KU LEUVEN

Scalability in Decision-Focused Learning: State of the Art, Challenges, and Beyond

Jayanta Mandi
Postdoctoral Researcher
DTAI Research Unit
Department of Computer Science
KU Leuven

Gratitude to My Wonderful Collaborators

Maxime Mulamba

Victor Bucarey Lopez

Rocsildes Canoy

Michele Lombardi

Senne Berden

Ferdinando Fioretto

Tias Guns

Jury Committee Members

Bistra Dilkina

Mathias Niepert

Ann Nowé

Vincent Ginis

Sam Verboven

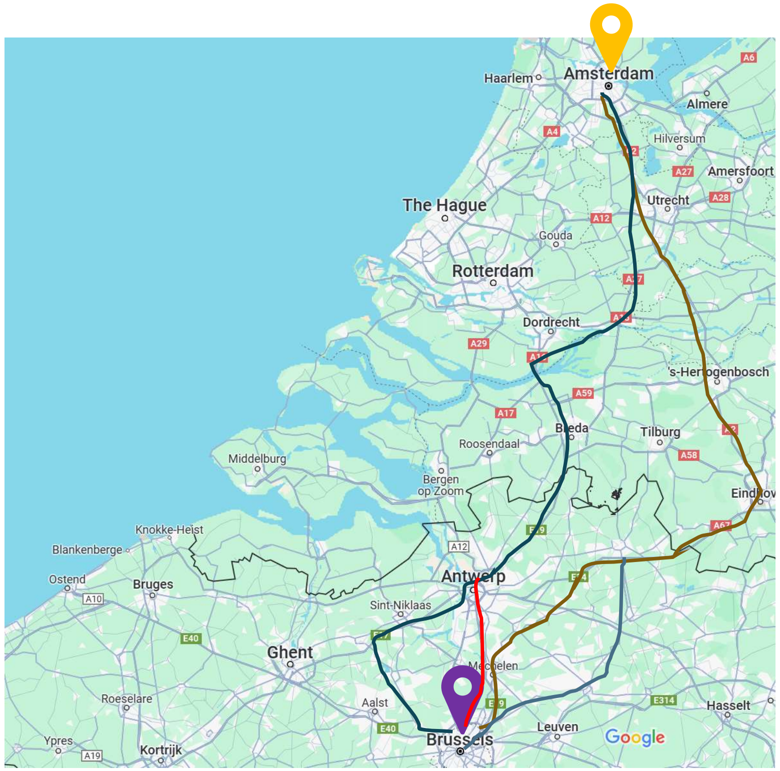

- What is the quickest route to go from A to B?
- A very simple CP problem if **the traffic congestion is known.**
- But as the future traffic congestion is unknown, it must be estimated using *contextual* features.

(Combinatorial) Optimization (CO)

$$min_{w \in \mathcal{F}(k)} f(c, w)$$

- f: objective function to be minimized
- $\mathcal{F}$: The set of feasible points
- w: decision variable

- **k**: parameters, defining the set of **feasible points(constraints)**
- **c**: parameters, defining the **objective function (cost parameter)**
- $\mathbf{w}^*(\mathbf{c}, \mathbf{k})$: **a parameteric solution** to the optimization for the parameter set $(\mathbf{c}, \mathbf{k})$

Solving a Single Instance of a Combinatorial Optimization Problem

Known at the time of Solving the problem

For one problem instance, one can solve it using a CP solver (or MIP, SAT).

- Job-shop Scheduling Problem
- Portfolio Optimization Problem
- Vehicle Routing Problem
- Bipartite Matching Problem
- Bin Packing Problem

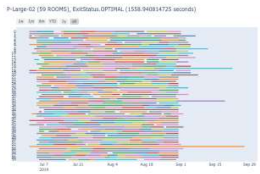

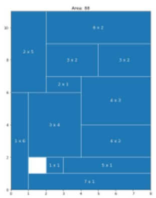

Predict-then-Optimize (PtO) Problem

- The parameter (c, k) is ***not*** known at the time of solving the optimization problem.

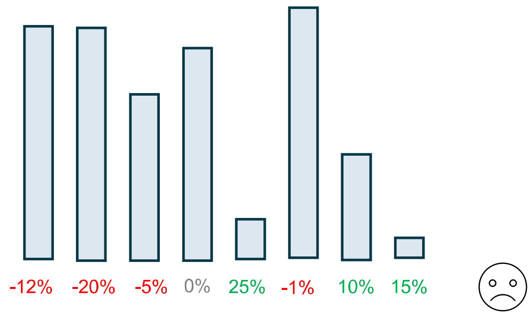

Asset Allocation for Portfolio Optimization
(Predicting c)

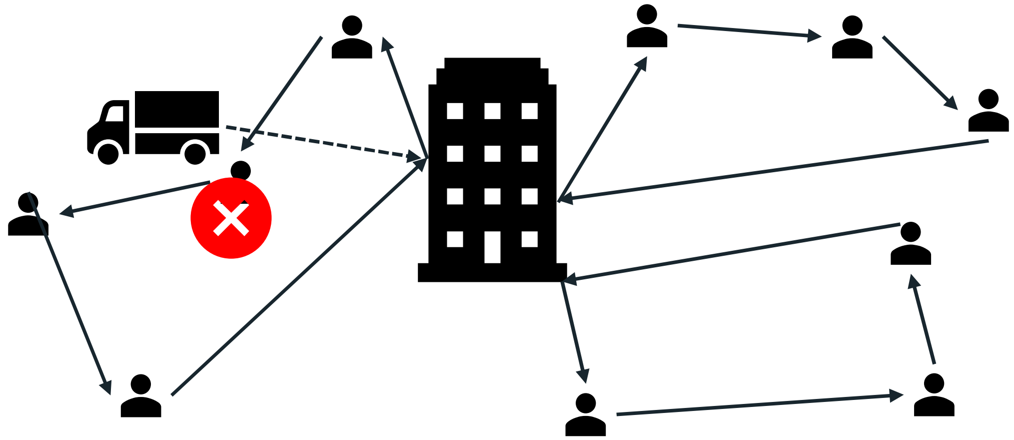

Vehicle Routing with Unknown Customer Demands
(Predicting k)

- Such problems are framed as stochastic optimization problem in the OR community.
  [Birge, J. R., & Louveaux, F.]

Predict-then-Optimize Problem Setup

- In PtO problems, the unknown parameter is predicted using contextual information (features).

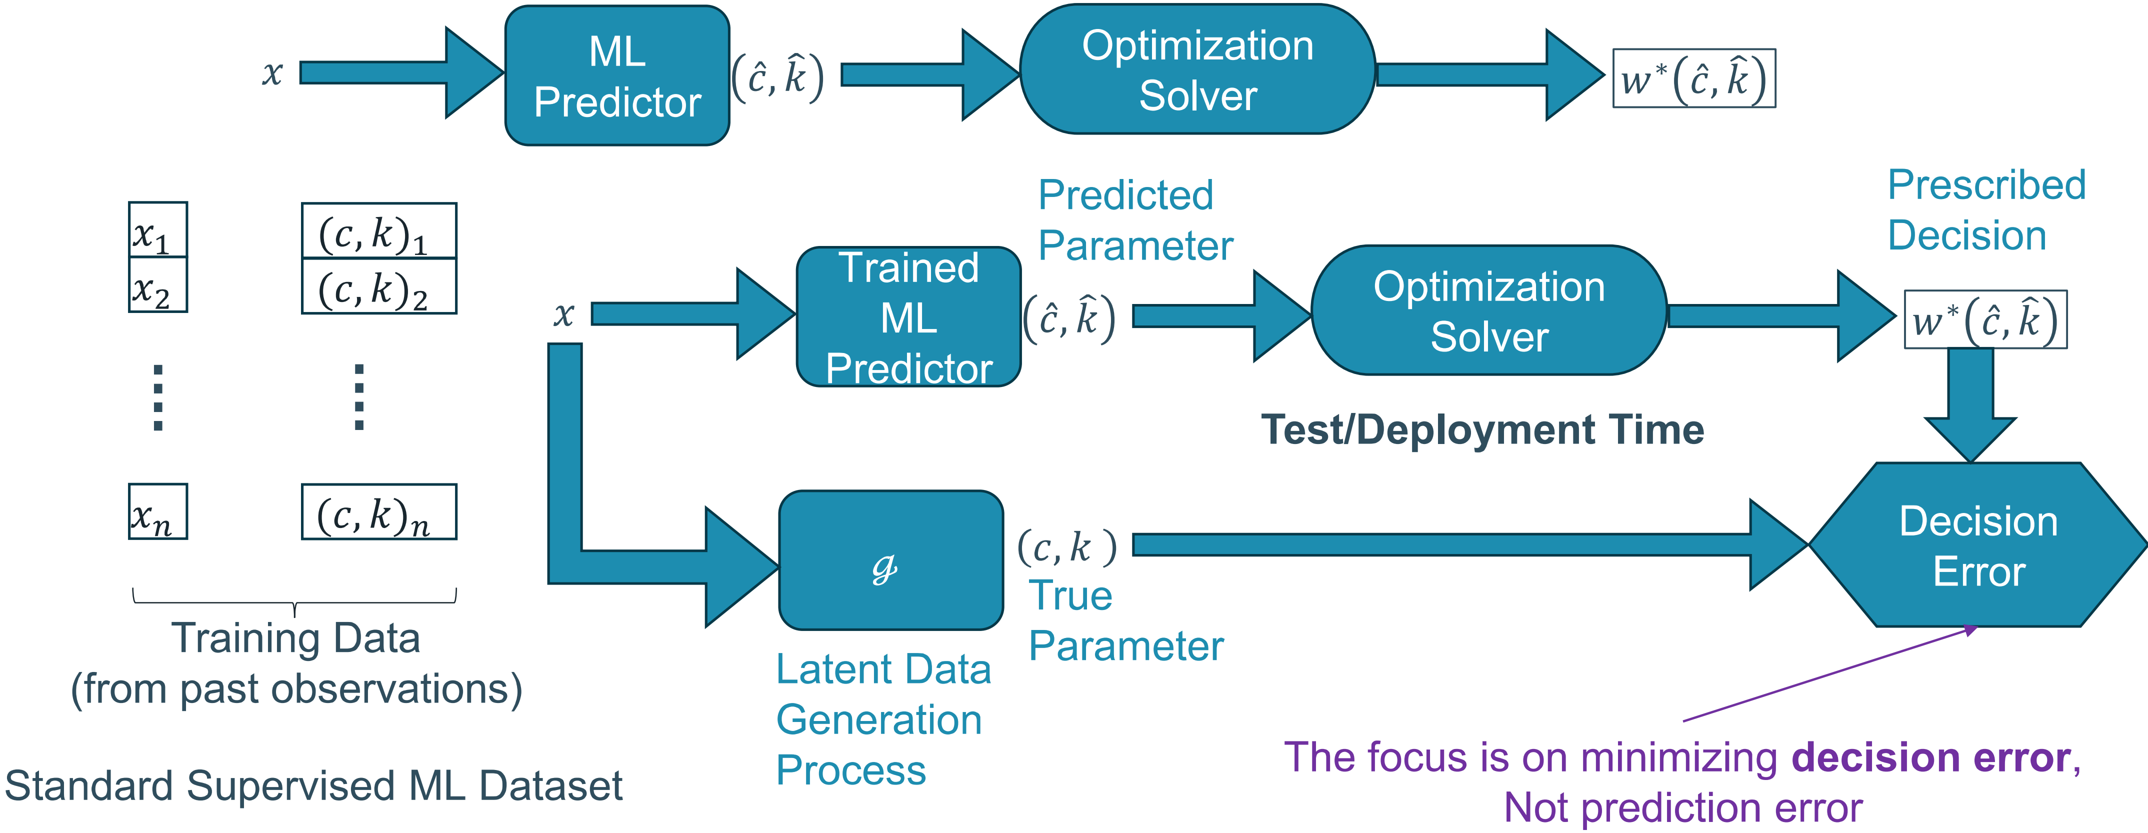

Decision Error

- While Predicting ***only*** c:
  - No uncertainty is associated with $\mathcal{F}$, the feasible set
    - The decision error is relatively easy to evaluate
      - i. *Regret*,
      - ii. Squared error between prescribed and true optimal decision

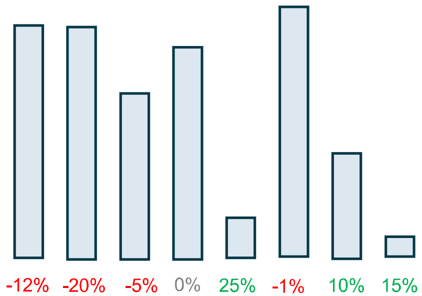

Allocation made using the predictions

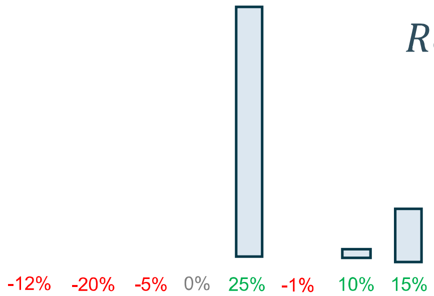

Optimal Allocation if the returns were known

$$Regret = \underbrace{f(c, w^*(\hat{c}, \hat{k}))}_{\text{Objective value with the decision made}} - \underbrace{f(c, w^*(c, k))}_{\text{Objective value if the true parameters were known}}$$

- While Predicting k: Post-hoc Regret[2], Mismatch function[3]

2: Hu, X., Lee, J. C., & Lee, J. H.
3: Paulus, A., Rolínek, M., Musil, V., Amos, B., & Martius, G.

Decision-Focused Learning: The Motivation

Cannot we minimize decision error by minimizing prediction error?

A Very Simple Knapsack Problem

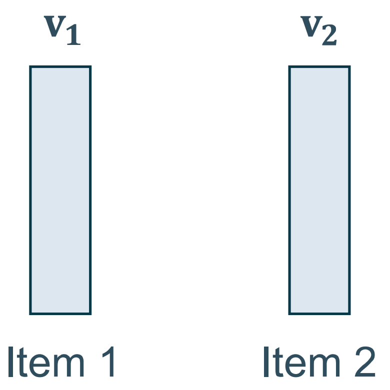

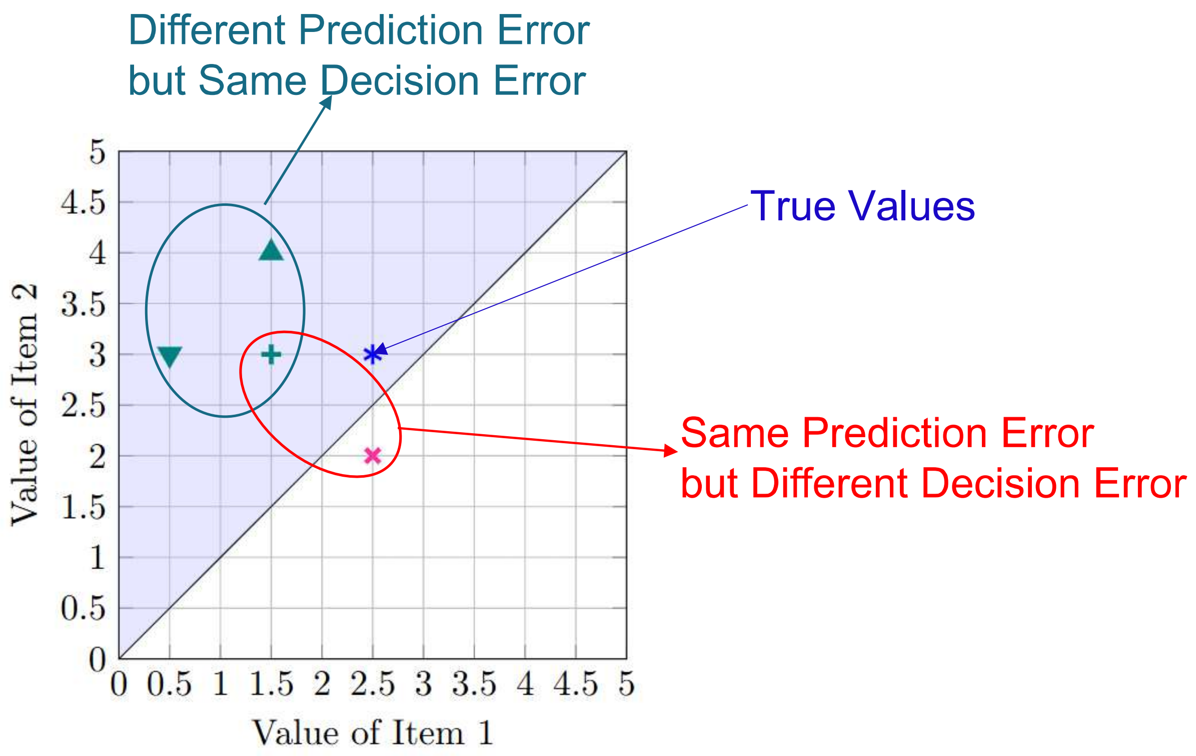

Decision-Focused Learning (for predicting c)

- Decision-focused learning (DFL)[4] ***directly*** trains the ML model to minimize the decision error.
- Due to the recent success of gradient descent-based ML, most of the focus has been to DFL using gradient descent.

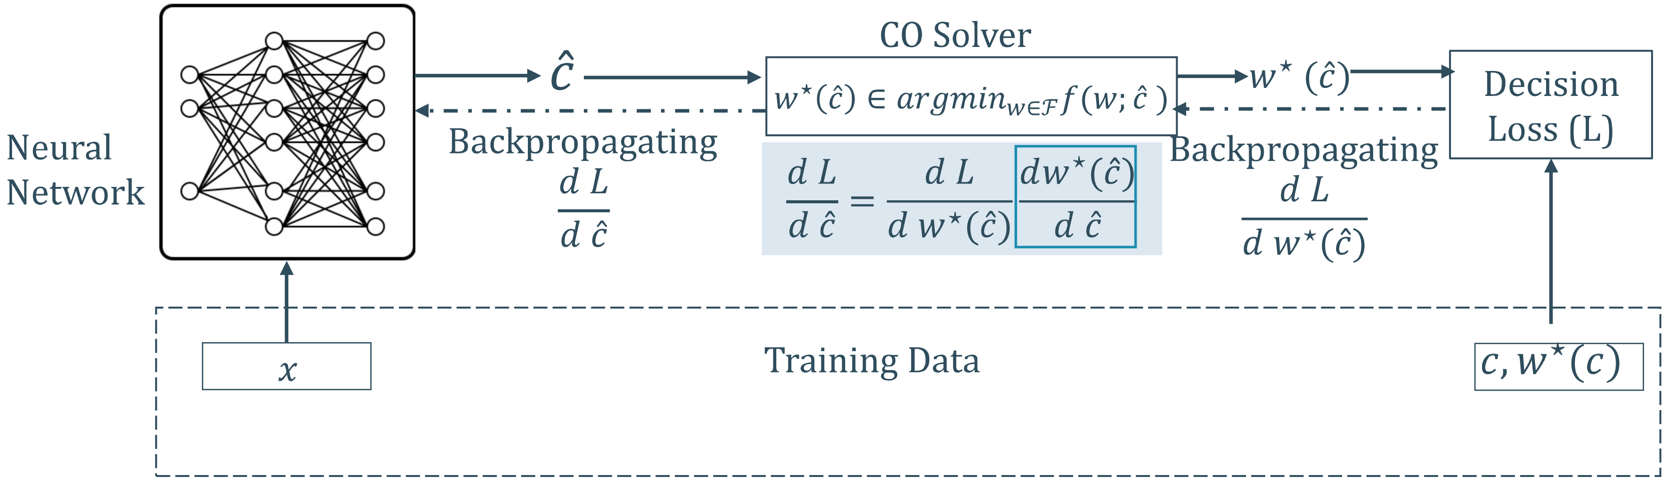

4: Wilder, B., Dilkina, B., & Tambe, M.

Challenges

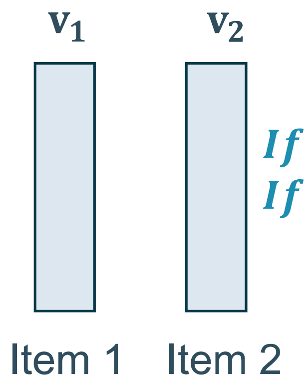

$If\ v_2 < v_1;\ Choose\ Item\ 1$
$If\ v_2 > v_1;\ Choose\ Item\ 2$

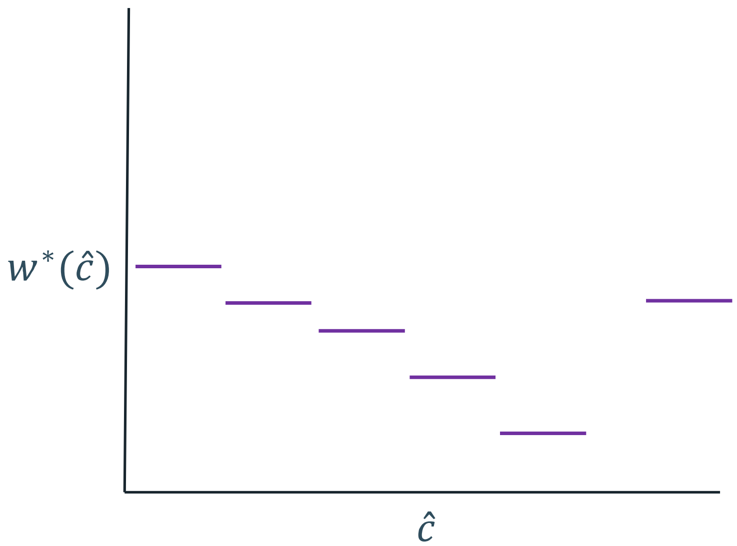

- For combinatorial optimization $\frac{dw^\star(\hat{c})}{d\,\hat{c}}$ is *zero almost everywhere* and does *not exist at the transition points*.

Two DFL Approaches

- Differentiable Layer
- Surrogate Loss

Linear Programs

$$\min c^T w$$
$$s.t.\ Aw = b; w \geq 0$$

Linear Programs (LPs)

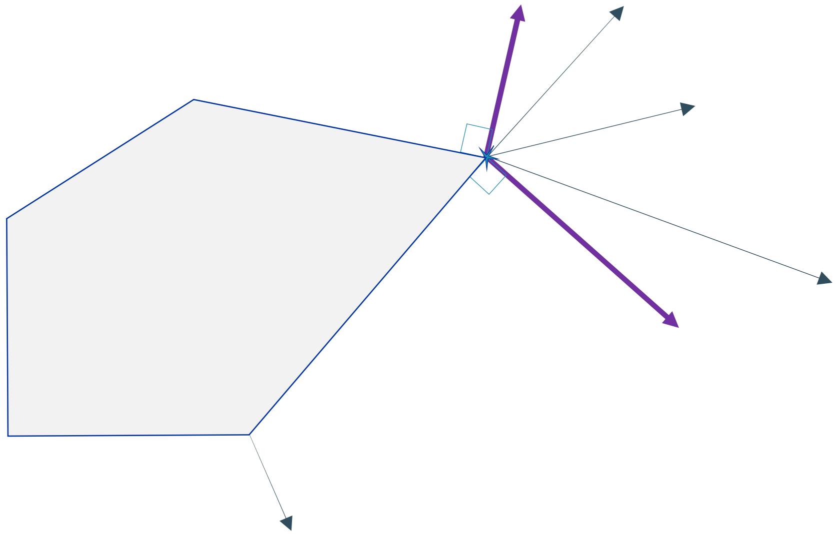

- The LP solution always lies at one of the vertices.

Analytical Smoothing of LP

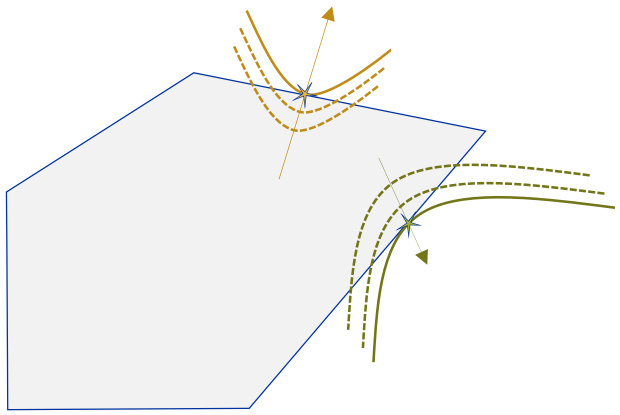

$$\min c^T x + \lambda \, ||x||^2$$
$$s.t. \; Ax = b; x \; \geq 0$$

QP Smoothing

(Wilder, B., Dilkina, B., & Tambe, M., AAAI 2019)

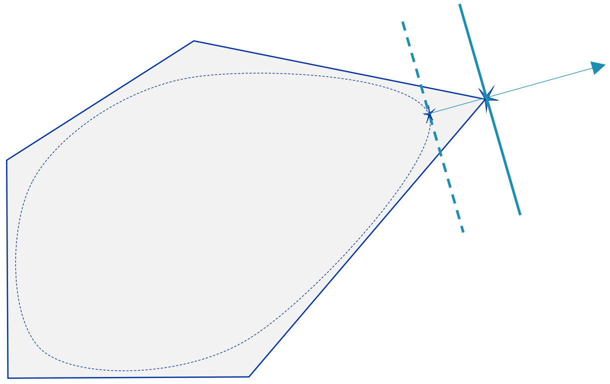

$$\min c^T x - \lambda \sum \ln x_i$$
$$s.t. \; Ax = b; x \; \geq 0$$

Log Barrier Smoothing

Mandi, J., & Guns, T., Neurips 2020)

Integer LP (ILP)

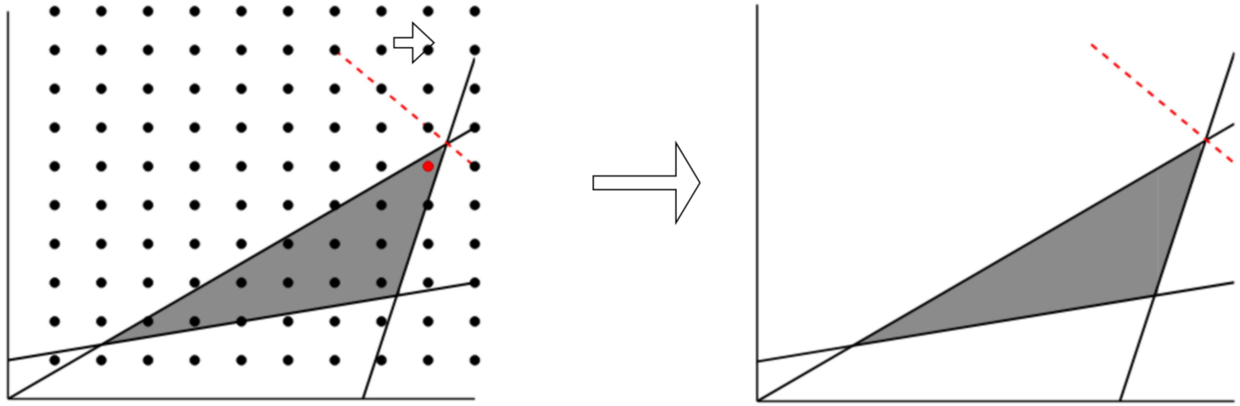

Surrogate Loss

- This approach minimizes a differentiable surrogate loss to reduce expected regret, as the derivative of regret is zero almost everywhere.

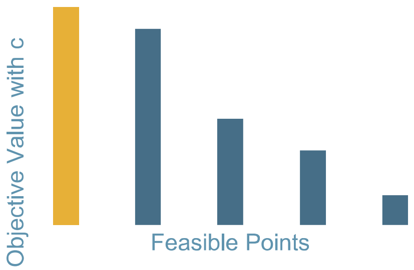

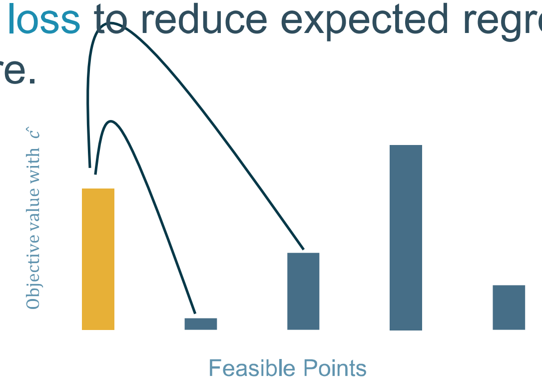

Contrastive Estimation (CE) [Mulamba, Mandi et al., IJCAI 2021]

Predict $\hat{c}$ so that:

$$f\left(\hat{c}, w^*(c)\right) \leq f(\hat{c}, w') \;\; \forall w' \in \mathcal{F}$$

$$\Rightarrow f\left(\hat{c}, w^*(c)\right) \leq \; min_{w'} f(\hat{c}, w') = f(\hat{c}, w^*(\hat{c}))$$

Contrastive Loss: $\boldsymbol{f\left(\hat{c}, w^*(c)\right) - f(\hat{c}, w^*(\hat{c}))}$

This idea has further been extended to develop learning-to-rank loss [Mandi et al, ICML 2022].

Scalability of DFL

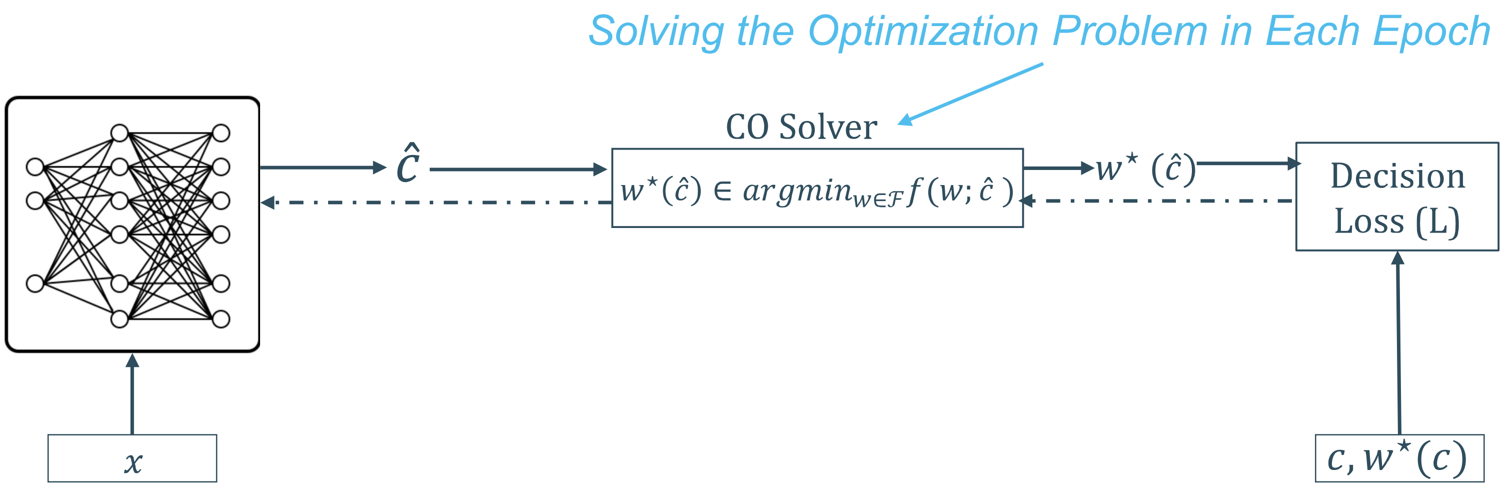

Scalability of DFL

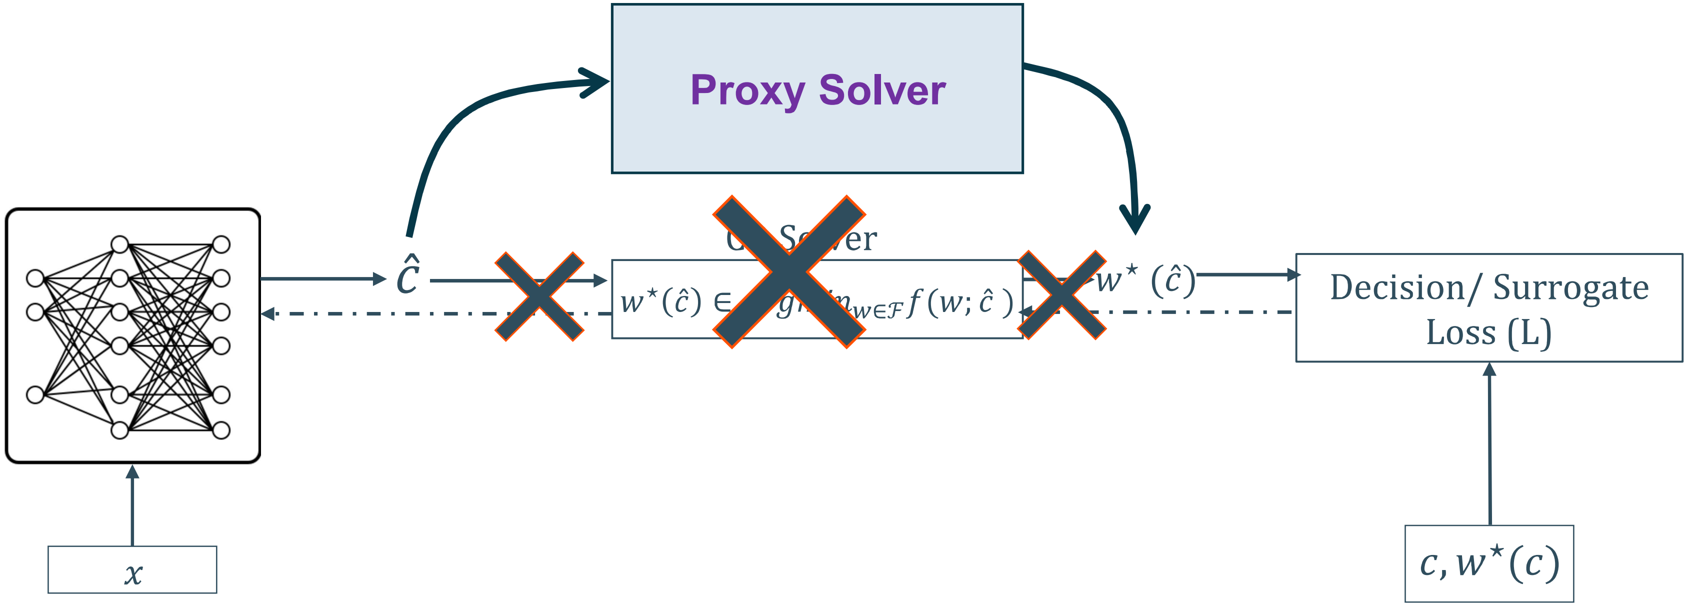

Solution Caching

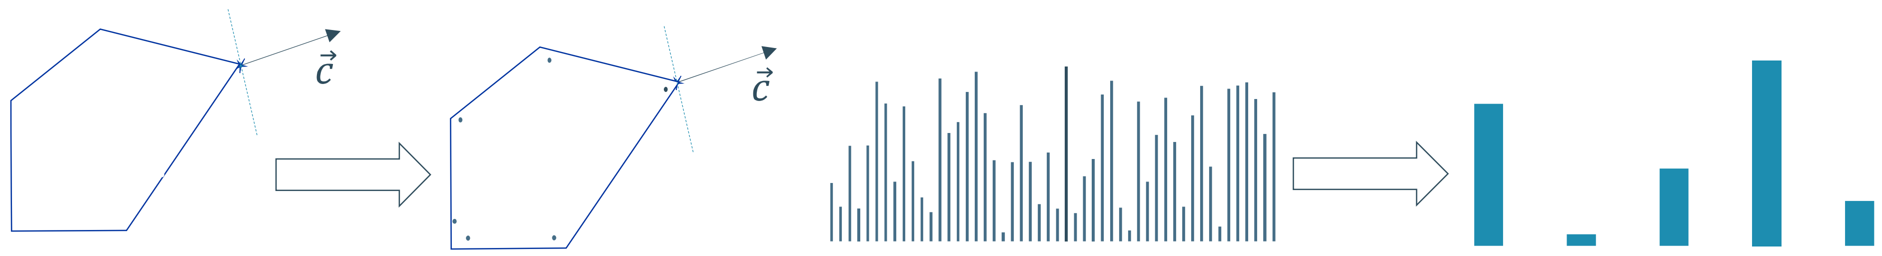

[Mulamba, Mandi et al., IJCAI 2021]

Replacing Solving Optimization problem with a lookup in finite dimensional cache.

Solution caching proves out to be effective in other domains such as planning [Mandi et al., ECAI 2024].

Conclusion

Two broad categories of gradient-based DFL :

i. Differentiable optimization by smoothing
ii. Differentiating surrogate loss function

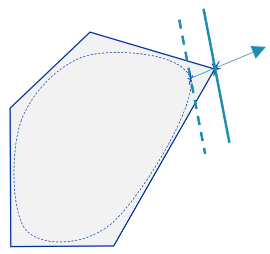

Noise contrastive estimation and learning-to-rank
have been used to devise surrogate loss function.

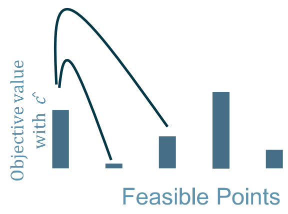

DFL generates predictions with **lower regret** compared to the prediction-focused approach for predicting $\mathbf{c}$.

Solution caching proves out to be effective in addressing scalability.

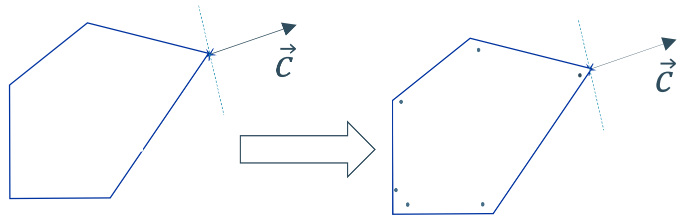

Looking Forward…

- Learning-to-Solve as an optimization proxy
- DFL for uncertainty in the constraints
- Risk-sensitive DFL
- Generalize DFL for related problem
- Real-world applications (more)

Decision-Focused Learning: Foundations, State of the Art, Benchmark and Future Opportunities

Jayanta Mandi JAYANTA.MANDI@KULEUVEN.BE
KU Leuven, Belgium

James Kotary JK4PN@VIRGINIA.EDU
University of Virginia, USA

Senne Berden SENNE.BERDEN@KULEUVEN.BE
KU Leuven, Belgium

Maxime Mulamba MAXIME.MULAMBA@VUB.BE
Vrije Universiteit Brussel, Belgium

Víctor Bucarey VICTOR.BUCAREY@UOH.CL
Universidad de O'Higgins, Chile

Tias Guns TIAS.GUNS@KULEUVEN.BE
KU Leuven, Belgium

Ferdinando Fioretto FIORETTO@VIRGINIA.EDU
University of Virginia, USA

Link to the Paper:

https://doi.org/10.1613/jair.1.15320

Link to the Source code:

References

2. Hu, Xinyi, Jasper CH Lee, and Jimmy HM Lee. "Predict+ Optimize for packing and covering LPs with unknown parameters in constraints." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. No. 4. 2023.
3. Paulus, Anselm, et al. "Comboptnet: Fit the right np-hard problem by learning integer programming constraints." *International Conference on Machine Learning*. PMLR, 2021.
4. Wilder, Bryan, Bistra Dilkina, and Milind Tambe. "Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. No. 01. 2019.
5. Mandi, Jayanta, and Tias Guns. "Interior point solving for lp-based prediction+ optimisation." *Advances in Neural Information Processing Systems* 33 (2020): 7272-7282.
6. Berthet, Quentin, et al. "Learning with differentiable pertubed optimizers." *Advances in neural information processing systems* 33 (2020): 9508-9519.
7. Domke, Justin. "Implicit differentiation by perturbation." *Advances in Neural Information Processing Systems* 23 (2010).
8. Papandreou, George, and Alan L. Yuille. "Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models." *2011 international conference on computer vision*. IEEE, 2011.
9. Niepert, Mathias, Pasquale Minervini, and Luca Franceschi. "Implicit MLE: backpropagating through discrete exponential family distributions." *Advances in Neural Information Processing Systems* 34 (2021): 14567-14579.
10. Pogančić, Marin Vlastelica, et al. "Differentiation of blackbox combinatorial solvers." *International Conference on Learning Representations*. 2020.
11. Sahoo, Subham Sekhar, et al. "Backpropagation through Combinatorial Algorithms: Identity with Projection Works." *The Eleventh International Conference on Learning Representations*.
12. Elmachtoub, Adam N., and Paul Grigas. "Smart “predict, then optimize”." *Management Science* 68.1 (2022): 9-26.
13. Mulamba, Maxime, et al. "Contrastive Losses and Solution Caching for Predict-and-Optimize." *30th International Joint Conference on Artificial Intelligence (IJCAI-21).*
14. Shah, Sanket, et al. "Decision-focused learning without decision-making: Learning locally optimized decision losses." *Advances in Neural Information Processing Systems* 35 (2022): 1320-1332.
15. Mandi, Jayanta, et al. "Decision-focused learning: Through the lens of learning to rank." *International conference on machine learning*. PMLR, 2022.
16. Zharmagambetov, Arman, et al. "Landscape surrogate: Learning decision losses for mathematical optimization under partial information." *Advances in Neural Information Processing Systems* 36 (2024).
17. Kotary, James, et al. "Predict-Then-Optimize by Proxy: Learning Joint Models of Prediction and Optimization." *arXiv preprint arXiv:2311.13087* (2023)
18. Guler, Ali Ugur, et al. "A divide and conquer algorithm for predict+ optimize with non-convex problems." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. No. 4. 2022.
19. Ferber, Aaron, et al. "Mipaal: Mixed integer program as a layer." *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. No. 02. 2020.
20. Elmachtoub, Adam N., Jason Cheuk Nam Liang, and Ryan McNellis. "Decision trees for decision-making under the predict-then-optimize framework." *International conference on machine learning*. PMLR, 2020.
21. Hu, Xinyi, Jasper CH Lee, and Jimmy HM Lee. "Branch & Learn with Post-hoc Correction for Predict+ Optimize with Unknown Parameters in Constraints." *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*.
22. Hu, Xinyi, et al. "Branch & Learn for recursively and iteratively solvable problems in Predict+ Optimize." *Advances in Neural Information Processing Systems* 35 (2022): 25807-25817.
23. Tang, Bo, and Elias B. Khalil. "Multi-task Predict-then-Optimize." *International Conference on Learning and Intelligent Optimization*. Cham: Springer International Publishing, 2023.
24. Silvestri, Mattia, et al. "Score Function Gradient Estimation to Widen the Applicability of Decision-Focused Learning." *arXiv preprint arXiv:2307.05213* (2023).
25. Mandi, Jayanta, et al. "Decision-Focused Learning to Predict Action Costs for Planning." *(ECAI 2024)*.

Thank You.
Questions?