent Set Problem (GISP) (Colombi, Mansini, and Savelsbergh 2017) and show that our models achieve faster solve times than Gurobi.

## Learning Pseudo-Backdoors for MIP

Our goal in finding pseudo-backdoors is to quickly solve MIPs. In a MIP we are asked to find real-valued settings for  $n$  decision variables  $x \in \mathbb{R}^n$ , which maximize a linear objective function  $c^T x$ , subject to  $m$  linear constraints  $Ax \leq b$ , and with a subset  $\mathcal{I} \subseteq [n]$  of the decision variables required to be integral  $x_i \in \mathbb{Z} \forall i \in \mathcal{I}$ . The problem can be written as  $\max_x \{c^T x : Ax \leq b, x_i \in \mathbb{Z} \forall i \in \mathcal{I}\}$ .

Given a MIP, specified by  $P = (c, A, b, \mathcal{I})$ , we want to find a pseudo-backdoor subset  $\mathcal{B} \subseteq \mathcal{I}$ , of the integral decision variables such that prioritizing branching on these decision variables yields fast solve times. We consider a distributional setting of MIP solving where we train a model on several MIP instances from a distribution and deploy it on unseen instances from the same distribution.

Our approach uses two learned models: one that scores subsets of variables according to performance as pseudo-backdoors, and another that classifies whether to use a given pseudo-backdoor over a standard solver. The score model  $S(P, \mathcal{B}; \theta_S)$  is parametrized by neural network parameters  $\theta_S$  which takes as input the MIP specification  $P$ , and a candidate subset  $\mathcal{B}$ , then predicts a score that characterizes if  $\mathcal{B}$  is a good pseudo-backdoor. The classifier  $C(P, \mathcal{B}; \theta_C)$  is parametrized by neural network parameters  $\theta_C$  and predicts whether prioritizing  $\mathcal{B}$  in branching would produce a smaller runtime than a standard solver. Our models use Graph Attention Network (Veličković et al. 2018) to perform message passing followed by global attention pooling (Li et al. 2016) to represent the graph as a single feature vector, which is

dataset	solver	mean	stdev	25 pct	median	75 pct	win / tie / loss vs grb
gisp easy	grb	611	<b>182</b>	488	580	681	0 / 100 / 0
gisp easy	scorer	960	755	515	649	915	41 / 0 / 59
gisp easy	scorer + cls	<b>601</b>	247	<b>481</b>	<b>568</b>	<b>663</b>	24 / 70 / 6
gisp hard	grb	2533	939	1840	2521	2976	0 / 100 / 0
gisp hard	scorer	2373	<b>855</b>	1721	2262	2926	47 / 0 / 53
gisp hard	scorer + cls	<b>2326</b>	<b>855</b>	<b>1654</b>	<b>2215</b>	<b>2866</b>	47 / 27 / 26

Table 1: Runtime comparison in seconds of standard gurobi (grb), the score model (scorer), and the score model with subsequent classification (scorer+cls) on the test set for 2 hardness settings of gisp.

then fed into a fully connected network to produce scalar predictions. The MIP is represented as a bipartite graph as in (Gasse et al. 2019) for permutation invariance and parameter sharing. Each MIP has two sets of nodes, one with nodes representing variables and another representing constraints. Variable nodes have features of the variable’s objective coefficient, and root LP status. Constraint nodes have features of the right hand side constant  $b_j$ , root LP dual variables, and sense ( $\leq$ ,  $\geq$  or  $=$ ). There is an edge between a variable  $i$  and a constraint  $j$  with attribute  $A_{ij}$  if variable  $i$  appears in constraint  $j$ . The candidate pseudo-backdoor set  $\mathcal{B}$  is represented as an additional binary feature for each variable node with value 1 if the variable is in  $\mathcal{B}$  and 0 otherwise. Encoding the input  $(P, \mathcal{B})$  as a graph allows us to leverage state-of-the-art techniques in making predictions on graphs.

We train the score model  $S$  by learning to score subsets of integer variables based on their quality (i.e. runtime) as pseudo-backdoors. The training data contains multiple pseudo-backdoor candidates for each train MIP instance. For a MIP  $P$  and two candidate subsets of integer variables  $\mathcal{B}_1, \mathcal{B}_2$  of  $P$ , we compute the marginal ranking loss (Tsochantaridis et al. 2005)  $\text{loss}(s_1, s_2, y) = \max(0, -y(s_1 - s_2) + m)$  for a given margin value  $m$ , where  $s_1 = S(P, \mathcal{B}_1; \theta_S)$ ,  $s_2 = S(P, \mathcal{B}_2; \theta_S)$  are the scalar score estimates and  $y$  is the ranking label ( $-1$  if  $\mathcal{B}_1$  leads to a smaller runtime than  $\mathcal{B}_2$ ,  $1$  otherwise). Using the ranking loss trains the model to focus on distinguishing between relative performance per-instance rather than accurately modeling the absolute performance. At test time, we sample pseudo-backdoor candidate sets, score them using our model, and use the one with the best score to set branching priorities.

We learn a subsequent classifier module to determine whether to use the candidate subset or the default MIP solver. The classifier has the same architecture as the scoring model, taking as input the bipartite graph representation of the MIP  $P$  and a candidate subset  $\mathcal{B}$ , and outputting a scalar logit for binary classification. The training data contains the most promising pseudo-backdoor for each MIP instance, with a binary label indicating whether the pseudo-backdoor solved the MIP faster than the standard solver.

## Experiment Results

We evaluate our method’s two components on MIP instances sampled from two hardness settings of the Generalized Independent Set Problem (GISP) (Hochbaum and Pathria 1997).

Each hardness setting has 3 sets of 100 MIPs each,  $\mathcal{D}_s$  for fitting the score model,  $\mathcal{D}_c$  for the classification model, and  $\mathcal{D}_t$  for testing results. For each of the 300 MIPs, we sample 50 candidate pseudo-backdoors containing 1% of the integer variables in a given MIP, with variable selection probability proportional to the variables’ root LP fractionality as in (Dilkina et al. 2009).

Table 1 shows that the score model alone performs well on many instances. On the hard instances, scorer performs well, having 6% faster runtimes than gurobi on average, and winning in 47 instances. However, scorer has 57% slower runtimes on the easy instances, but still outperformed gurobi on 41 instances, demonstrating that while it has poorer performance on average, it has potential for yielding fast solve times on many instances. Additionally, we can see that the score model alone has much higher variability than Gurobi on the easy instances. This undesired property further motivates the inclusion of the classifier model to improve the overall performance. The scorer + cls model outperforms Gurobi across both MIP distributions in terms of the time distribution for MIP solving. It performs well across distributions of instances, having the lowest solve times on average and at different quantiles. The score + cls pipeline outperforms gurobi by 2%, 8% on average for easy, and hard instances respectively. In terms of win / tie / loss, the scorer’s losses against gurobi are reduced by the classifier while its wins are retained. Furthermore, the variability on the gisp easy instances is reduced to be on par with Gurobi.

## Conclusion

In this extended abstract, we have presented a flexible method for learning and using pseudo-backdoors to improve MIP solving. Inspired by previous work on backdoors in MIPs, our method is comprised of two parts: an initial ranking model which scores candidate pseudo-backdoors according to how fast they will solve a given MIP, and a subsequent classification module to determine whether the selected pseudo-backdoor will indeed result in faster runtime compared with a base solver. Across varying difficulty levels of GISP distributions, our method finds high-quality pseudo-backdoors on unseen MIP instances. Furthermore, we find that the combination of the score model and the classification model yields consistently faster solution times compared to a state-of-the-art solver.

## References

- Colombi, M.; Mansini, R.; and Savelsbergh, M. 2017. The generalized independent set problem: Polyhedral analysis and solution approaches. *European Journal of Operational Research* 260(1): 41–55.
- Conforti, M.; Cornuéjols, G.; Zambelli, G.; et al. 2014. *Integer programming*, volume 271. Springer.
- Dilkina, B.; Gomes, C. P.; Malitsky, Y.; Sabharwal, A.; and Sellmann, M. 2009. Backdoors to combinatorial optimization: Feasibility and optimality. In *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, 56–70. Springer.
- Dilkina, B.; Gomes, C. P.; and Sabharwal, A. 2009. Backdoors in the context of learning. In *International Conference on Theory and Applications of Satisfiability Testing*, 73–79. Springer.
- Fischetti, M.; and Monaci, M. 2011. Backdoor branching. In *International Conference on Integer Programming and Combinatorial Optimization*, 183–191. Springer.
- Gasse, M.; Chételat, D.; Ferroni, N.; Charlin, L.; and Lodi, A. 2019. Exact Combinatorial Optimization with Graph Convolutional Neural Networks. In *Advances in Neural Information Processing Systems* 32.
- Hochbaum, D. S.; and Pathria, A. 1997. Forest harvesting and minimum cuts: a new approach to handling spatial constraints. *Forest Science* 43(4): 544–554.
- Järvisalo, M.; and Junttila, T. 2007. Limitations of restricted branching in clause learning. In *International Conference on Principles and Practice of Constraint Programming*, 348–363. Springer.
- Järvisalo, M.; and Niemelä, I. 2008. The effect of structural branching on the efficiency of clause learning SAT solving: An experimental study. *Journal of Algorithms* 63(1-3): 90–113.
- Land, A. H.; and Doig, A. G. 2010. An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, 105–132. Springer.
- Li, Y.; Zemel, R.; Brockschmidt, M.; and Tarlow, D. 2016. Gated Graph Sequence Neural Networks. In *Proceedings of ICLR'16*.
- Semenov, A.; Zaikin, O.; Otpuschennikov, I.; Kochemazov, S.; and Ignatiev, A. 2018. On cryptographic attacks using backdoors for SAT. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Tsochantaridis, I.; Joachims, T.; Hofmann, T.; Altun, Y.; and Singer, Y. 2005. Large margin methods for structured and interdependent output variables. *Journal of machine learning research* 6(9).
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. *International Conference on Learning Representations* URL <https://openreview.net/forum?id=rJXMpikCZ>. Accepted as poster.
- Williams, R.; Gomes, C. P.; and Selman, B. 2003. Backdoors to typical case complexity. In *Proceedings of the 18th international joint conference on Artificial intelligence*, 1173–1178.