

ML4CO

Methodology of Machine Learning for Combinatorial Optimization

Antoine Prouvost, Andrea Lodi

`antoine.prouvost@polymtl.ca`

Polytechnique Montréal

28th of June 2018

CANADA
EXCELLENCE
RESEARCH
CHAIR



DATA SCIENCE
FOR REAL-TIME
DECISION-MAKING



Mila

POLYTECHNIQUE
MONTREAL



WORLD-CLASS
ENGINEERING



Motivation

IBM Deep Blue



v.s.



Deep Mind AlphaGo



Motivation

Traditional OR algorithms

- 👍 Hard rules & constraints
- 👎 Can have big complexity

Focuses on all possible cases

Deep Learning

- 👍 Perception: fast, intuitive approximate
- 👎 Integers & rules

Focuses on the cases shown

Objectives



Sometimes in OR, we just don't know.

→ Build mixed algorithm where **decisions** are arbitrary learned.



Sometimes in OR, the complexity is just too big.

→ Build a fast **approximator** using deep learning.

In this talk, we refer to both cases as **estimation**.

A 3-class taxonomy

End-to-End learning

Build a machine learning model to directly output solutions.

Offline estimation

Build a machine learning model to parametrize your OR algorithm.

Online estimation

The machine learning model is used repeatedly by the OR algorithm to assist it in its solving process.

End-to-end learned algorithms

- ▶ Given a problem, learn the **solution**.
*E.g. Given a weighted graph, learn a solution to the TSP.*¹
- ▶ Solutions may not be fully **detailed**.
*E.g. Predict only aggregation of MILP solution*².
- ▶ **Runtime**: input the problem to the ML model, get the solution.

¹Bello et al. 2016; Dai et al. 2017; Emami and Ranka 2018; Kool and Welling 2018; Nowak et al. 2017; Vinyals, Fortunato, and Jaitly 2015.

²Larsen et al. n.d.

End-to-end learned algorithms 🧠

Advantages 😊

- 👍 Fast inference (Polynomial complexity)
- 👍 Based on data appropriate to the task
- 👍 Repurposed by retraining
- 👍 Good default solution when nothing else is available

Disadvantages 😞

- 👎 Needs data (eventually labels)
- 👎 May need (very) expensive training
- 👎 **Very** heuristic:
 - ▶ No optimality guarantee
 - ▶ Little feasibility guarantee
 - Dedicated architectures
 - Must not break differentiability
- 👎 Generalization with size?

Offline estimation

- ▶ Predict a **property** of a problem instance that can get **exploited** by an OR algorithm.
E.g. Learn when to linearize MIQP³, when to apply DW decomposition⁴.
- ▶ May be seen as a more general case of the previous.
- ▶ **Runtime**: Run the prediction, then run the OR algorithm accordingly.

³Bonami, Lodi, and Zarpellon 2018.

⁴Kruber, Lübbecke, and Parmentier 2017.

- ▶ Throughout its **progress**, the OR algorithm **repeatedly** calls the ML model.
E.g. Learning to branch⁵, where to run heuristics⁶ in MILP. Learning to apply SGD updates⁷. Learning to select cutting plane in convex SDP relaxation of QP⁸.
- ▶ Can also build heuristics⁹.
- ▶ **Runtime**: The OR & the ML algorithms need to be deployed and run together → more specific code.

⁵Lodi and Zarpellon 2017.

⁶Shao et al. 2017.

⁷Andrychowicz et al. 2016; Li and Malik 2016, 2017; Wichrowska et al. 2017.

⁸Baltea-Lugojan and Misener n.d.

⁹Dai et al. 2017.

Representation challenge

Deep learning works well for natural signals (images, speech, etc).
Will it work well for OR problems?



- ▶ Are common architectures the best prior?
- ▶ Additional constraints: large sparse structured inputs/ outputs of variable sizes.
- ▶ Partial solutions: graph NN, attention, CNN, RNN and other parameter reuse.

Supervised vs reinforcement

Supervised

- 👍 Straightforward
- 👎 Expensive targets
- 👎 Not well suited when multiples targets are corrects (E.g. multiple solutions)

Reinforcement

- 👍 Natural formulation to the *online* setting
- 👍 Naturally **maximizes** sum of expected **future rewards**
- 👎 More complex to define and train

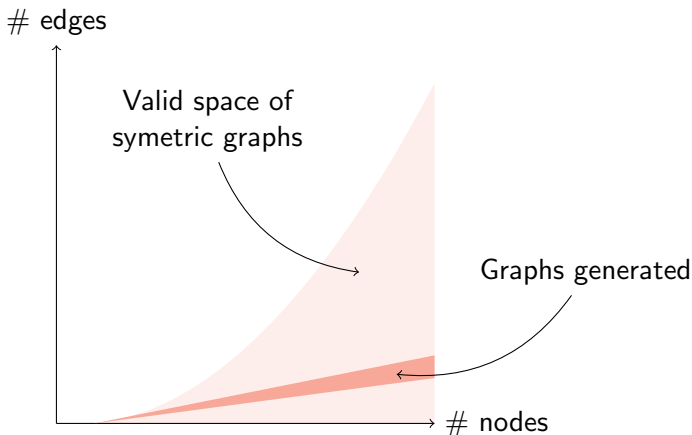
Supervised in the online setting?

- ▶ No apparent link between performance of ML and performance of overall algorithm.
- ▶ When estimating a *known* quantity, it's possible to link the performance fo ML to the overall performance¹⁰

¹⁰Baltean-Lugojan and Misener n.d.; Shao et al. 2017.

Data distributions

Easily fooled 🤖



Especially if the test set is generated from the same distribution as the training set 🧑🏫 🚚

Historical data?

“[S]ampling from historical data is appropriate when attempting to mimic a behavior reflected in such data.”¹¹

¹¹Larsen et al. n.d.

Other motivations

- ▶ ML to model uncertainty¹².
- ▶ Extract knowledge out of learned models¹³


¹²Larsen et al. n.d.

¹³Bonami, Lodi, and Zarpellon 2018; Dai et al. 2017.

Questions



References I

-  Andrychowicz, Marcin et al. (2016). “Learning to Learn by Gradient Descent by Gradient Descent”. In: arXiv: 1606.04474 [cs].
-  Baltean-Lugojan, Radu and Ruth Misener (n.d.). “Learning-Based Cutting Plane Approximation of QP Convex (SDP) Relaxations”. In: p. 12.
-  Bello, Irwan et al. (2016). “Neural Combinatorial Optimization with Reinforcement Learning”. In: arXiv: 1611.09940 [cs, stat].
-  Bonami, Pierre, Andrea Lodi, and Giulia Zarpellon (2018). “Learning a Classification of Mixed-Integer Quadratic Programming Problems”. In: Integration of Constraint Programming, Artificial Intelligence, and Operations Research International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research. Lecture Notes in Computer Science. Springer, Cham, pp. 595–604.
-  Dai, Hanjun et al. (2017). “Learning Combinatorial Optimization Algorithms over Graphs”. In: arXiv: 1704.01665 [cs, stat].
-  Emami, Patrick and Sanjay Ranka (2018). “Learning Permutations with Sinkhorn Policy Gradient”. In: arXiv: 1805.07010 [cs, stat].
-  Kool, W. W. M. and M. Welling (2018). “Attention Solves Your TSP”. In: arXiv: 1803.08475 [cs, stat].

References II



Kruber, Markus, Marco E. Lübbecke, and Axel Parmentier (2017). “Learning When to Use a Decomposition”. In: Integration of AI and OR Techniques in Constraint Programming. International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems. Lecture Notes in Computer Science. Springer, Cham, pp. 202–210.



Larsen, Eric et al. (n.d.). “A Machine Learning Approximation Algorithm for Fast Prediction of Solutions to Discrete Optimization Problems”. In: p. 22.



Li, Ke and Jitendra Malik (2016). “Learning to Optimize”. In: arXiv: 1606.01885 [cs, math, stat].



– (2017). “Learning to Optimize Neural Nets”. In: arXiv: 1703.00441 [cs, math, stat].



Lodi, Andrea and Giulia Zarpellon (2017). “On Learning and Branching: A Survey”. In: TOP 25.2, pp. 207–236.



Nowak, Alex et al. (2017). “A Note on Learning Algorithms for Quadratic Assignment with Graph Neural Networks”. In: arXiv: 1706.07450 [cs, stat].



Shao, Yufen et al. (2017). “Learning to Run Heuristics in Tree Search”. In: pp. 659–666.

References III



Vinyals, Oriol, Meire Fortunato, and Navdeep Jaitly (2015). “Pointer Networks”. In: Advances in Neural Information Processing Systems 28. Ed. by C. Cortes et al. Curran Associates, Inc., pp. 2692–2700.



Wichrowska, Olga et al. (2017). “Learned Optimizers That Scale and Generalize”. In:

Emoji artwork provided by EmojiOne under free license.