Modeling Robustness in Decision-Focused Learning as a Stackelberg Game

Harvard John A. Paulson School of Engineering and Applied Sciences

Sonja Johnson-Yu[1], Kai Wang[1], Jessie Finocchiaro[1], Aparna Taneja[2], Milind Tambe[1,2]

[1]*Harvard University*, [2]*Google Research*

Introduction

- *Predict-then-Optimize* allows smart decision-making by first estimating missing parameters and then optimizing over these predicted parameters.

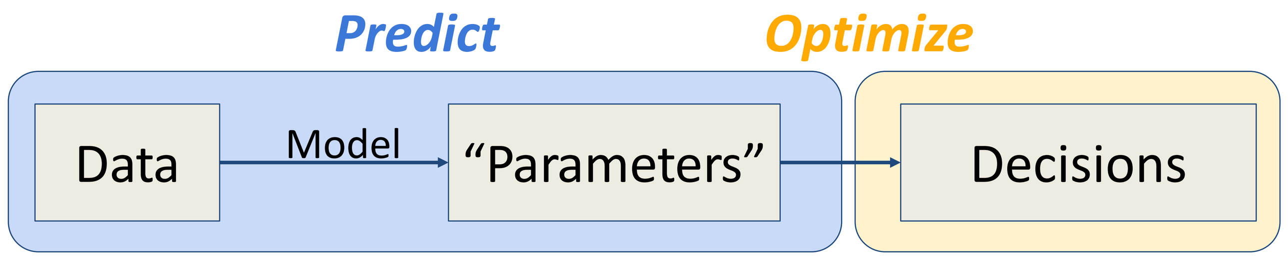

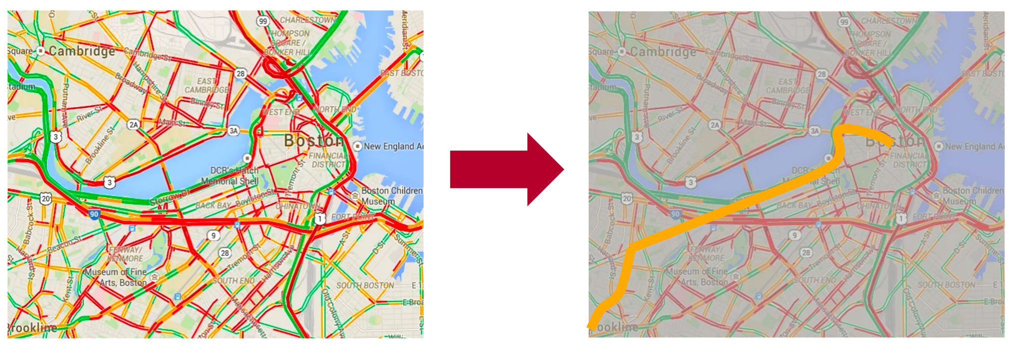

e.g. route planning: *predict* travel times and then *optimize* for shortest path

- *Decision-focused learning* makes predict-then-optimize an end-to-end learning pipeline by differentiating through the optimization to train the predictive model.

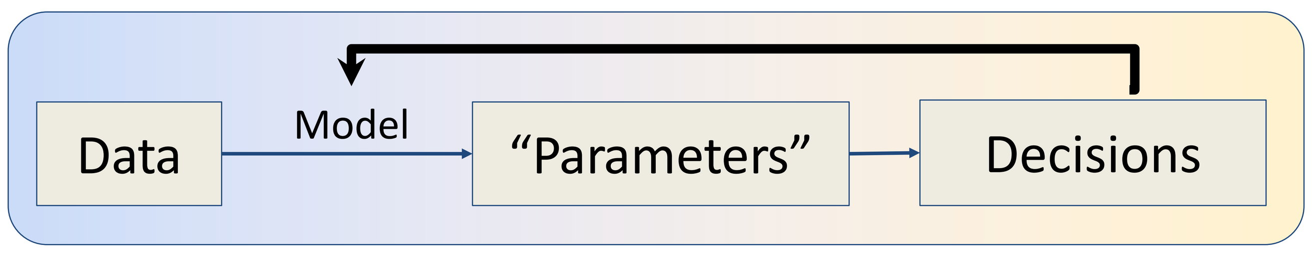

Traditional: Get as many predictions as right as possible! (MSE)

But: It's okay to have lower accuracy in these predictions…

As long as you get these predictions correct!

- What happens when there is **noise in the labels**?
→ can hurt the downstream optimization!

Optimizing over inaccurate predictions results in suboptimal decisions!

Background

"Two-Stage" Traditional Predict-then-Optimize

$$\min_{w} \sum_{(x,y)\in\mathcal{D}_{\text{train}}} \|m_w(x) - y\|_2$$

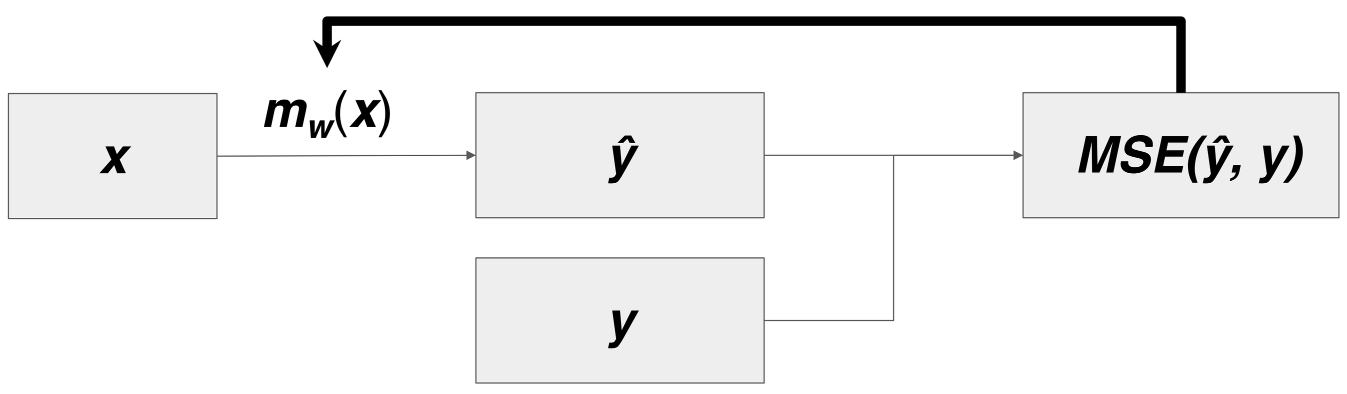

Decision-Focused Learning

$$\max_{w} \sum_{(x,y)\in\mathcal{D}_{\text{train}}} f(z^*(m_w(x)), y)$$

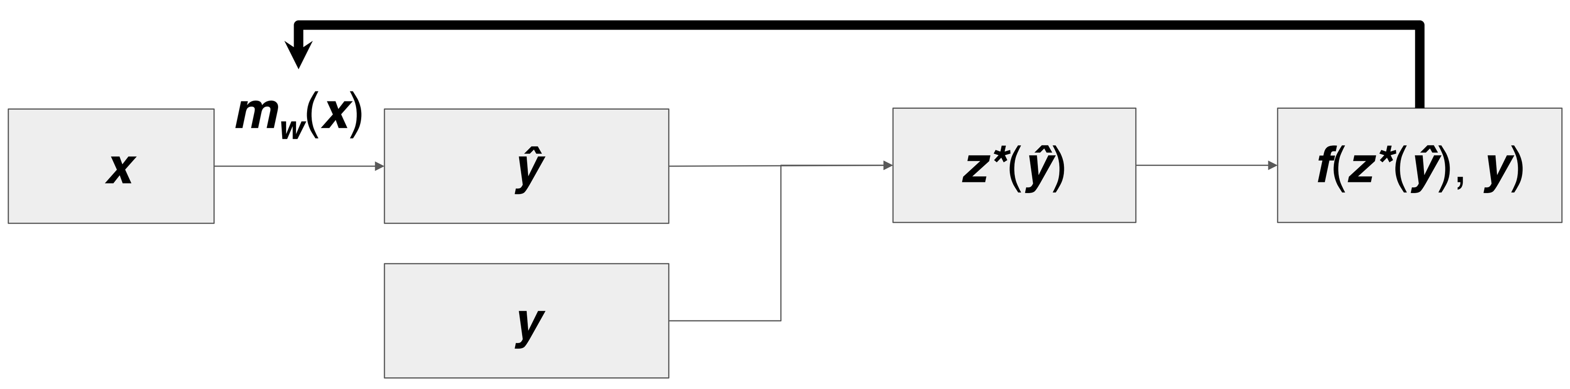

Contributions

- We propose robust variants of traditional two-stage Predict-then-Optimize and decision-focused learning.

Robust Two-Stage

$$\min_{w} \sum_{(x,y)\in\mathcal{D}} \|m_w(x) - (y + \epsilon^*_{x,y})\|_2$$
$$\text{s.t.} \quad \epsilon^*_{x,y} = \arg\min_{\epsilon:\|\epsilon\|\le r} \quad f(z^*(m_w(x)), y + \epsilon)$$

Robust Decision-Focused Learning

$$\max_{w} \sum_{(x,y)\in\mathcal{D}_{\text{train}}} f(z^*(m_w(x)), y + \epsilon^*_{x,y})$$
$$\text{s.t.} \quad \epsilon^*_{x,y} = \arg\min_{\epsilon:\|\epsilon\|\le r} \quad f(z^*(m_w(x)), y + \epsilon)$$

- We show that these robust algorithms can be viewed through the lens of Stackelberg games.
- We use this framework to provide regret bounds for the robust algorithms.
- We show that robust decision-focused learning outperforms robust traditional Predict-then-Optimize.

Acknowledgements. S.J.Y. was sponsored by the ARO under Grant Number: W911NF-18-1-0208. J.F. was supported by the National Science Foundation under Award No. 2202898.