Lookback for Learning to Branch

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

Huawei London Research Center, Jan. 10th 2023



The Alan Turing Institute













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

Problem formulation

Our solution

Problem formulation

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

Our solution

Problem formulation

Discrete Optimization

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

Our solution

$$\underset{x}{\operatorname{arg\,min}} \quad \mathbf{c}^{\top}\mathbf{x}$$

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

$$\begin{array}{ll} \arg\min_{x} & \mathsf{c}^{\top}\mathsf{x}\\ \text{subject to} & \mathsf{A}\mathsf{x} \leq \mathsf{b}, \end{array}$$

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

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

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & c^{\top}x\\ \text{subject to} & Ax \leq b,\\ & I \leq x \leq u, \end{array}$$

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

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & \mathsf{c}^\top \mathsf{x} \\ \text{subject to} & \mathsf{A}\mathsf{x} \leq \mathsf{b}, \\ & \mathsf{I} \leq \mathsf{x} \leq \mathsf{u}, \\ & \mathsf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{array}$$

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

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & \mathsf{c}^\top \mathsf{x} \\ \text{subject to} & \mathsf{A}\mathsf{x} \leq \mathsf{b}, \\ & \mathsf{I} \leq \mathsf{x} \leq \mathsf{u}, \\ & \mathsf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{array}$$

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

NP-hard problem.

Applications

Combinatorial Auctions

Facility location-Allocation

Maximum Indendent Set

Set Covering

and many more ...





$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & c^{\top}x\\ \text{subject to} & Ax \leq b,\\ & I \leq x \leq u,\\ & \hline x \in \mathbb{Z}^{p} \times \mathbb{R}^{n-p} \end{array}$$

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

NP-hard problem.



Image credit: Maxime Gasse

Linear Program (LP)

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & c^{\top}x\\ \text{subject to} & Ax \leq b,\\ & \mathsf{I} \leq x \leq \mathsf{u},\\ & \overline{x \in \mathbb{R}^n} \end{array}$$

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

Linear Program (LP)

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & c^{\top}x\\ \text{subject to} & Ax \leq b,\\ & \mathsf{I} \leq x \leq \mathsf{u},\\ & \overline{x \in \mathbb{R}^n} \end{array}$$

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

Linear Program (LP)

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & \mathsf{c}^{\top}\mathsf{x}\\ \text{subject to} & \mathsf{A}\mathsf{x} \leq \mathsf{b},\\ & \mathsf{I} \leq \mathsf{x} \leq \mathsf{u},\\ & & \overline{\mathsf{x} \in \mathbb{R}^n} \end{array}$$

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

LP Relaxation of a $\ensuremath{\mathsf{MILP}}$



Problem formulation

Discrete Optimization Branch-and-Bound

The Branching Problem Learning to branch Lookback property

Our solution

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

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
- 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

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

LP Relaxation of a $\ensuremath{\mathsf{MILP}}$











Branch: Split the LP recursively over a non-integral variable, i.e. $\exists i \leq p \mid x_i^* \notin \mathbb{Z}$

 $x_i \leq \lfloor x_i^\star \rfloor \quad \lor \quad x_i \geq \lceil x_i^\star \rceil.$

Branch: Split the LP recursively over a non-integral variable, i.e. $\exists i \leq p \mid x_i^* \notin \mathbb{Z}$ $x_i \leq |x_i^*| \quad \forall \quad x_i \geq \lceil x_i^* \rceil.$

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

Branch: Split the LP recursively over a non-integral variable, i.e. $\exists i \leq p \mid x_i^* \notin \mathbb{Z}$ $x_i \leq |x_i^*| \quad \forall \quad x_i \geq \lceil x_i^* \rceil.$

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

Stopping criterion:

- \blacktriangleright L = U (optimality certificate)
- $L = \infty$ (infeasibility certificate)
- L U < threshold (early stopping)</p>

Branch: Split the LP recursively over a non-integral variable, i.e. $\exists i \leq p \mid x_i^* \notin \mathbb{Z}$ $x_i \leq |x_i^*| \quad \forall \quad x_i \geq \lceil x_i^* \rceil.$

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

Stopping criterion:

- \blacktriangleright L = U (optimality certificate)
- $L = \infty$ (infeasibility certificate)
- L U < threshold (early stopping)</p>

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



Problem formulation

Discrete Optimization Branch-and-Bound The Branching Problem

The Branching Problem

Learning to branch Lookback property

Our solution

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
Strong branching score for a variable i at a node n

Let L be the value of LP relaxation of the MILP

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 ≥ [x^{*}_i] constraint

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 ≥ [x^{*}_i] constraint
- Similarly, denote L_i^- for the other half

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 ≥ [x^{*}_i] constraint
- Similarly, denote L_i^- for the other half

Strong branching score

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

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 ≥ [x^{*}_i] constraint
- Similarly, denote L_i^- for the other half

Strong branching score

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

Strong branching decision

$$i_{SB}^{\star} = \arg \max_{i} \operatorname{score}_{SB,i}$$

Outline

Problem formulation

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

Our solution

Conclusion

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

Objective: Given a dataset of MILPs

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

Objective: Given a dataset of MILPs

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

$$\begin{split} i^{\star}_{SB} &= \arg\max_{i\in\mathcal{C}} \operatorname{score}_{SB,i} \qquad i^{\star}_{f} = \arg\max_{i\in\mathcal{C}} \operatorname{score}_{f_{\theta},i}, \\ \text{where } s^{i}_{f_{\theta}} \text{ is the score for } i \leq p \text{ variable as estimated by } f_{\theta}. \end{split}$$

Objective: Given a dataset of MILPs

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

$$\begin{split} i^{\star}_{SB} &= \arg\max_{i\in\mathcal{C}} \operatorname{score}_{SB,i} \qquad i^{\star}_{f} = \arg\max_{i\in\mathcal{C}} \operatorname{score}_{f_{\theta},i}, \\ \text{where } s^{i}_{f_{\theta}} \text{ is the score for } i \leq p \text{ variable as estimated by } f_{\theta}. \end{split}$$

$$heta^* = rgmin_{ heta} \mathcal{L}(f_{ heta}(MILP), i^{\star}_{SB})$$

Objective: Given a dataset of MILPs

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

$$\begin{split} i^{\star}_{SB} &= \arg\max_{i \in \mathcal{C}} \operatorname{score}_{SB,i} \qquad i^{\star}_{f} &= \arg\max_{i \in \mathcal{C}} \operatorname{score}_{f_{\theta},i}, \\ \text{where } s^{i}_{f_{\theta}} \text{ is the score for } i \leq p \text{ variable as estimated by } f_{\theta}. \end{split}$$

$$heta^{*} = rgmin_{ heta} \mathcal{L}(f_{ heta}(MILP), i^{\star}_{SB})$$

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 ⇒ studies the limitations of existing GNNs to represent MILPs

- + superior representation power
- + best overall accuracy

- + superior representation power
- + best overall accuracy

Model inputs

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

Natural representation : variable / constraint bipartite graph

$$\begin{array}{ll} \underset{x}{\operatorname{arg\,min}} & \mathsf{c}^\top \mathsf{x} \\ \text{subject to} & \mathsf{A}\mathsf{x} \leq \mathsf{b}, \\ & \mathsf{I} \leq \mathsf{x} \leq \mathsf{u}, \\ & \mathsf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p}. \end{array}$$

Natural representation : variable / constraint bipartite graph

$$\begin{array}{ccc} \underset{x}{\operatorname{arg\,min}} & c^{\top}x & & \overbrace{v_{0}} \\ \text{subject to} & Ax \leq b, & & \overbrace{v_{1}} \\ & I \leq x \leq u, & & \\ & x \in \mathbb{Z}^{p} \times \mathbb{R}^{n-p}. & & & \overbrace{v_{2}} \end{array}$$

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

Natural representation : variable / constraint bipartite graph



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

Natural representation : variable / constraint bipartite graph



v_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

- + superior representation power
- + best overall accuracy

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

- + 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

We consider two types of targets

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

Original one-hot encoded target, y

 $\mathbf{y}_i = egin{cases} 1, & i = i^*_{SB} \\ 0, & \text{otherwise} \end{cases}$

We consider two types of targets

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

Original one-hot encoded target, y

 $\mathbf{y}_i = egin{cases} 1, & i = i^*_{SB} \\ 0, & \text{otherwise} \end{cases}$

$$heta_y^{\star} = rgmin_{ heta} rac{1}{N} \sum_{k=1}^{N} CE(f_{ heta}(\mathcal{G}_k), \mathsf{y}_k)$$

We consider two types of targets

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



We consider two types of targets

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



Outline

Problem formulation

Our solution Loss target Regularizer Evaluation

Conclusion

We consider a regularizer to encourage the lookback proprety in GNNs

We consider a regularizer to encourage the lookback proprety in GNNs

 $loss_{PAT} = 1{Lookback_i}$.

We consider a regularizer to encourage the lookback proprety in GNNs

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

We consider a regularizer to encourage the lookback proprety in GNNs

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

Outline

Problem formulation

Our solution

Loss target Regularizer Evaluation

Conclusion
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, $\lambda_{l2} \in \{0.0, 0.01, 0.1, 1.0\}$

Performance evaluation

$$\theta_{y} = \operatorname*{arg\,min}_{\theta,\lambda_{l2}} \frac{1}{N} \sum_{k=1}^{N} CE(f_{\theta}(\mathcal{G}_{k}), \mathsf{y}_{k}) + \lambda_{l2} \cdot ||\theta||_{2}$$

Performance evaluation

$$\theta_{y} = \operatorname*{arg\,min}_{\theta,\lambda_{l2}} \frac{1}{N} \sum_{k=1}^{N} CE(f_{\theta}(\mathcal{G}_{k}), y_{k}) + \lambda_{l2} \cdot ||\theta||_{2}$$

$$\theta_z = \operatorname*{arg\,min}_{\theta,\lambda_{l2}} \frac{1}{N} \sum_{k=1}^{N} CE(f_{\theta}(\mathcal{G}_k), \mathsf{z}_k) + \lambda_{l2} \cdot ||\theta||_2$$

Performance evaluation

$$\theta_{y} = \operatorname*{arg\,min}_{\theta,\lambda_{l2}} \frac{1}{N} \sum_{k=1}^{N} CE(f_{\theta}(\mathcal{G}_{k}), y_{k}) + \lambda_{l2} \cdot ||\theta||_{2}$$

$$\theta_z = \operatorname*{arg\,min}_{\theta,\lambda_{l2}} \frac{1}{N} \sum_{k=1}^{N} CE(f_{\theta}(\mathcal{G}_k), \mathsf{z}_k) + \lambda_{l2} \cdot ||\theta||_2$$

$$\theta_{PAT} = \underset{\theta, \mathsf{v}, \lambda_{I2}, \lambda_{PAT}}{\arg\min} \frac{1}{N} \sum_{k=1}^{N} CE(f_{\theta}(\mathcal{G}_{k}), \mathsf{v}) + \lambda_{I2} \cdot ||\theta||_{2} + \lambda_{PAT} \cdot \mathsf{loss}_{PAT}$$

Performance evaluation: Instances

Small instances are used to <u>collect training data</u> 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 <u>collect training data</u> 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 <u>hyperparameter selection</u> incorporating harder-to-formulate criterion in the objective function

Performance evaluation: Instances

- Small instances are used to <u>collect training data</u> 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 <u>hyperparameter selection</u> 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.)

Model	Time	Time (c)	Wins	Solved	Nodes (c)
${ m FSB}^{*}$	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17979
TUNEDRPB	644.20	450.06	0	80	18104
GNN	507.06	333.59	14	80	17145
GNN-PAT (ours)	477.26	310.22	69	84	16388

Combinatorial Auction (Bigger)

Model	Time	Time (c)	Wins	Solved	Nodes (c)
${\rm FSB}^*$	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17979
TUNEDRPB	644.20	450.06	0	80	18104
GNN	507.06	333.59	14	80	17145
gnn-PAT (ours)	477.26	310.22	69	84	16388

Combinatorial Auction (Bigger)

Model	Time	Time (c)	Wins	Solved	Nodes (c)
${ m FSB}^{*}$	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17979
TUNEDRPB	644.20	450.06	0	80	18104
GNN	507.06	333.59	14	80	17145
GNN-PAT (ours)	477.26	310.22	69	84	16388

Combinatorial Auction (Bigger)

Model	Time	Time (c)	Wins	Solved	Nodes (c)
${ m FSB}^{*}$	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17979
TUNEDRPB	644.20	450.06	0	80	18104
GNN	507.06	333.59	14	80	17145
gnn-PAT (ours)	477.26	310.22	69	84	16388

Combinatorial Auction (Bigger)

Model	Time	Time (c)	Wins	Solved	Nodes (c)
${ m FSB}^*$	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17979
TUNEDRPB	644.20	450.06	0	80	18104
GNN	507.06	333.59	14	80	17145
GNN-PAT (ours)	477.26	310.22	69	84	16388

Combinatorial Auction (Bigger)

Model	Time	Time (c)	Wins	Solved	Nodes (c)
${ m FSB}^{*}$	n/a	n/a	n/a	n/a	n/a
RPB	626.81	434.92	1	80	17979
TUNEDRPB	644.20	450.06	0	80	18104
GNN	507.06	333.59	14	80	17145
GNN-PAT (ours)	477.26	310.22	69	84	16388

Combinatorial Auction (Bigger)

Optimality gap on commonly unsolved instances



Figure: Mean optimality gap of the commonly unsolved instances

Outline

Problem formulation

Our solution

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

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

- We discover lookback phenomenon in the gold-standard (by tree size) variable-selection heuristic
- We proposed second-best e-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

- We discover lookback phenomenon in the gold-standard (by tree size) variable-selection heuristic
- We proposed second-best e-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

Discovery of more inductive biases

QR Codes generated via https://www.qr-code-generator.com/

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

QR Codes generated via https://www.qr-code-generator.com/

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

QR Codes generated via https://www.qr-code-generator.com/

- 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

QR Codes generated via https://www.qr-code-generator.com/

Lookback for Learning to Branch

Thank you!

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

