

Lookback for Learning to Branch

Prateek Gupta*, Elias B. Khalil, Didier Chételat, Maxime Gasse,
M. Pawan Kumar, Andrea Lodi, Yoshua Bengio

SIAM Conference on Optimization, June 2nd 2023



To **improve** the extent to which neural networks can **imitate** a computationally expensive but accurate heuristic to solve mixed-integer linear programming (MILP) problems.

Outline

Problem formulation

Our solution

Conclusion

Outline

Problem formulation

- Discrete Optimization
- Branch-and-Bound
- The Branching Problem
- Learning to branch
- Lookback property

Our solution

Conclusion

Outline

Problem formulation

Discrete Optimization

Branch-and-Bound

The Branching Problem

Learning to branch

Lookback property

Our solution

Conclusion

Mixed-Integer Linear Program (MILP)

$$\arg \min_x \quad c^\top x$$

- ▶ $c \in \mathbb{R}^n$ the objective coefficients

Mixed-Integer Linear Program (MILP)

$$\begin{array}{ll} \arg \min_{\mathbf{x}} & \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \mathbf{Ax} \leq \mathbf{b}, \end{array}$$

- ▶ $\mathbf{c} \in \mathbb{R}^n$ the objective coefficients
- ▶ $\mathbf{A} \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $\mathbf{b} \in \mathbb{R}^m$ the constraint right-hand-sides

Mixed-Integer Linear Program (MILP)

$$\begin{array}{ll}\arg \min_{\mathbf{x}} & \mathbf{c}^T \mathbf{x} \\ \text{subject to} & \mathbf{A}\mathbf{x} \leq \mathbf{b}, \\ & \mathbf{l} \leq \mathbf{x} \leq \mathbf{u},\end{array}$$

- ▶ $\mathbf{c} \in \mathbb{R}^n$ the objective coefficients
- ▶ $\mathbf{A} \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $\mathbf{b} \in \mathbb{R}^m$ the constraint right-hand-sides
- ▶ $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ the lower and upper variable bounds

Mixed-Integer Linear Program (MILP)

$$\begin{aligned} \arg \min_x \quad & c^T x \\ \text{subject to} \quad & Ax \leq b, \\ & l \leq x \leq u, \\ & x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{aligned}$$

- ▶ $c \in \mathbb{R}^n$ the objective coefficients
- ▶ $A \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $b \in \mathbb{R}^m$ the constraint right-hand-sides
- ▶ $l, u \in \mathbb{R}^n$ the lower and upper variable bounds
- ▶ $p \leq n$ integer variables

Mixed-Integer Linear Program (MILP)

$$\begin{aligned} \arg \min_{\mathbf{x}} \quad & \mathbf{c}^T \mathbf{x} \\ \text{subject to} \quad & \mathbf{Ax} \leq \mathbf{b}, \\ & \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \\ & \mathbf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{aligned}$$

- ▶ $\mathbf{c} \in \mathbb{R}^n$ the objective coefficients
- ▶ $\mathbf{A} \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $\mathbf{b} \in \mathbb{R}^m$ the constraint right-hand-sides
- ▶ $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ the lower and upper variable bounds
- ▶ $p \leq n$ integer variables

NP-hard problem.

Applications


Combinatorial Auctions

Facility location-Allocation

Maximum Independent Set

Set Covering

and many more ...



Mixed-Integer Linear Program (MILP)

$$\begin{aligned} \arg \min_x \quad & c^\top x \\ \text{subject to} \quad & Ax \leq b, \\ & l \leq x \leq u, \\ & \boxed{x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}}. \end{aligned}$$

- ▶ $c \in \mathbb{R}^n$ the objective coefficients
- ▶ $A \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $b \in \mathbb{R}^m$ the constraint right-hand-sides
- ▶ $l, u \in \mathbb{R}^n$ the lower and upper variable bounds
- ▶ $p \leq n$ integer variables

NP-hard problem.

Mixed-Integer Linear Program (MILP)




Image credit: Maxime Gasse

Linear Program (LP)

$$\begin{aligned} \arg \min_x \quad & c^\top x \\ \text{subject to} \quad & Ax \leq b, \\ & l \leq x \leq u, \\ & \boxed{x \in \mathbb{R}^n}. \end{aligned}$$

- ▶ $c \in \mathbb{R}^n$ the objective coefficients
- ▶ $A \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $b \in \mathbb{R}^m$ the constraint right-hand-sides
- ▶ $l, u \in \mathbb{R}^n$ the lower and upper variable bounds

Linear Program (LP)

$$\begin{aligned} \arg \min_x \quad & c^\top x \\ \text{subject to} \quad & Ax \leq b, \\ & l \leq x \leq u, \\ & \boxed{x \in \mathbb{R}^n}. \end{aligned}$$


- ▶ $c \in \mathbb{R}^n$ the objective coefficients
- ▶ $A \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
- ▶ $b \in \mathbb{R}^m$ the constraint right-hand-sides
- ▶ $l, u \in \mathbb{R}^n$ the lower and upper variable bounds
- Polynomially solvable

Linear Program (LP)

$$\begin{aligned} & \arg \min_x && c^\top x \\ & \text{subject to} && Ax \leq b, \\ & && l \leq x \leq u, \\ & && \boxed{x \in \mathbb{R}^n}. \end{aligned}$$

- ▶ $c \in \mathbb{R}^n$ the objective coefficients
 - ▶ $A \in \mathbb{R}^{m \times n}$ the constraint coefficient matrix
 - ▶ $b \in \mathbb{R}^m$ the constraint right-hand-sides
 - ▶ $l, u \in \mathbb{R}^n$ the lower and upper variable bounds
-
- Polynomially solvable
 - Yields lower bounds to the original MILP

LP Relaxation of a MILP



Outline

Problem formulation

Discrete Optimization

Branch-and-Bound

The Branching Problem

Learning to branch

Lookback property

Our solution

Conclusion

Branch-and-Bound (B&B)

B&B (Land et al., 1960) is the widely used framework to solve MILPs.

It consists of two steps



Each node in branch-and-bound is a new MIP

Image source: <https://www.gurobi.com/resource/mip-basics/>

Branch-and-Bound (B&B)

B&B (Land et al., 1960) is the widely used framework to solve MILPs.

It consists of two steps

- **Branching** - Select variable to split the problem into two



Each node in branch-and-bound is a new MIP

Branch-and-Bound (B&B)

B&B (Land et al., 1960) is the widely used framework to solve MILPs.

It consists of two steps

- ▶ **Branching** - Select variable to split the problem into two
- ▶ **Bounding** - Solve the LP relaxation of resulting problem to obtain optimization guarantees on the solution



Each node in branch-and-bound is a new MIP

Branch-and-Bound

Branch: Split the LP recursively over a non-integral variable, i.e.

$$\exists i \leq p \mid x_i^* \notin \mathbb{Z}$$

$$x_i \leq \lfloor x_i^* \rfloor \quad \vee \quad x_i \geq \lceil x_i^* \rceil.$$

Branch-and-Bound

Branch: Split the LP recursively over a non-integral variable, i.e.

$$\exists i \leq p \mid x_i^* \notin \mathbb{Z}$$

$$x_i \leq \lfloor x_i^* \rfloor \quad \vee \quad x_i \geq \lceil x_i^* \rceil.$$

Lower bound (L): minimal among leaf nodes.

Upper bound (U): minimal among leaf nodes with integral solution.

Branch-and-Bound

Branch: Split the LP recursively over a non-integral variable, i.e.

$$\exists i \leq p \mid x_i^* \notin \mathbb{Z}$$

$$x_i \leq \lfloor x_i^* \rfloor \quad \vee \quad x_i \geq \lceil x_i^* \rceil.$$

Lower bound (L): minimal among leaf nodes.

Upper bound (U): minimal among leaf nodes with integral solution.

Stopping criterion:

- ▶ **L** = **U** (optimality certificate)
- ▶ **L** = ∞ (infeasibility certificate)
- ▶ **L** - **U** < threshold (early stopping)

Branch-and-Bound

Branch: Split the LP recursively over a non-integral variable, i.e.

$$\exists i \leq p \mid x_i^* \notin \mathbb{Z}$$

$$x_i \leq \lfloor x_i^* \rfloor \quad \vee \quad x_i \geq \lceil x_i^* \rceil.$$

Lower bound (L): minimal among leaf nodes.

Upper bound (U): minimal among leaf nodes with integral solution.

Stopping criterion:


- ▶ **L** = **U** (optimality certificate)
- ▶ **L** = ∞ (infeasibility certificate)
- ▶ **L** - **U** < threshold (early stopping)

Note: A time limit is used to ensure termination.

Branch-and-bound: a sequential process

Sequential decisions:


- ▶ variable selection (branching)
- ▶ node selection
- ▶ *cutting plane selection*
- ▶ *primal heuristic selection*
- ▶ *simplex initialization*
- ▶ ...



Branch-and-bound: a sequential process

Sequential decisions:

- ▶ variable selection (branching)
- ▶ node selection
- ▶ *cutting plane selection*
- ▶ *primal heuristic selection*
- ▶ *simplex initialization*
- ▶ ...



Outline

Problem formulation

Discrete Optimization

Branch-and-Bound

The Branching Problem

Learning to branch

Lookback property

Our solution

Conclusion

Branching Policy

It is also called as **variable selection policy**.

Policy Objective: Given a B&B node i.e. MILP, select a variable $i \leq p \mid x_i^* \notin \mathbb{Z}$ so that the final size of the tree is minimum (a proxy for running time).

A gold standard: Strong Branching (impractical)

Strong branching¹: one-step forward looking (greedy)

- ▶ solve both LPs **for each candidate variable**
- ▶ select the variable resulting in tightest relaxation
- + small trees
- computationally expensive

¹D. Applegate et al. (1995). Finding cuts in the TSP. [Tech. rep. DIMACS](#); J. Linderoth et al. (May 1999). A Computational Study of Search Strategies for Mixed Integer Programming.

A gold standard: Strong Branching (impractical)

Strong branching score for a variable i at a node n

A gold standard: Strong Branching (impractical)

Strong branching score for a variable i at a node n

- ▶ Let L be the value of LP relaxation of the MILP

A gold standard: Strong Branching (impractical)

Strong branching score for a variable i at a node n

- ▶ Let L be the value of LP relaxation of the MILP
- ▶ Denote L_i^+ as the value of LP relaxation of the MILP *after adding* $x_i \geq \lceil x_i^* \rceil$ *constraint*

A gold standard: Strong Branching (impractical)

Strong branching score for a variable i at a node n

- ▶ Let L be the value of LP relaxation of the MILP
- ▶ Denote L_i^+ as the value of LP relaxation of the MILP *after adding* $x_i \geq \lceil x_i^* \rceil$ *constraint*
- ▶ Similarly, denote L_i^- for the other half

A gold standard: Strong Branching (impractical)

Strong branching score for a variable i at a node n

- ▶ Let L be the value of LP relaxation of the MILP
- ▶ Denote L_i^+ as the value of LP relaxation of the MILP *after adding* $x_i \geq \lceil x_i^* \rceil$ *constraint*
- ▶ Similarly, denote L_i^- for the other half

Strong branching score

$$\text{score}_{SB,i} = \max(L - L_i^+, \epsilon) \times \max(L - L_i^-, \epsilon)$$

A gold standard: Strong Branching (impractical)

Strong branching score for a variable i at a node n

- ▶ Let L be the value of LP relaxation of the MILP
- ▶ Denote L_i^+ as the value of LP relaxation of the MILP *after adding* $x_i \geq \lceil x_i^* \rceil$ *constraint*
- ▶ Similarly, denote L_i^- for the other half

Strong branching score

$$\text{score}_{SB,i} = \max(L - L_i^+, \epsilon) \times \max(L - L_i^-, \epsilon)$$

Strong branching decision

$$i_{SB}^* = \arg \max_i \text{score}_{SB,i}$$

Outline

Problem formulation

Discrete Optimization

Branch-and-Bound

The Branching Problem

Learning to branch

Lookback property

Our solution

Conclusion

Learning to branch

Objective:

Given a distribution of problem sets, find a branching policy that yields a shortest tree on an average.

Exploits statistical correlation across problem sets.




Figure: Application specific distribution

Learning to branch

Objective: Given a dataset of MILPs

- ▶ learn an inexpensive function f
- ▶ that imitates strong branching decisions (computationally expensive)

Learning to branch

Objective: Given a dataset of MILPs

- ▶ learn an inexpensive function f
- ▶ that imitates strong branching decisions (computationally expensive)

$$i_{SB}^* = \arg \max_{i \in \mathcal{C}} \text{score}_{SB,i} \quad i_f^* = \arg \max_{i \in \mathcal{C}} \text{score}_{f_\theta,i},$$

where $s_{f_\theta}^i$ is the score for $i \leq p$ variable as estimated by f_θ .

Learning to branch

Objective: Given a dataset of MILPs

- ▶ learn an inexpensive function f
- ▶ that imitates strong branching decisions (computationally expensive)

$$i_{SB}^* = \arg \max_{i \in \mathcal{C}} \text{score}_{SB,i} \quad i_f^* = \arg \max_{i \in \mathcal{C}} \text{score}_{f_\theta,i},$$

where $s_{f_\theta}^i$ is the score for $i \leq p$ variable as estimated by f_θ .

$$\theta^* = \arg \min_{\theta} \mathcal{L}(f_\theta(MILP), i_{SB}^*)$$

Learning to branch

Objective: Given a dataset of MILPs

- ▶ learn an inexpensive function f
- ▶ that imitates strong branching decisions (computationally expensive)

$$i_{SB}^* = \arg \max_{i \in \mathcal{C}} \text{score}_{SB,i} \quad i_f^* = \arg \max_{i \in \mathcal{C}} \text{score}_{f_\theta,i},$$

where $s_{f_\theta}^i$ is the score for $i \leq p$ variable as estimated by f_θ .

$$\theta^* = \arg \min_{\theta} \mathcal{L}(f_\theta(MILP), i_{SB}^*)$$

Well studied problem (not an exhaustive list)

- ▶ Gasse et al., 2019 \implies offline imitation learning using GCNN
- ▶ Nair et al., 2020 \implies uses GCNNs to design other heuristics
- ▶ Chen et al., 2022 \implies studies the limitations of existing GNNs to represent MILPs

Learning to branch: GNNs

Gasse et al., 2019 uses Graph Neural Networks to imitate the strong branching policy *through classification framework*

- + superior representation power
- + best overall accuracy

Learning to branch: GNNs

Gasse et al., 2019 uses Graph Neural Networks to imitate the strong branching policy *through classification framework*

- + superior representation power
- + best overall accuracy

Model inputs

Inputs to the GNN is a **bipartite-representation of MILP**: G

GNNs: Bipartite Representation of MILPs

Natural representation : variable / constraint bipartite graph

$$\begin{aligned} & \arg \min_x \quad c^\top x \\ & \text{subject to} \quad Ax \leq b, \\ & \quad \quad \quad l \leq x \leq u, \\ & \quad \quad \quad x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{aligned}$$

GNNs: Bipartite Representation of MILPs

Natural representation : variable / constraint bipartite graph

$$\begin{array}{ll} \arg \min_x & c^\top x \\ \text{subject to} & Ax \leq b, \\ & l \leq x \leq u, \\ & x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{array}$$

$$x_0$$

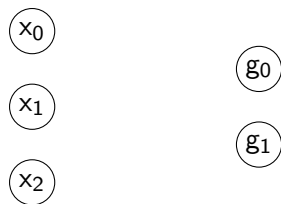
$$x_1$$

$$x_2$$

- x_i : variable features (type, coef., bounds, LP solution...)

GNNs: Bipartite Representation of MILPs

Natural representation : variable / constraint bipartite graph


$$\begin{array}{ll} \arg \min_x & c^\top x \\ \text{subject to} & Ax \leq b, \\ & l \leq x \leq u, \\ & x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{array}$$


- ▶ x_i : variable features (type, coef., bounds, LP solution...)
- ▶ g_j : constraint features (right-hand-side, LP slack...)

GNNs: Bipartite Representation of MILPs

Natural representation : variable / constraint bipartite graph

$$\begin{array}{ll}\arg \min_x & c^\top x \\ \text{subject to} & Ax \leq b, \\ & l \leq x \leq u, \\ & x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}.\end{array}$$



- ▶ x_i : variable features (type, coef., bounds, LP solution...)
- ▶ g_j : constraint features (right-hand-side, LP slack...)
- ▶ $e_{i,j}$: non-zero coefficients in A

Learning to branch: GNNs

Gasse et al., 2019 uses Graph Neural Networks to imitate the strong branching policy *through classification framework*

- + superior representation power
- + best overall accuracy

Learning to branch: GNNs

Gasse et al., 2019 uses Graph Neural Networks to imitate the strong branching policy *through classification framework*

- + superior representation power
- + best overall accuracy
- requires GPUs for best running times (Gupta et al., 2020 addresses this drawback)

Learning to branch: GNNs

Gasse et al., 2019 uses Graph Neural Networks to imitate the strong branching policy *through classification framework*

- + superior representation power
- + best overall accuracy
- requires GPUs for best running times (Gupta et al., 2020 addresses this drawback)
- ? Can we further improve the performance?

Outline

Problem formulation

- Discrete Optimization
- Branch-and-Bound
- The Branching Problem
- Learning to branch
- Lookback property**


Our solution

Conclusion

Lookback condition in strong branching

Strong branching heuristic exhibits the following condition:
Parent's second best choice is *often* the child's best choice.

Frequency of Lookback condition



Frequency of Lookback condition

Instances	Description	number of parent-child pairs collected	number of parent-child pairs exhibiting the lookback property	Frequency of the lookback property
CORLAT	Corridor planning in wildlife management	5082	1765	34.73%
RCW	Red-cockaded woodpecker diffusion conservation	5115	1952	38.16%

Frequency of the lookback property in the real-world instances is as prevalent as in the synthetic instances considered in the main paper. These instances are made available by [Dilkina et al., 2017](#).

Outline

Problem formulation

Our solution

- Loss target

- Regularizer

- Evaluation

Conclusion

Outline

Problem formulation

Our solution

- Loss target

- Regularizer

- Evaluation

Conclusion

Loss targets

We consider two types of targets

(\mathcal{Z} is the set of all the second best branching variables)

Original one-hot encoded target,

y

$$y_i = \begin{cases} 1, & i = i_{SB}^* \\ 0, & \text{otherwise} \end{cases}$$

Loss targets

We consider two types of targets

(\mathcal{Z} is the set of all the second best branching variables)

Original one-hot encoded target,

y

$$y_i = \begin{cases} 1, & i = i_{SB}^* \\ 0, & \text{otherwise} \end{cases}$$

$$\theta_y^* = \arg \min_{\theta} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), y_k)$$

Loss targets

We consider two types of targets

(\mathcal{Z} is the set of all the second best branching variables)

Original one-hot encoded target,
 y

$$y_i = \begin{cases} 1, & i = i_{SB}^* \\ 0, & \text{otherwise} \end{cases}$$

Second-best ϵ -smoothed target,
 z

$$z_i = \begin{cases} 1 - \epsilon, & i = i_{SB}^* \\ \frac{\epsilon}{|\mathcal{Z}|}, & i \in \mathcal{Z} \\ 0, & \text{otherwise} \end{cases}$$

$$\theta_y^* = \arg \min_{\theta} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), y_k)$$

Loss targets

We consider two types of targets

(\mathcal{Z} is the set of all the second best branching variables)

Original one-hot encoded target,

y

$$y_i = \begin{cases} 1, & i = i_{SB}^* \\ 0, & \text{otherwise} \end{cases}$$

Second-best ϵ -smoothed target,

z

$$z_i = \begin{cases} 1 - \epsilon, & i = i_{SB}^* \\ \frac{\epsilon}{|\mathcal{Z}|}, & i \in \mathcal{Z} \\ 0, & \text{otherwise} \end{cases}$$

$$\theta_y^* = \arg \min_{\theta} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), y_k)$$

$$\theta_z^* = \arg \min_{\theta} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), z_k)$$

Outline

Problem formulation

Our solution

Loss target

Regularizer

Evaluation

Conclusion

Parent-As-Target (PAT) regularizer

We consider a regularizer to encourage the lookback property in GNNs

Parent-As-Target (PAT) regularizer

We consider a regularizer to encourage the lookback property in GNNs

$$\text{loss}_{PAT} = 1\{\textit{Lookback}_i\}.$$

Parent-As-Target (PAT) regularizer

We consider a regularizer to encourage the lookback property in GNNs

$$\text{loss}_{PAT} = 1\{Lookback_i\} \cdot CE(f_{\theta}(\mathcal{G}_i), ??),$$

Parent-As-Target (PAT) regularizer

We consider a regularizer to encourage the lookback property in GNNs

$$\text{loss}_{PAT} = 1\{\text{Lookback}_i\} \cdot CE(f_{\theta}(\mathcal{G}_i), f_{\theta}(\mathcal{G}_i^{\text{parent}})[\mathcal{C}_i]),$$

Outline

Problem formulation

Our solution

Loss target

Regularizer

Evaluation

Conclusion

Performance evaluation

We will consider three different set of parameters

- ▶ Choice of the target:
 - ▶ One-hot encoded, y
 - ▶ Second-best ϵ -smoothed, z
- ▶ Strength of the PAT regularizer, $\lambda_{PAT} \in \{0, 0.01, 0.1, 0.2, 0.3\}$
- ▶ Strength of the l_2 -regularizer, $\lambda_{l_2} \in \{0.0, 0.01, 0.1, 1.0\}$

Performance evaluation

$$\theta_y = \arg \min_{\theta, \lambda_{l2}} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), y_k) + \lambda_{l2} \cdot \|\theta\|_2$$

Performance evaluation

$$\theta_y = \arg \min_{\theta, \lambda_{l2}} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), y_k) + \lambda_{l2} \cdot \|\theta\|_2$$

$$\theta_z = \arg \min_{\theta, \lambda_{l2}} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), z_k) + \lambda_{l2} \cdot \|\theta\|_2$$

Performance evaluation

$$\theta_y = \arg \min_{\theta, \lambda_{l2}} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), y_k) + \lambda_{l2} \cdot \|\theta\|_2$$

$$\theta_z = \arg \min_{\theta, \lambda_{l2}} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), z_k) + \lambda_{l2} \cdot \|\theta\|_2$$

$$\theta_{PAT} = \arg \min_{\theta, v, \lambda_{l2}, \lambda_{PAT}} \frac{1}{N} \sum_{k=1}^N CE(f_{\theta}(\mathcal{G}_k), v) + \lambda_{l2} \cdot \|\theta\|_2 + \lambda_{PAT} \cdot \text{loss}_{PAT}$$

Performance evaluation: Instances

- ▶ **Small** instances are used to collect training data of parent-child nodes by solving these instances using the strong branching heuristic as the variable selection policy in the solver


Performance evaluation: Instances

- ▶ **Small** instances are used to collect training data of parent-child nodes by solving these instances using the strong branching heuristic as the variable selection policy in the solver
- ▶ **Medium** instances are used for hyperparameter selection incorporating harder-to-formulate criterion in the objective function

Performance evaluation: Instances

- ▶ **Small** instances are used to collect training data of parent-child nodes by solving these instances using the strong branching heuristic as the variable selection policy in the solver
- ▶ **Medium** instances are used for hyperparameter selection incorporating harder-to-formulate criterion in the objective function
- ▶ **Big** instances are used to report performance evaluation

Model selection criterion: Validation accuracy




Top-1 accuracy (1-standard deviation) on validation dataset.

Model selection criterion: Out-of-distribution performance


We solve 100 **medium** instances and collect the following metrics

- ▶ Wins: Number of times a model solved the instance fastest
- ▶ Time: 1-shifted geometric mean of time taken to solve each instance
- ▶ Nodes: 1-shifted geometric mean of nodes taken in the B&B tree of the *commonly solved instances*

Model selection criterion: Out-of-distribution performance



Model selection criterion: Out-of-distribution performance



We plot the range-normalized (range is specified in parenthesis) Time and Node performance of the selected models. The centered "X" black mark shows the final models that were selected to be used for evaluating the performance on Big instances. The points with a red outline show the performance of the models selected according to the best validation accuracy (Note that we omit such models for indset as it distorts the scale of the plot.)

Final performance

Model	Time	Time (c)	Wins	Solved	Nodes (c)
FSB*	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17 979
TUNEDRPB	644.20	450.06	0	80	18 104
GNN	507.06	333.59	14	80	17 145
GNN-PAT (ours)	477.26	310.22	69	84	16 388

Combinatorial Auction (Bigger)

Figure: Evaluation metrics on Big instances with a time budget of 30 minutes per instance

Final performance

Model	Time	Time (c)	Wins	Solved	Nodes (c)
FSB*	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17 979
TUNEDRPB	644.20	450.06	0	80	18 104
GNN	507.06	333.59	14	80	17 145
GNN-PAT (ours)	477.26	310.22	69	84	16 388

Combinatorial Auction (Bigger)

Figure: Evaluation metrics on Big instances with a time budget of 30 minutes per instance

Final performance

Model	Time	Time (c)	Wins	Solved	Nodes (c)
FSB*	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17 979
TUNEDRPB	644.20	450.06	0	80	18 104
GNN	507.06	333.59	14	80	17 145
GNN-PAT (ours)	477.26	310.22	69	84	16 388

Combinatorial Auction (Bigger)

Figure: Evaluation metrics on Big instances with a time budget of 30 minutes per instance

Final performance

Model	Time	Time (c)	Wins	Solved	Nodes (c)
FSB*	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17 979
TUNEDRPB	644.20	450.06	0	80	18 104
GNN	507.06	333.59	14	80	17 145
GNN-PAT (ours)	477.26	310.22	69	84	16 388

Combinatorial Auction (Bigger)

Figure: Evaluation metrics on Big instances with a time budget of 30 minutes per instance

Final performance

Model	Time	Time (c)	Wins	Solved	Nodes (c)
FSB*	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17 979
TUNEDRPB	644.20	450.06	0	80	18 104
GNN	507.06	333.59	14	80	17 145
GNN-PAT (ours)	477.26	310.22	69	84	16 388

Combinatorial Auction (Bigger)

Figure: Evaluation metrics on Big instances with a time budget of 30 minutes per instance

Final performance

Model	Time	Time (c)	Wins	Solved	Nodes (c)
FSB*	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17 979
TUNEDRPB	644.20	450.06	0	80	18 104
GNN	507.06	333.59	14	80	17 145
GNN-PAT (ours)	477.26	310.22	69	84	16 388

Combinatorial Auction (Bigger)

Figure: Evaluation metrics on Big instances with a time budget of 30 minutes per instance

Optimality gap on commonly unsolved instances




Figure: Mean optimality gap of the *commonly unsolved instances*

Outline

Problem formulation

Our solution

Conclusion

Conclusion

- ▶ We **discover** *lookback* phenomenon in the gold-standard (by tree size) variable-selection heuristic

Conclusion

- ▶ We **discover** *lookback* phenomenon in the gold-standard (by tree size) variable-selection heuristic
- ▶ We **proposed** second-best ϵ -smoothed target and a PAT regularizer term to incorporate lookback phenomenon in deep learning models

Conclusion

- ▶ We **discover** *lookback* phenomenon in the gold-standard (by tree size) variable-selection heuristic
- ▶ We **proposed** second-best ϵ -smoothed target and a PAT regularizer term to incorporate lookback phenomenon in deep learning models
- ▶ We **proposed** a model selection scheme to incorporate final utility of these models in the objective function

Conclusion

- ▶ We **discover** *lookback* phenomenon in the gold-standard (by tree size) variable-selection heuristic
- ▶ We **proposed** second-best ϵ -smoothed target and a PAT regularizer term to incorporate lookback phenomenon in deep learning models
- ▶ We **proposed** a model selection scheme to incorporate final utility of these models in the objective function
- ▶ Our proposed models outperform the SOTA results

Open questions

- ▶ Discovery of more inductive biases

Open questions

- ▶ Discovery of more inductive biases
- ▶ Designing better ways to incorporate lookback property

Open questions

- ▶ Discovery of more inductive biases
- ▶ Designing better ways to incorporate lookback property
- ▶ Improve reinforcement learning solutions using the lookback property

Open questions

- ▶ Discovery of more inductive biases
- ▶ Designing better ways to incorporate lookback property
- ▶ Improve reinforcement learning solutions using the lookback property



Paper: <https://arxiv.org/abs/2006.15212>

Lookback for Learning to Branch

Thank you!

Prateek Gupta*, Elias B. Khalil, Didier Chételat, Maxime Gasse,
M. Pawan Kumar, Andrea Lodi, Yoshua Bengio

