Exact Combinatorial Optimization with Graph Convolutional Neural Networks

Laurent Charlin
with Maxime Gasse, Didier Chételat, Nicola Ferroni, Andrea Lodi

February 24, 2021
(slide deck adapted from M. Gasse)

POLYTECHNIQUE MONTRÉAL
HEC MONTRÉAL
DATA SCIENCE FOR REAL-TIME DECISION-MAKING
Mila

Problem Setting

- Goal: Solve combinatorial optimization
  - Solvers provide exact solutions
- Use case:
  - Regularly solve problems from the same family (e.g. TSP)
    - Learn offline from a collected dataset (train)
    - Help solve new problems from the same family (test)
- Machine learning as a heuristic inside the CO solver
- Focuss on the problem of branching

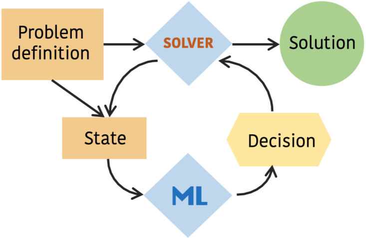

Overview

- Brief introduction to Branching
- Machine learning modelling
  - MDP
  - GCNN
- Empirical Study

Brief introduction to Branching

Mixed-Integer Linear Program (MILP)

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

- $c \in \mathbb{R}^n$ 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.

Linear Program (LP) relaxation

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

Convex problem, efficient algorithms (e.g., simplex).

- $x^\star \in \mathbb{Z}^p \times \mathbb{R}^{n-p}$ (lucky) $\rightarrow$ solution to the original MILP
- $x^\star \notin \mathbb{Z}^p \times \mathbb{R}^{n-p}$ $\rightarrow$ lower bound to the original MILP

Branch-and-Bound

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

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

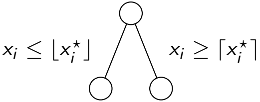

Lower bound (L): minimal among leaf nodes.
Upper bound (U): minimal among leaf nodes with integral solution.

Stopping criterion:

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

Branch-and-Bound

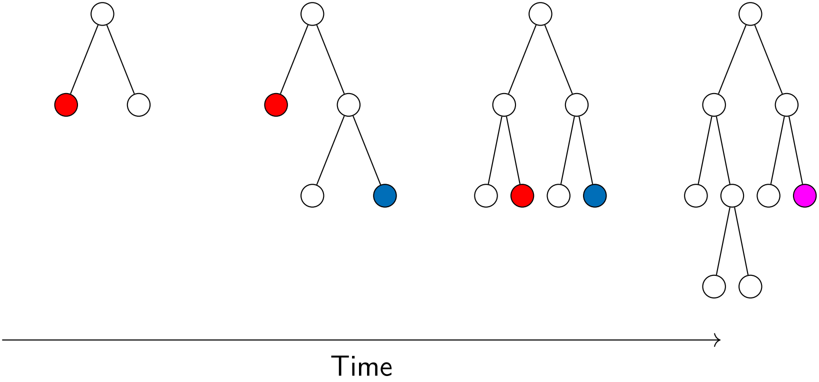

Branch-and-Bound: a sequential process

Sequential decisions:

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

State-of-the-art in B&B solvers: expert rules

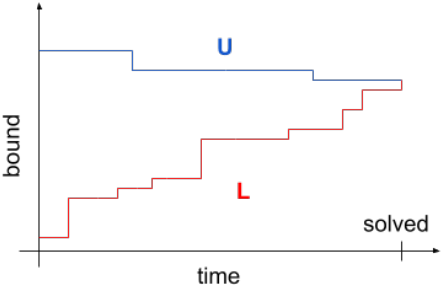

Objective: no clear consensus

- L = U fast ?
- U - L ↘ fast ?
- L ↗ fast ?
- U ↘ fast ?

Expert branching rules: state-of-the-art

Strong branching (SB): one-step forward looking (greedy)

- solve both LPs for each candidate variable
- pick the variable resulting in tightest relaxation
- \+ small trees
- − computationally expensive

Pseudo-cost: backward looking

- keep track of tightenings in past branchings
- pick the most promising variable
- \+ very fast, almost no computations
- − cold start

Reliability pseudo-cost: best of both worlds

- compute SB scores at the beginning
- gradually switches to pseudo-cost (+ other heuristics)
- \+ best overall solving time trade-off (on MIPLIB)

Machine learning modelling

Markov Decision Process

Objective: take actions which maximize the long-term reward

$$\sum_{t=0}^{\infty} \gamma^t r(s_t),$$

with $r : \mathcal{S} \to \mathbb{R}$ a reward function and discount factor γ

Branching as a Markov Decision Process

State: the whole internal state of the solver, s.
Action: a branching variable, $a \in \{1, \dots, p\}$.

Trajectory: $\tau = (s_0, \dots, s_T)$

- initial state s_0: a MILP $\sim p(s_0)$;
- terminal state s_T: the MILP is solved;
- intermediate states: branching

$$s_{t+1} \sim p_\pi(s_{t+1}|s_t) = \sum_{a \in \mathcal{A}} \underbrace{\pi(a|s_t)}_{\text{branching policy}} \underbrace{p(s_{t+1}|s_t, a)}_{\text{solver internals}}.$$

<u>Branching problem</u>: solve

$$\pi^\star = \arg\max_\pi \mathop{\mathbb{E}}_{\tau \sim p_\pi} [r(\tau)],$$

with $r(\tau) = \sum_{s \in \tau} r(s)$.

Challenges

MDP $\implies$ Reinforcement learning (RL) ?

State representation: s

- ▶ global level: original MILP, tree, bounds, focused node...
- ▶ node level: variable bounds, LP solution, simplex statistics...
- – dynamically growing structure (tree)
- – variable-size instances (cols, rows) $\implies$ Graph Neural Network

Sampling trajectories: $\tau \sim p_\pi$

- ▶ collect one τ = solving a MILP (with π likely not optimal)
- – expensive $\implies$ train on small instances

Reward function: r

- ▶ no consensus
- + a strong expert exists $\implies$ imitation learning

Machine learning approaches (Pre-2020)

Node selection

- He et al., 2014
- Song et al., 2018

Variable selection (branching)

- Khalil, Le Bodic, et al., 2016 $\implies$ "online" imitation learning
- Hansknecht et al., 2018 $\implies$ offline imitation learning
- Balcan et al., 2018 $\implies$ theoretical results

Cut selection

- Baltean-Lugojan et al., 2018
- Tang et al., 2019

Primal heuristic selection

- Khalil, Dilkina, et al., 2017
- Hendel et al., 2018

State encoding

Natural representation : variable / constraint bipartite graph

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

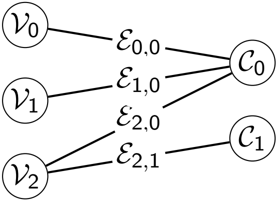

- $\mathcal{V}_i$: variable features (type, coef. c, bounds, LP solution...)
- $\mathcal{C}_j$: constraint features (right-hand-side, LP slack...)
- $\mathcal{E}_{i,j}$: non-zero coefficients in A

Graph structure is fixed through time.

D. Selsam et al. (2019). Learning a SAT Solver from Single-Bit Supervision.

Branching Policy as a GCNN Model

Neighbourhood-based updates: $\mathrm{v}_i \leftarrow \sum_{j \in \mathcal{N}_i} \mathrm{f}_\theta(\mathrm{v}_i, \mathrm{e}_{i,j}, \mathrm{c}_j)$

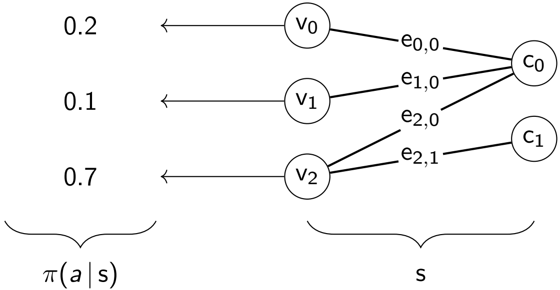

Natural model choice for graph-structured data

- permutation-invariance
- benefits from sparsity

T. N. Kipf et al. (2016). Semi-Supervised Classification with Graph Convolutional Networks.

Implementation Details

Interleaved half convolutions

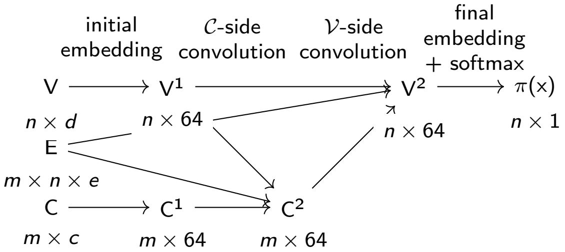

Prenorm-layer

- standardize conv. operations

Empirical Study

Imitation learning

Full Strong Branching (FSB): good branching rule, but expensive.
Can we learn a fast, good-enough approximation ?

Behavioural cloning

- collect $\mathcal{D} = \{(\mathrm{s}, a^\star), \dots\}$ from the expert agent (FSB)
- estimate $\pi^\star(a \mid \mathrm{s})$ from $\mathcal{D}$
- \+ no reward function, supervised learning, well-behaved
- − will never surpass the expert...

Implementation with the open-source solver SCIP[1]

Not a new idea

- Alvarez et al., 2017 predict SB scores, XTrees model
- Khalil, Le Bodic, et al., 2016 predict SB rankings, SVMrank model
- Hansknecht et al., 2018 do the same, λ-MART model

[1] A. Gleixner et al. (July 2018). The SCIP Optimization Suite 6.

Minimum set covering[2]

- Comparing the accuracy of machine learning models

Set Covering			
model	acc@1	acc@5	acc@10
XTrees	51.8±0.3	80.5±0.1	91.4±0.2
SVMrank	57.6±0.2	84.7±0.1	94.0±0.1
λ-MART	57.4±0.2	84.5±0.1	93.8±0.1
GCNN	**65.5**±0.1	**92.4**±0.1	**98.2**±0.0

[2] E. Balas et al. (1980). Set covering algorithms using cutting planes, heuristics, and subgradient optimization: a computational study.

Minimum set covering[3]

Model	Easy Time	Easy Wins	Easy Nodes	Medium Time	Medium Wins	Medium Nodes	Hard Time	Hard Wins	Hard Nodes
FSB	17.30	0 / 100	17	411.34	0 / 90	171	3600.00	0 / 0	n/a
RPB	8.98	0 / 100	**54**	60.07	0 / 100	1741	1677.02	4 / 65	47 299
XTrees	9.28	0 / 100	187	92.47	0 / 100	2187	2869.21	0 / 35	59 013
SVMrank	8.10	1 / 100	165	73.58	0 / 100	1915	2389.92	0 / 47	42 120
λ-MART	7.19	14 / 100	167	59.98	0 / 100	1925	2165.96	0 / 54	45 319
GCNN	**6.59**	**85** / 100	134	**42.48**	**100** / 100	**1450**	**1489.91**	**66** / 70	**29 981**

3 problem sizes

- 500 rows, 1000 cols (easy), training distribution
- 1000 rows, 1000 cols (medium)
- 2000 rows, 1000 cols (hard)

Pays off: better than SCIP's default in terms of solving time.
Generalizes to harder problems !

[3] E. Balas et al. (1980). Set covering algorithms using cutting planes, heuristics, and subgradient optimization: a computational study.

Combinatorial auction[4]

	Easy			Medium			Hard		
Model	Time	Wins	Nodes	Time	Wins	Nodes	Time	Wins	Nodes
FSB	4.11	0 / 100	6	86.90	0 / 100	72	1813.33	0 / 68	400
RPB	2.74	0 / 100	10	17.41	0 / 100	689	136.17	13 / 100	5511
XTrees	2.47	0 / 100	86	23.70	0 / 100	976	451.39	0 / 95	10 290
SVMrank	2.31	0 / 100	77	23.10	0 / 100	867	364.48	0 / 98	6329
λ-MART	**1.79**	**75** / 100	77	14.42	1 / 100	873	222.54	0 / 100	7006
GCNN	1.85	25 / 100	**70**	**10.29**	**99** / 100	**657**	**114.16**	**87** / 100	**5169**

3 problem sizes

- 100 items, 500 bids (easy), training distribution
- 200 items, 1000 bids (medium)
- 300 items, 1500 bids (hard)

[4] K. Leyton-Brown et al. (2000). Towards a Universal Test Suite for Combinatorial Auction Algorithms.

Capacitated facility location[5]

Model	Easy			Medium			Hard		
	Time	Wins	Nodes	Time	Wins	Nodes	Time	Wins	Nodes
FSB	30.36	4 / 100	14	214.25	1 / 100	76	742.91	15 / 90	55
RPB	26.55	9 / 100	**22**	156.12	8 / 100	**142**	631.50	14 / 96	**110**
XTrees	28.96	3 / 100	135	159.86	3 / 100	401	671.01	1 / 95	381
SVMrank	23.58	11 / 100	117	130.86	13 / 100	348	586.13	21 / 95	321
λ-MART	23.34	16 / 100	117	128.48	23 / 100	349	582.38	15 / 95	314
GCNN	**22.10**	**57** / 100	107	**120.94**	**52** / 100	339	**563.36**	**30** / 95	338

3 problem sizes

- 100 facilities, 100 customers (easy), training distribution
- 100 facilities, 200 customers (medium)
- 100 facilities, 400 customers (hard)

[5] G. Cornuejols et al. (1991). A comparison of heuristics and relaxations for the capacitated plant location problem.

Maximum independent set[6]

Model	Easy Time	Easy Wins	Easy Nodes	Medium Time	Medium Wins	Medium Nodes	Hard Time	Hard Wins	Hard Nodes
FSB	23.58	9 / 100	7	1503.55	0 / 74	38	3600.00	0 / 0	n/a
RPB	8.77	7 / 100	**20**	**110.99**	41 / 100	729	2045.61	22 / 42	**2675**
XTrees	10.75	1 / 100	76	1183.37	1 / 47	4664	3565.12	0 / 3	38 296
SVMrank	8.83	2 / 100	46	242.91	1 / 96	**546**	2902.94	1 / 18	6256
λ-MART	7.31	30 / 100	52	219.22	15 / 91	747	3044.94	0 / 12	8893
GCNN	**6.43**	51 / 100	43	192.91	**42** / 82	1841	**2024.37**	**25** / 29	2997

3 problem sizes, Barabási-Albert graphs (affinity=4)

- 500 nodes (easy), training distribution
- 1000 nodes (medium)
- 1500 nodes (hard)

[6] D. Chalupa et al. (2014). On the Growth of Large Independent Sets in Scale-Free Networks.

Conclusion

Heuristic vs data-driven branching:

+ tune B&B to your problem of interest automatically
− no guarantees outside of the training distribution
− requires training instances

Paper: `https://arxiv.org/abs/1906.01629` M. Gasse et al. (2019). Exact Combinatorial Optimization with Graph Convolutional Neural Networks.

Code: `https://github.com/ds4dm/learn2branch` Since:

- Branching for reinforcement learning Sun et al., 2021
- Properties of GCNNs for CO Cappart et al., 2021
- GNN-like operations on the CPU Gupta et al., 2020
- Ecole.ai: API + benchmarks for MILP adaptive solving (based on the open-source SCIP solver) – See Maxime's talk later today.

Exact Combinatorial Optimization with Graph Convolutional Neural Networks

Thank you!

Laurent Charlin
with Maxime Gasse, Didier Chételat, Nicola Ferroni, Andrea Lodi