Decision-Focused Learning: Through the Lens of Learning to Rank

Jayanta Mandi, Víctor Bucarey, Maxime Mulamba, Tias Guns

VRIJE UNIVERSITEIT BRUSSEL · Universidad de O'Higgins · KU LEUVEN

ICML International Conference On Machine Learning

Problem Description

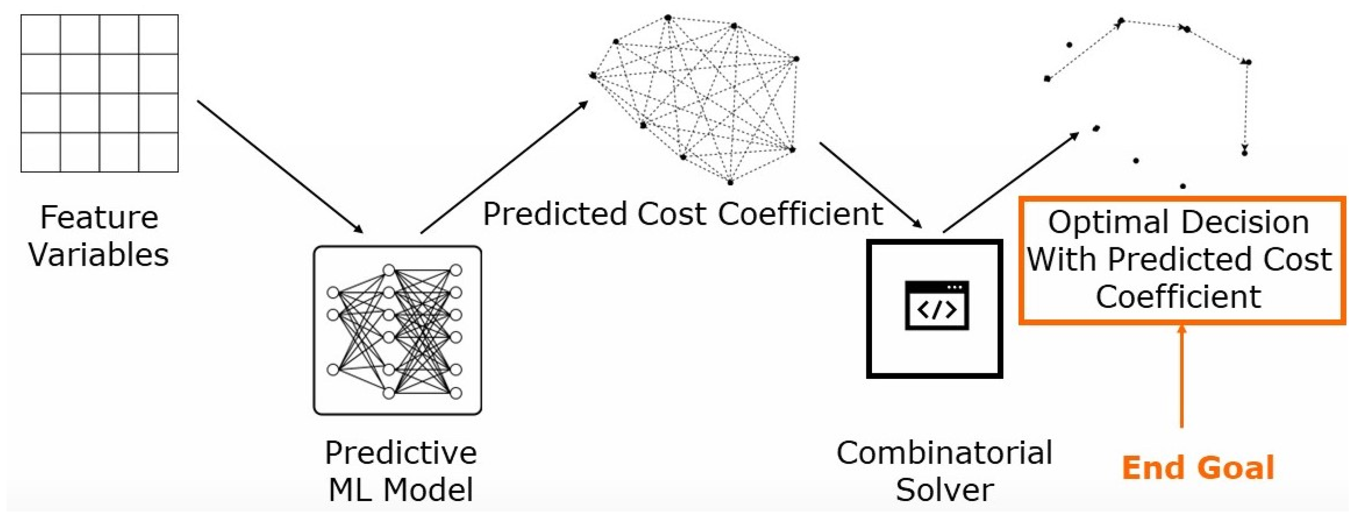

Decision-Focused Learning by Minimizing Regret

We consider combinatorial optimization problem $v^\star(c) = \text{argmin}_{v \in V} f(v, c)$.
Train to minimize regret of predicting $\hat{c}$: regret$(\hat{c}, c) = f(v^\star(\hat{c}), c) - f(v^\star(c), c)$

Training Loop in Decision-focused Learning

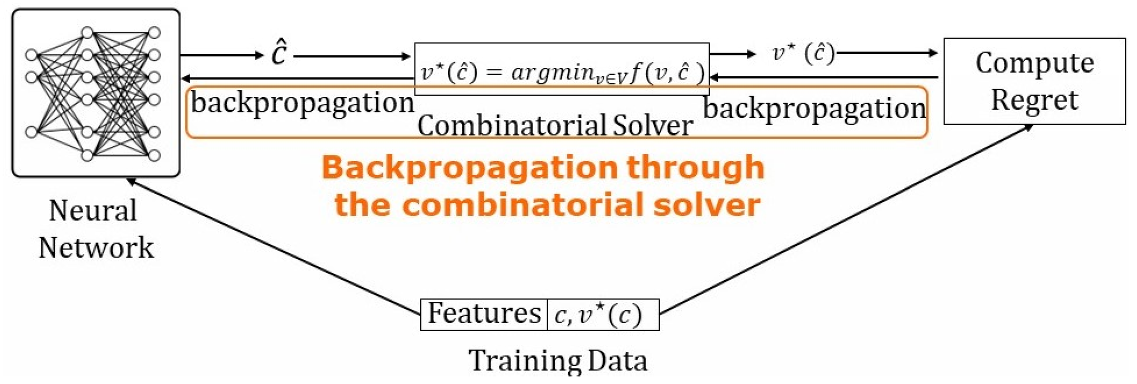

Challenge

- Repeatedly solving the combinatorial optimization problem for each instance in each training loop.
- We propose training with limited number of solving the optimzation problem.

Rank-based Loss function

- We want to predict $\hat{c}$ so that $f(v, c)$ and $f(v, \hat{c})$ follow same ordering in domain V.
- Formulated as learning to rank $v \in V$ w.r.t. objective function $f(.,.)$.
- In practice, as V is intractable, we learn to rank $v \in S \subset V$.

Proposed Training Loop with Rank-based loss

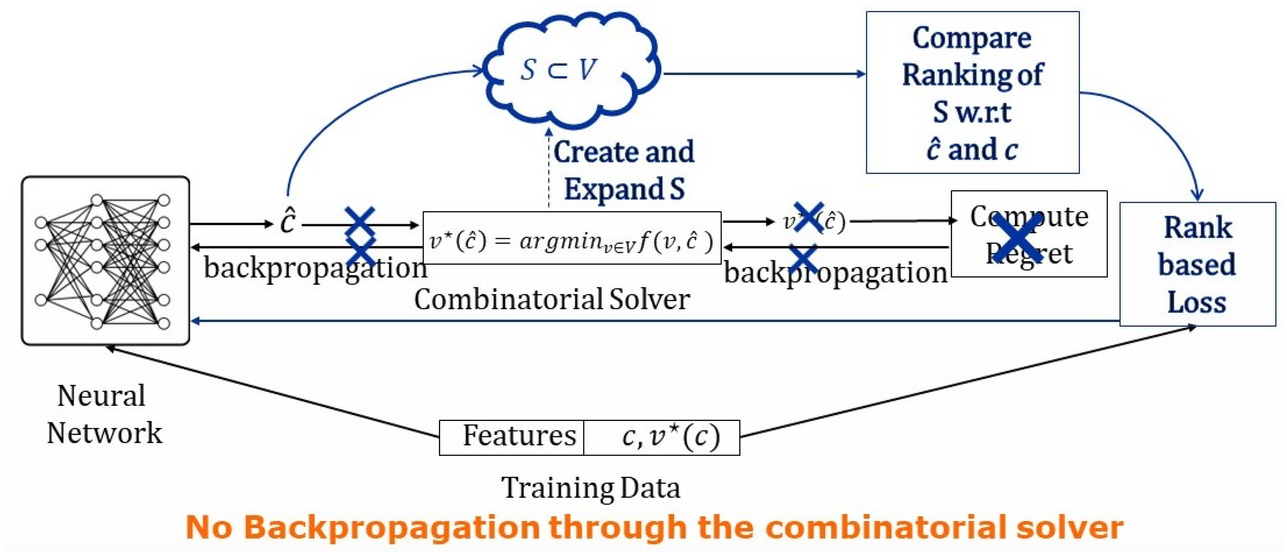

Pointwise Loss Functions

Regress predicted objective values $f(v, \hat{c})$ on the actual objective values $f(v, c)$ for $v \in S$:

$$\frac{1}{|S|} \sum_{v \in S} \left(f(v, \hat{c}) - f(v, c)\right)^2 \tag{1}$$

Pairwise Loss Functions

Instead of treating each $v \in S$ separately, here we do a pairwise comparison of $(v_p, v_q) \in S$:

$$\sum_{(v_p, v_q) \in \{(v_p, v_q) | f(v_p, c) < f(v_q, c)\}} \max\left(0, \nu + f(v_p, \hat{c}) - f(v_q, \hat{c})\right) \tag{2}$$

Generalization of NCE loss (Mulamba et al.)

Instead of all possible $\mathcal{O}(|S|^2)$ pairs, we propose *best-versus-rest* pair generation scheme, where we compare all other $v \in S$ with $v_{best} = \text{argmin}_{v \in S} f(v, c)$.

$$\sum_{v \in S} \max\left(0, (\nu + f(v_{best}, \hat{c}) - f(v, \hat{c}))\right) \tag{3}$$

The NCE loss proposed by [1] is particular case of this without the Relu operator and $\nu = 0$.

Pairwise Difference Loss Functions

Regress pairwise difference of predicted objective values on difference of true objective values.

$$\sum_{v \in S} \left(\left(f(v_{best}, \hat{c}) - f(v, \hat{c})\right) - \left(f(v_{best}, c) - f(v, c)\right) \right)^2 \tag{4}$$

Listwise Loss Function

Listwise loss is computed based on the ordering of the entire set S. We define the following discrete probability distribution of $v \in S$ being v_{best}

$$p(v; \tau, c) = \frac{1}{Z} \sum_{v' \in V} \exp\left(-\frac{f(v', c)}{\tau}\right) \forall v \in S \tag{5}$$

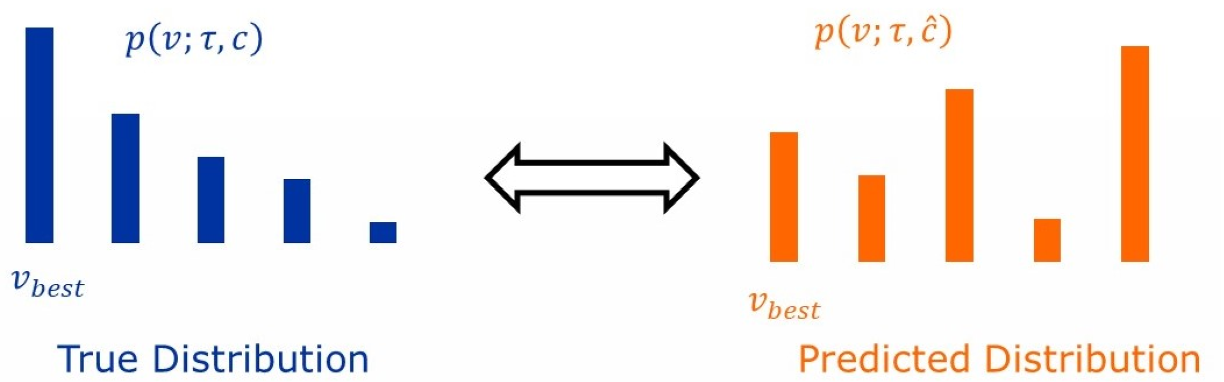

We define **listwise loss** as the cross entropy loss between the predicted and the true distribution:

$$-\frac{1}{|S|} \sum_{v \in S} p(v; \tau, c) \log p(v; \tau, \hat{c}) \tag{6}$$

We lower regret by minimizing rank-based loss

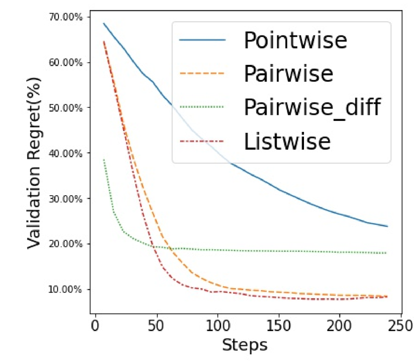

Degree- 4

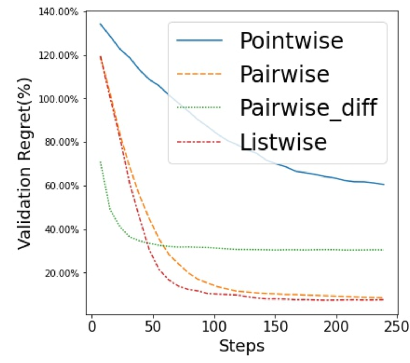

Degree- 6

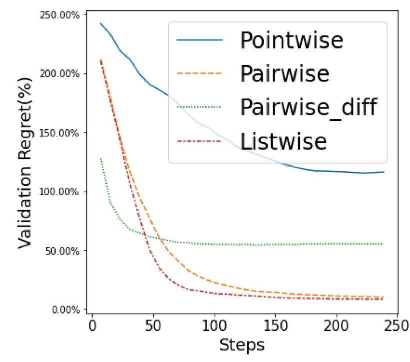

Degree- 8

Comparison with Other Approaches

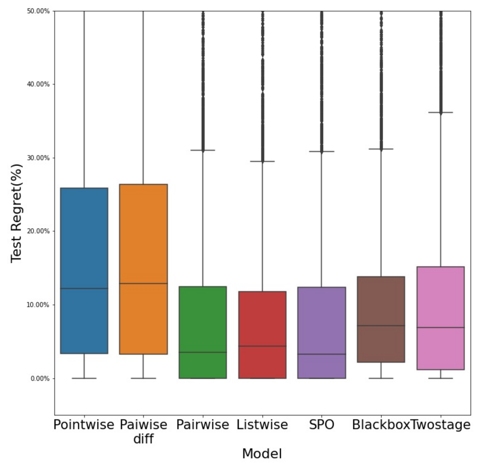

Degree- 4

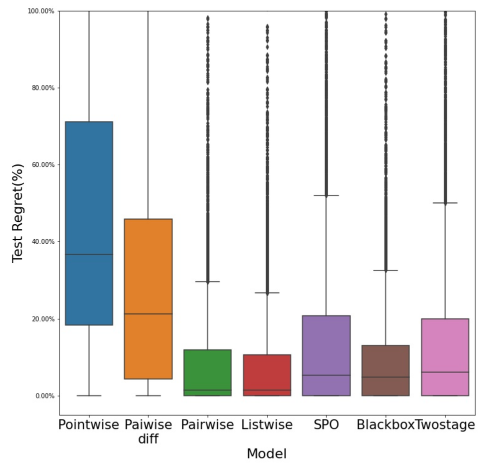

Degree- 6

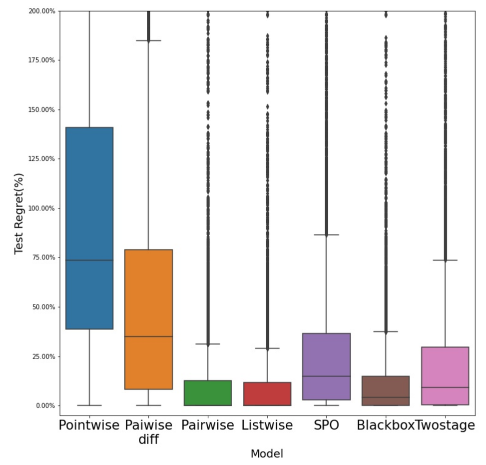

Degree- 8

Efficiency Gain in Training Time

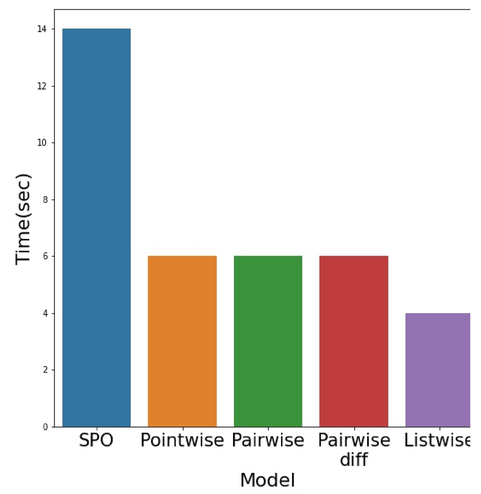

Degree- 4

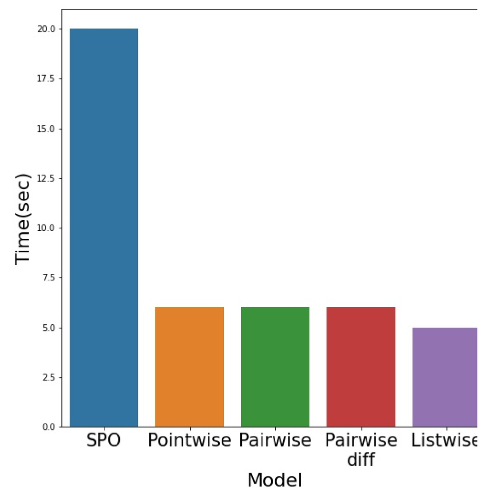

Degree- 6

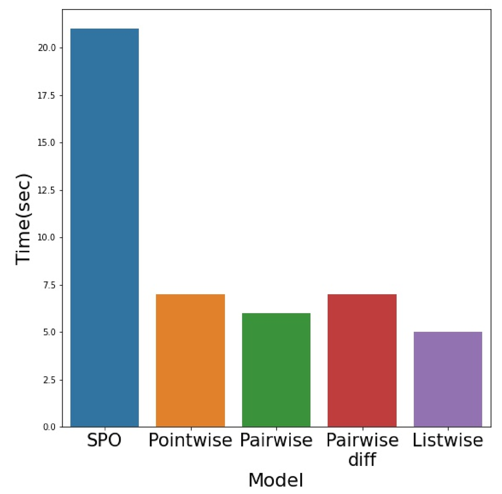

Degree- 8

References

References

[1] Maxime Mulamba, Jayanta Mandi, Michelangelo Diligenti, Michele Lombardi, Victor Bucarey, and Tias Guns. Contrastive losses and solution caching for predict-and-optimize. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 2833–2840. International Joint Conferences on Artificial Intelligence Organization, 8 2021.

[2] Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *AAAI-19*.

[3] Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *arXiv preprint arXiv:1710.08005*, 2017.

[4] Jayanta Mandi, Tias Guns, Emir Demirović, and Peter. J Stuckey. Smart predict-and-optimize for hard combinatorial optimization problems. In *AAAI 2020 : The Thirty-Fourth AAAI Conference on Artificial Intelligence*, volume 34, pages 1603–1610, 2020.