Inverse Optimization: An Efficient Learning Framework for Complex Behaviors

Peyman Mohajerin Esfahani

University of Toronto & TU Delft

IEEE Conference on Decision and Control

December 2025

Outline

- Supervised Learning
- Inverse Optimization
- Applications

Outline

- Supervised Learning
  - Two challenges in functions approximation
- Inverse optimization
- Applications

Supervised Learning

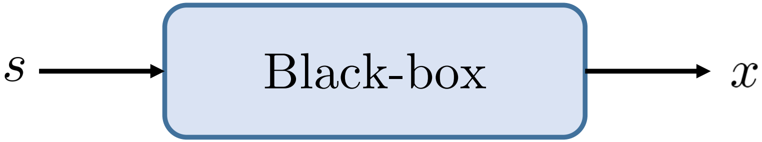

Supervised Learning

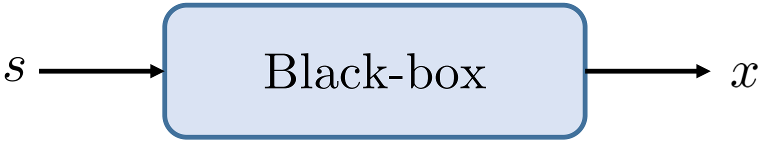

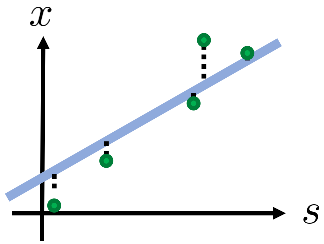

Supervised Learning

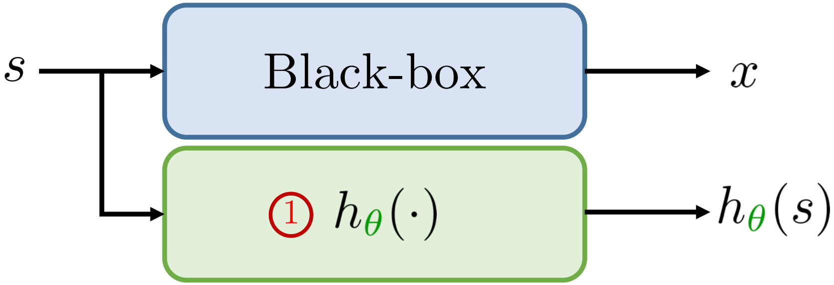

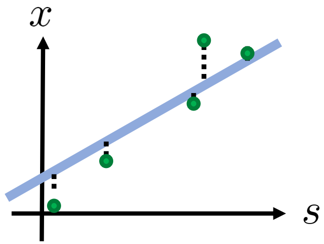

Supervised Learning

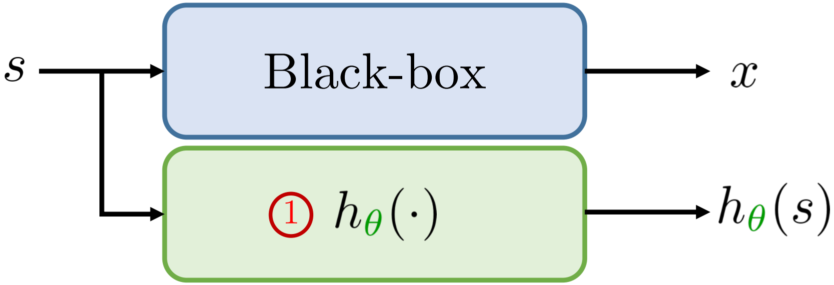

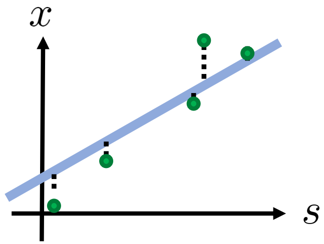

$$h_{\theta}(s) = \theta_1 s + \theta_0 1$$

Supervised Learning

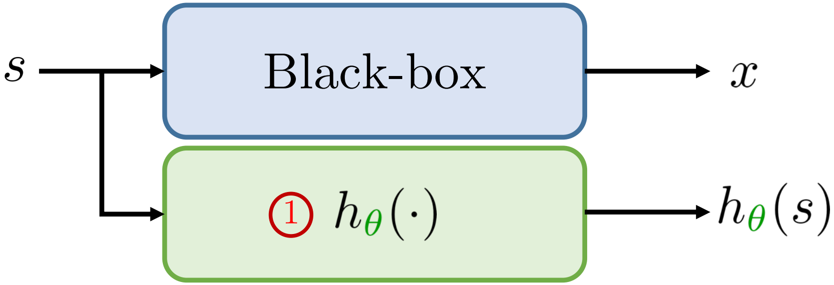

2 $\ell(x, h_\theta(s))$

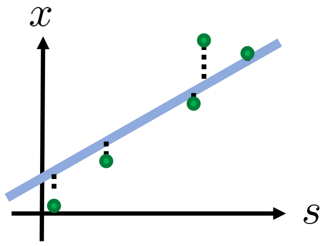

$$h_\theta(s) = \theta_1 s + \theta_0 1$$

Supervised Learning

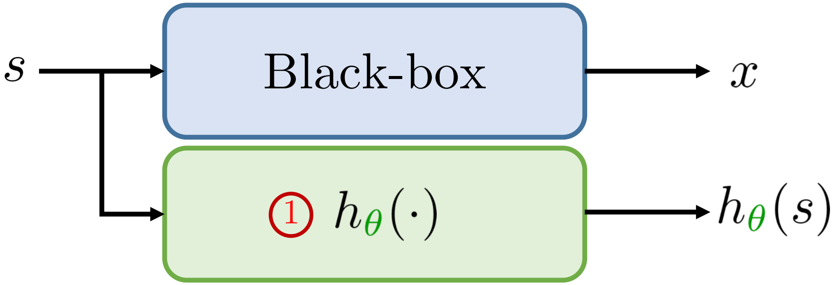

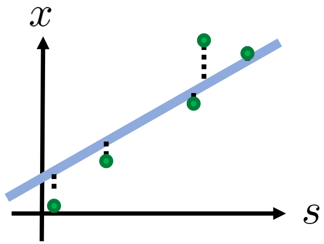

2 $\ell(x, h_\theta(s)) = \|x - h_\theta(s)\|^2$

$h_\theta(s) = \theta_1 s + \theta_0 1$

Supervised Learning

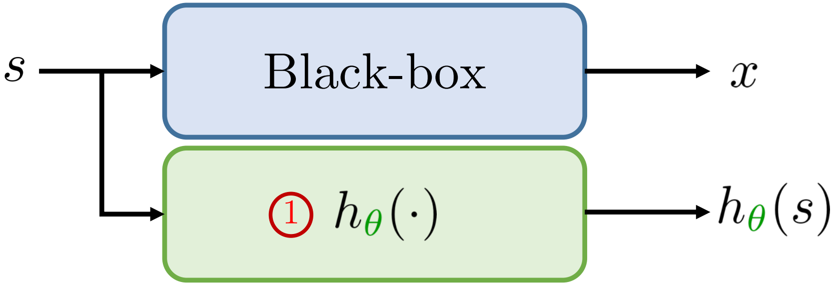

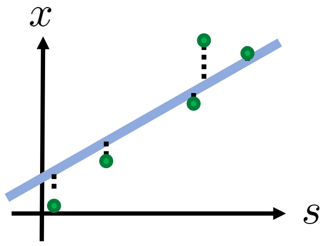

2 $\ell\big(x, h_\theta(s)\big) = \big\|x - h_\theta(s)\big\|^2$

$h_\theta(s) = \theta_1 s + \theta_0 1$

$$\text{Training} \begin{cases} \text{Data } \big\{(\widehat{s}_i, \widehat{x}_i)\big\}_{i\leq N} \\ \min\limits_{\theta} \sum\limits_{i\leq N} \ell\big(\widehat{x}_i, h_\theta(\widehat{s}_i)\big) \end{cases}$$

Supervised Learning

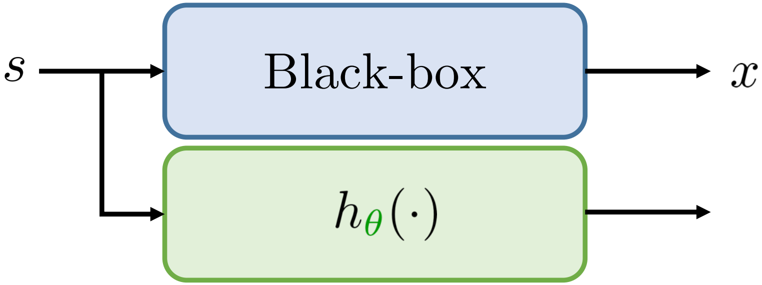

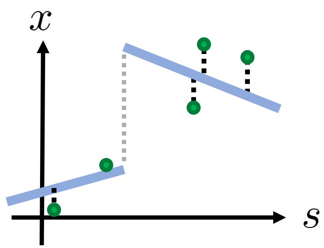

Supervised Learning

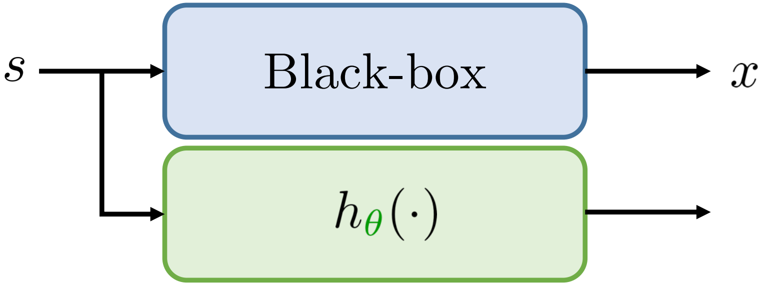

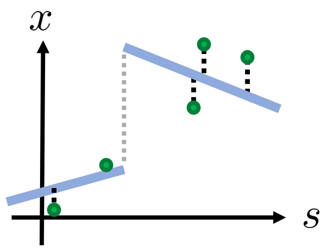

$$h_{\theta}(s) = \theta_1 \phi_1(s) + \cdots + \theta_k \phi_k(s)$$

Supervised Learning

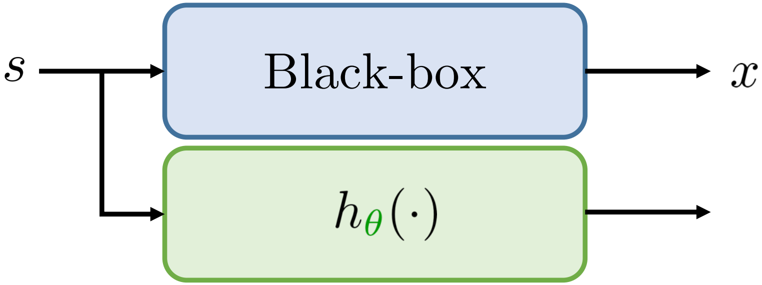

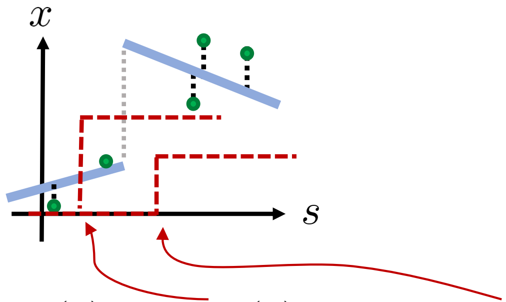

$$h_\theta(s) = \theta_1 \phi_1(s) + \cdots + \theta_k \phi_k(s)$$

Supervised Learning

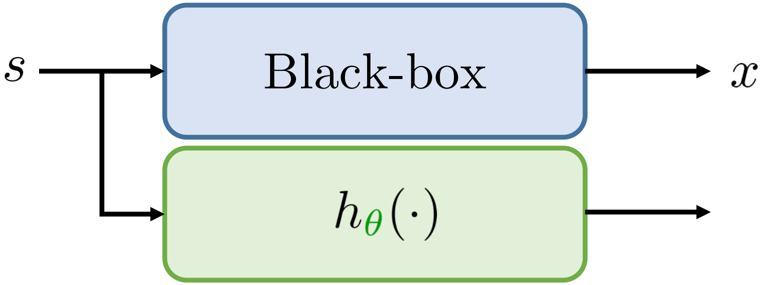

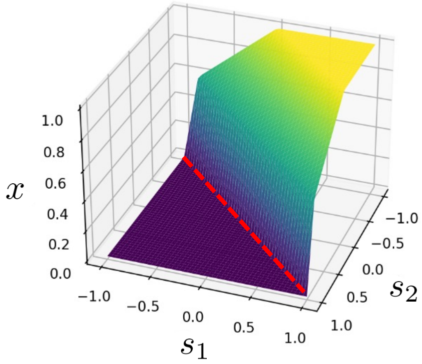

Supervised Learning

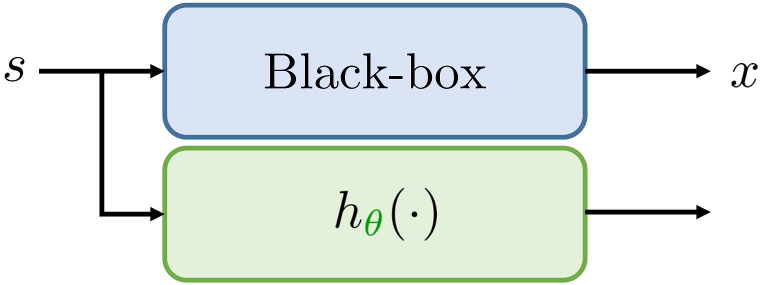

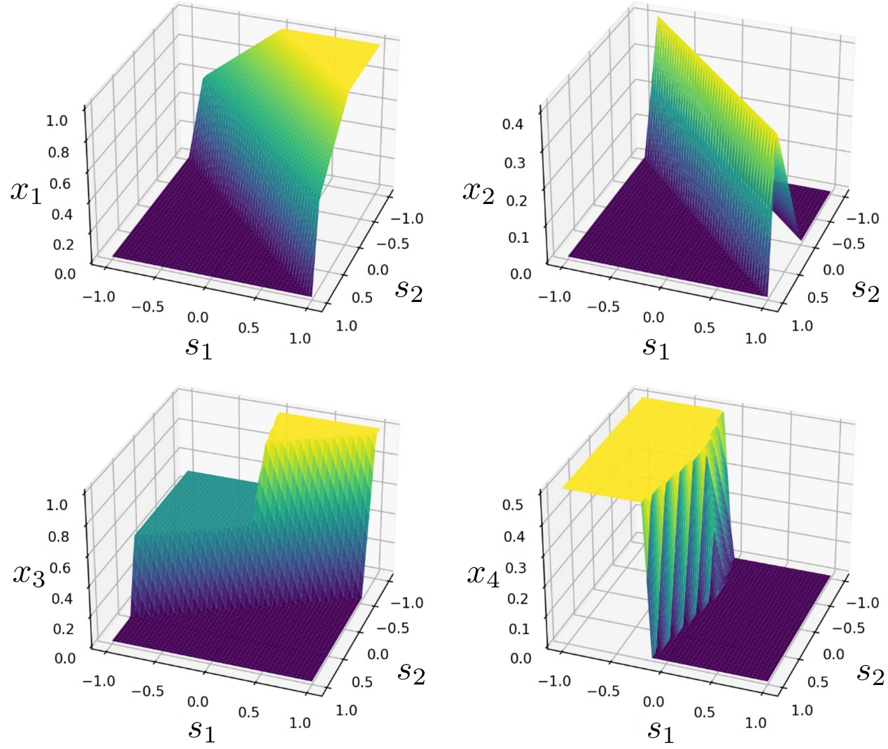

Outline

- Supervised learning

- Inverse Optimization

  - A <u>rich</u> model with a <u>convex</u> training loss

- Applications

Inverse Optimization

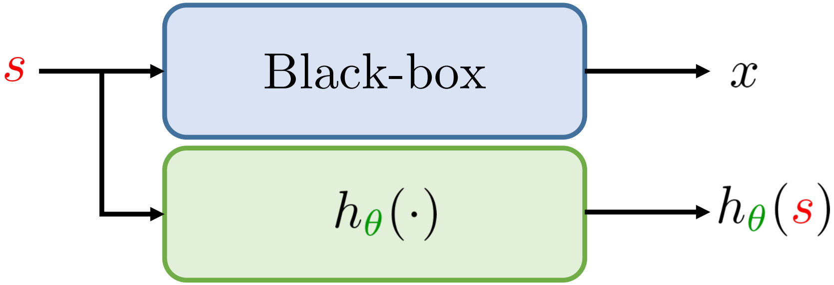

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

Inverse Optimization

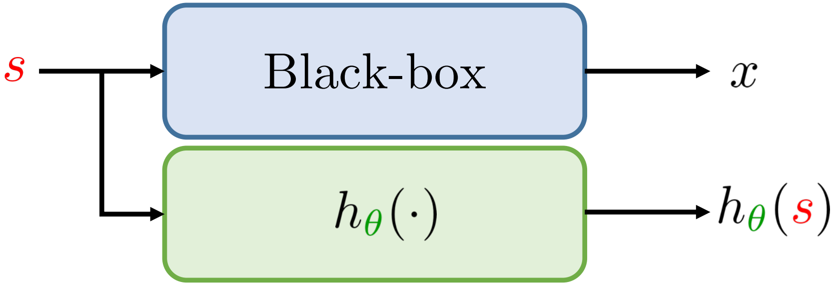

$$
\begin{aligned}
h_\theta(s) &= \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y) \\
&= \arg\min_{y \in \mathbb{X}(s)} y^\top \Theta_2 y + y^\top \Theta_1 s
\end{aligned}
$$

Inverse Optimization

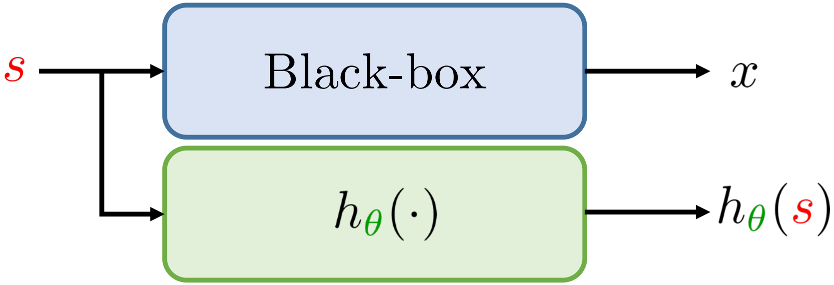

$$
\begin{aligned}
h_{\theta}(s) &= \arg\min_{y \in \mathbb{X}(s)} F_{\theta}(s, y) \\
&= \arg\min_{y \in \mathbb{X}(s)} y^{\top}\Theta_2 y + y^{\top}\Theta_1 s
\end{aligned}
$$

output
information

Inverse Optimization

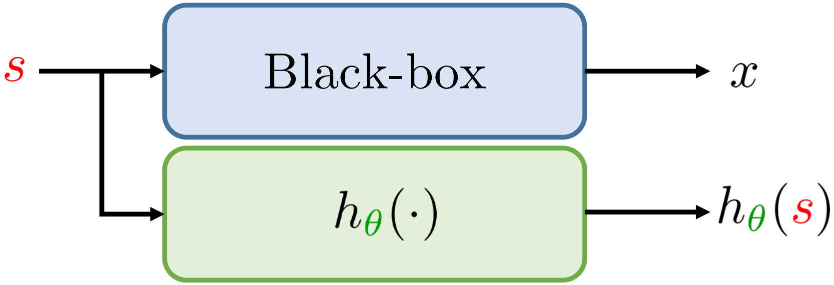

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

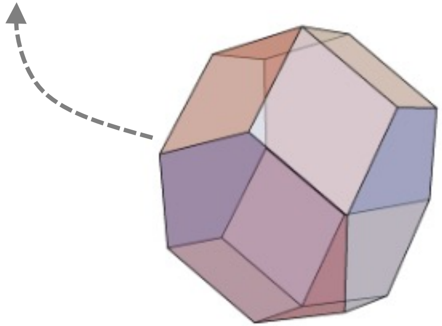

Inverse Optimization

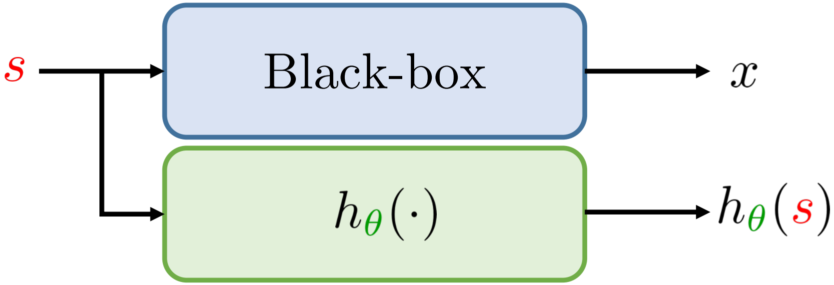

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}} \langle \Theta s, y \rangle$$

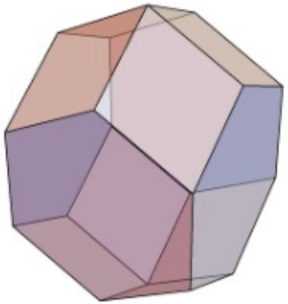

Inverse Optimization

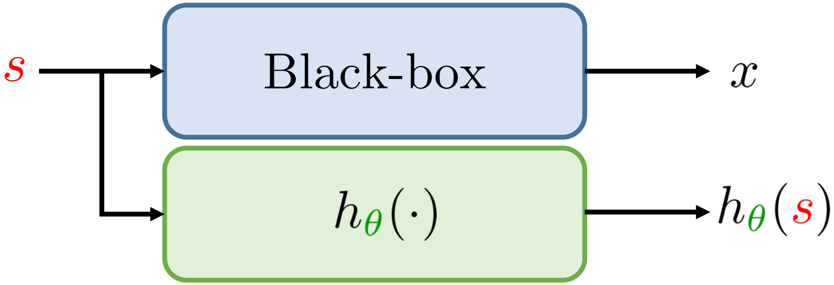

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}} \langle \Theta s, y \rangle$$

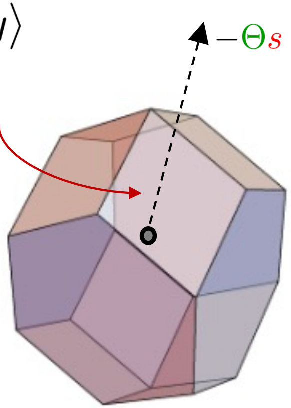

Inverse Optimization

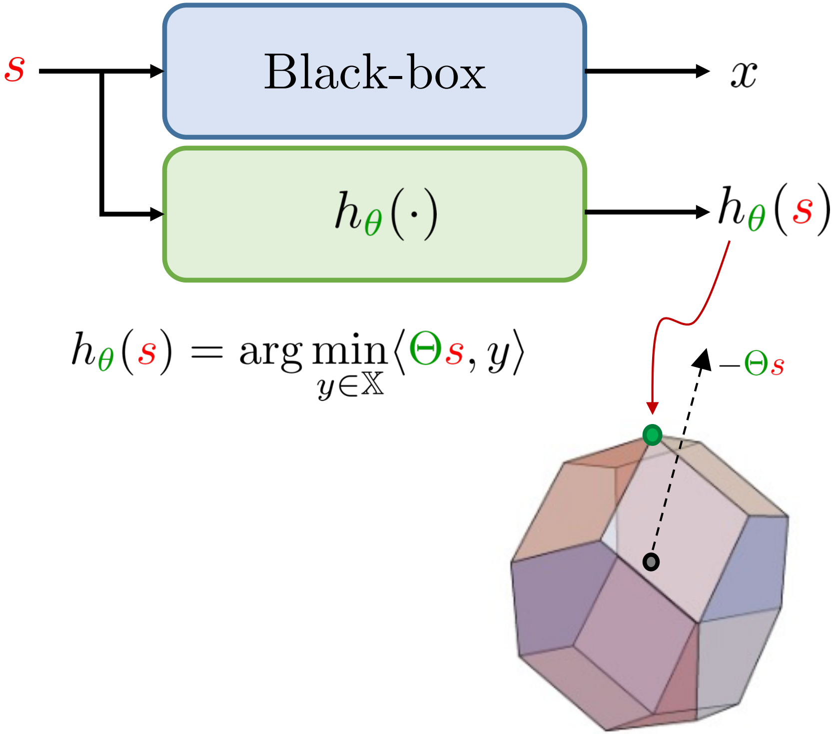

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}} \langle \Theta s, y \rangle$$

Inverse Optimization

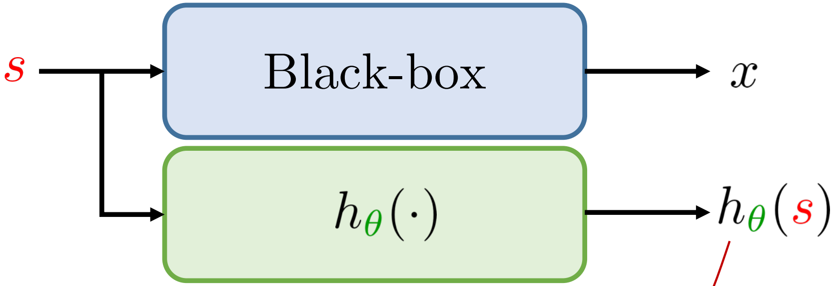

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}} \langle \Theta s, y \rangle$$

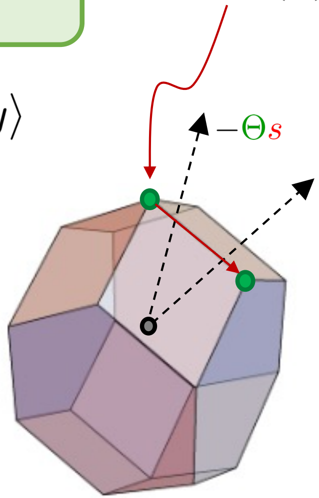

Inverse Optimization

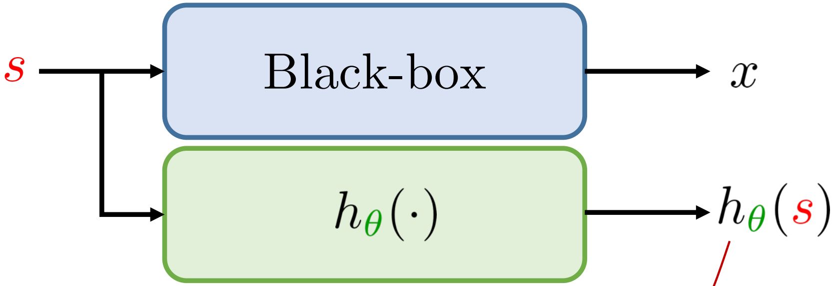

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}} \langle \Theta s, y \rangle$$

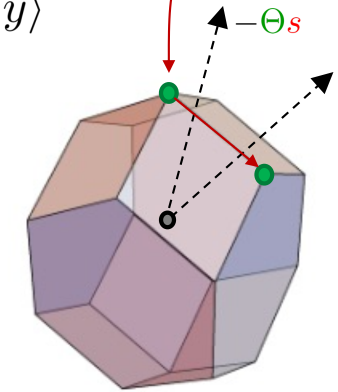

Polynomial representation
exponential vertices (discontinuities)

Inverse Optimization

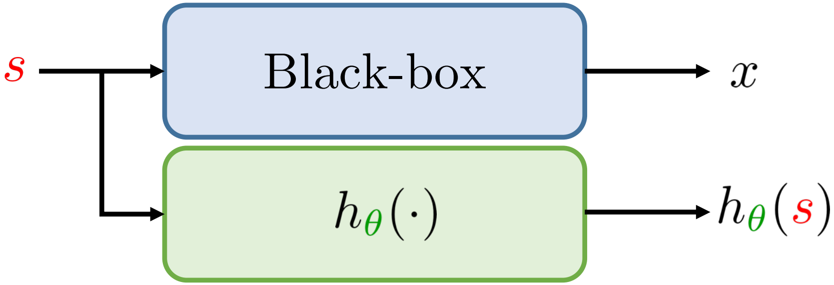

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\big(x, h_\theta(s)\big) = \big\|x - h_\theta(s)\big\|^2$$

Inverse Optimization

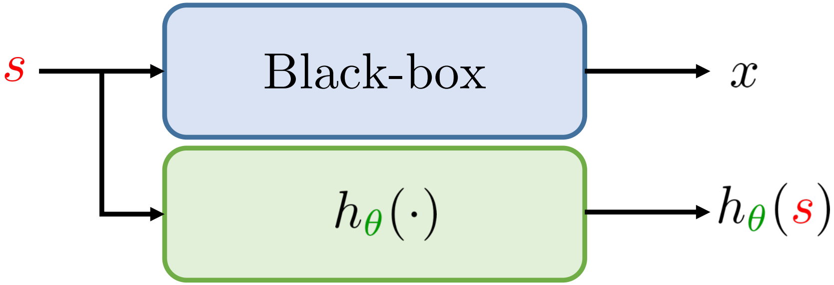

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\left(x, h_\theta(s)\right) = \left\|x - h_\theta(s)\right\|^2$$

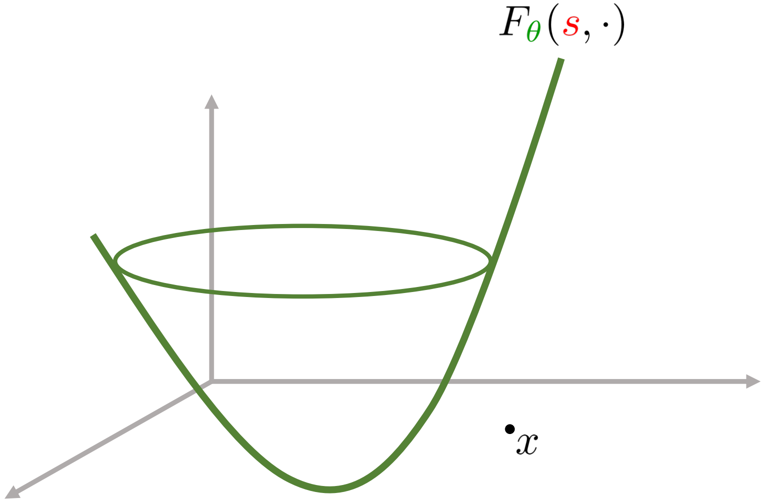

Inverse Optimization

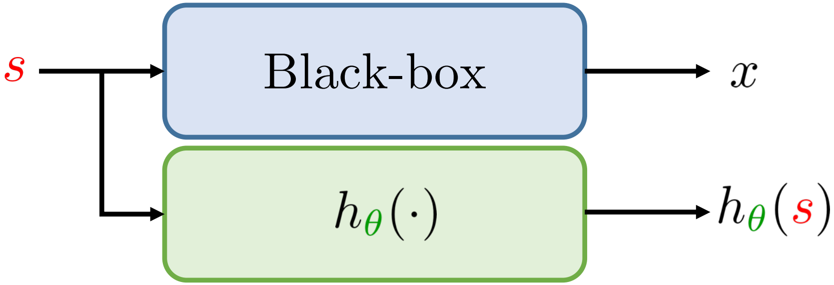

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\left(x, h_\theta(s)\right) = \left\|x - h_\theta(s)\right\|^2$$

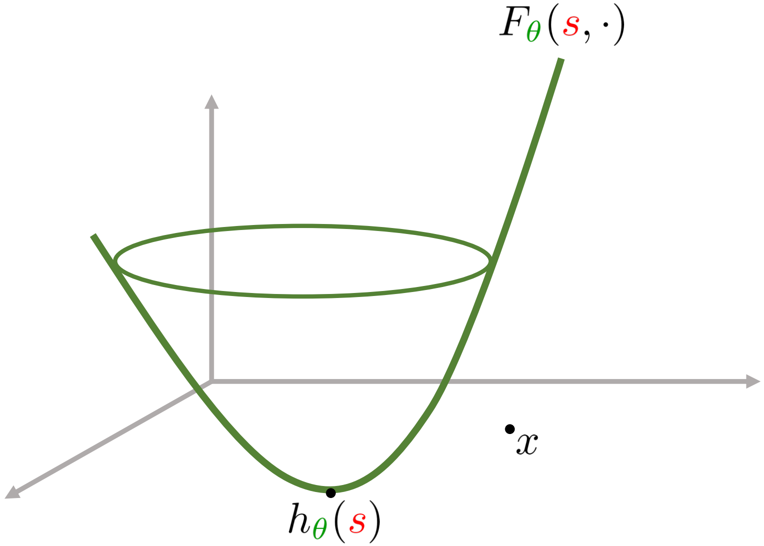

Inverse Optimization

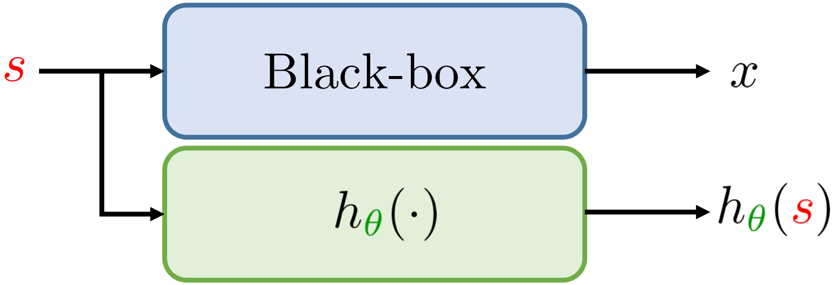

$$h_\theta(s) = \arg \min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\left(x, h_\theta(s)\right) = \left\|x - h_\theta(s)\right\|^2$$

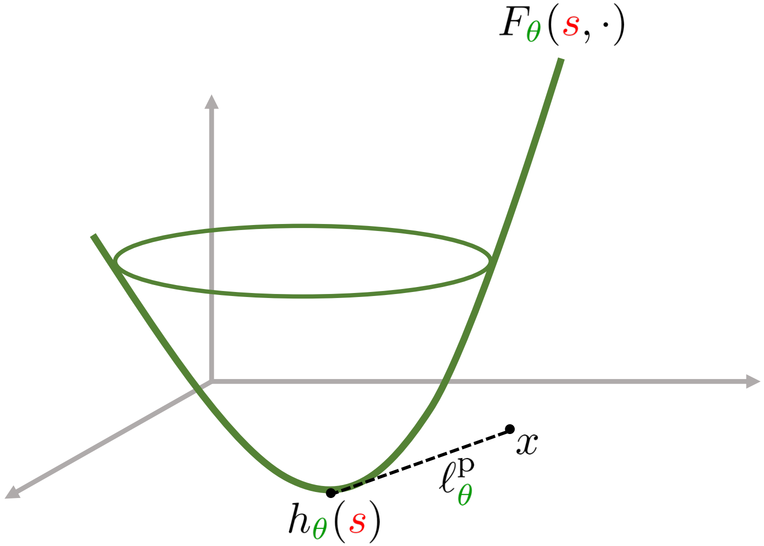

Inverse Optimization

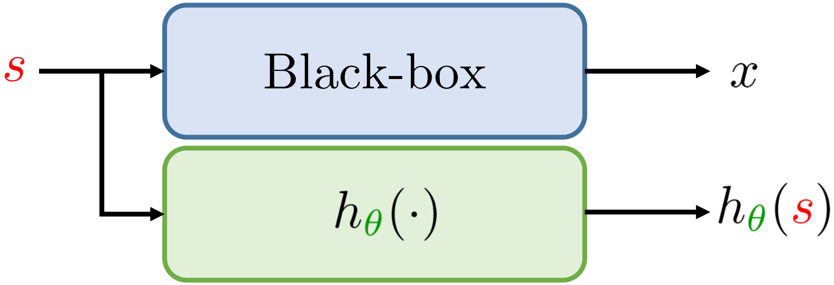

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\left(x, h_\theta(s)\right) = \left\| x - h_\theta(s) \right\|^2$$

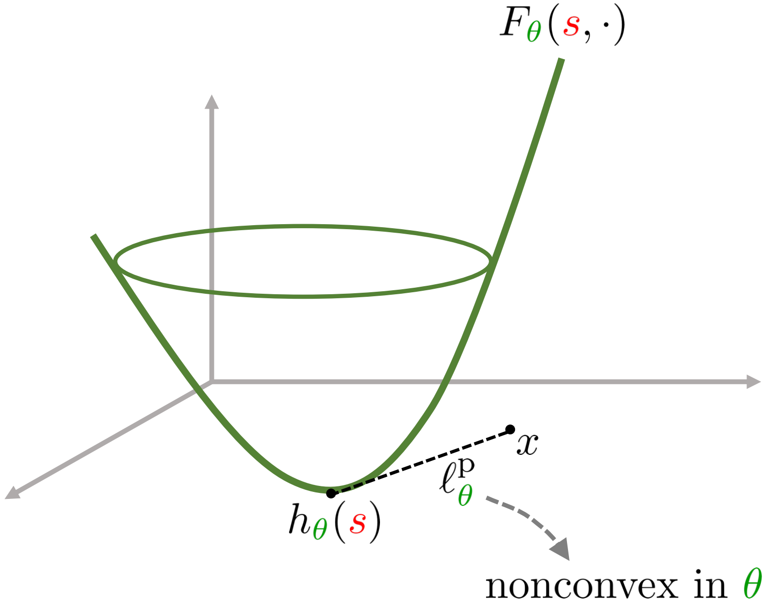

Inverse Optimization

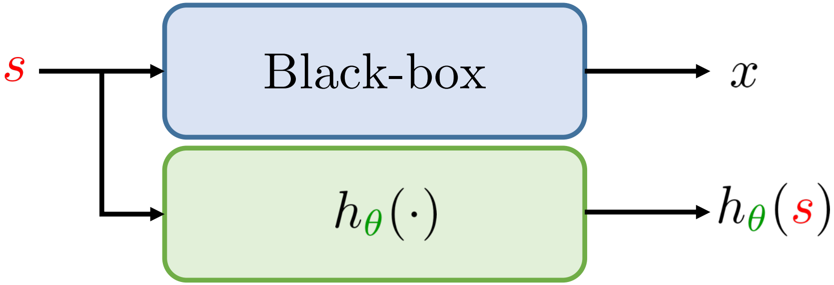

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\big(x, h_\theta(s)\big) = \big\|x - h_\theta(s)\big\|^2$$

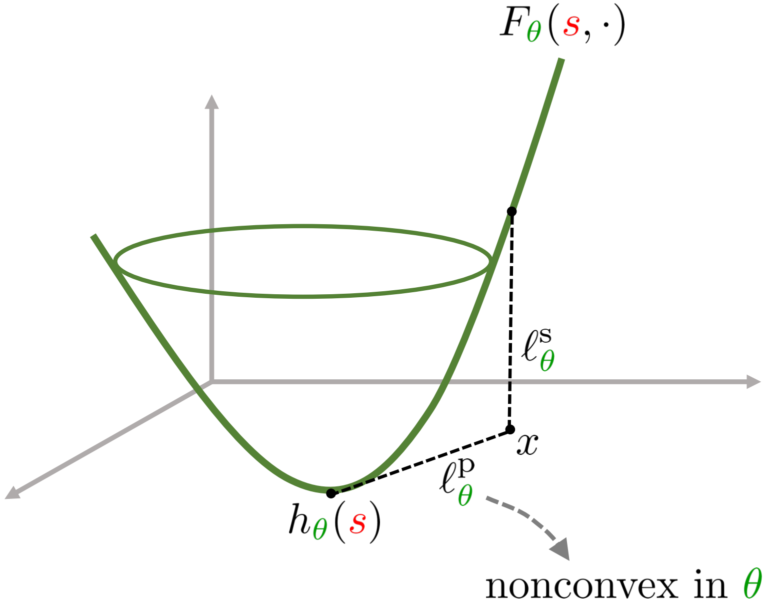

Inverse Optimization

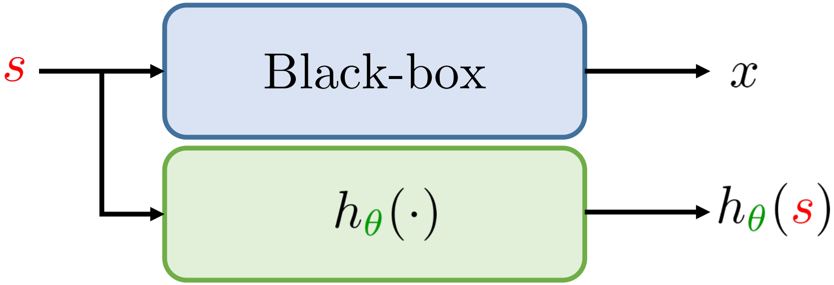

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\left(x, h_\theta(s)\right) = \left\|x - h_\theta(s)\right\|^2$$

$$\ell^{\mathrm{s}}\left(x, h_\theta(s)\right) = F_\theta(s, x) - F_\theta\big(s, h_\theta(s)\big)$$

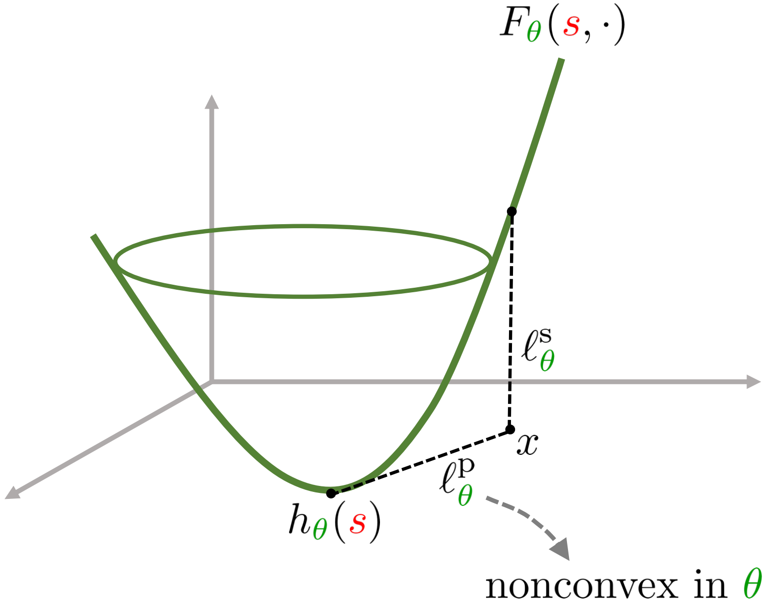

Inverse Optimization

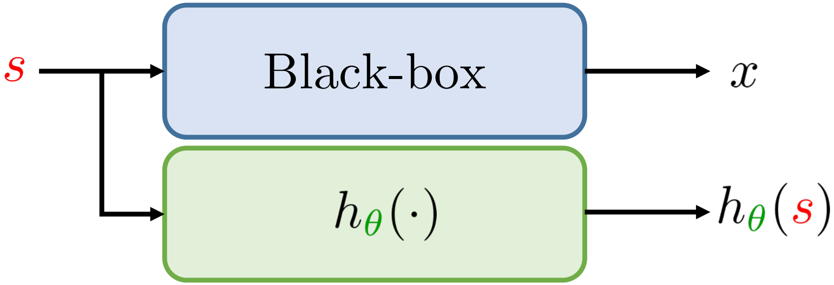

$$h_\theta(s) = \arg\min_{y\in\mathbb{X}(s)} F_\theta(s,y)$$

$$\ell^{\mathrm{p}}\big(x, h_\theta(s)\big) = \big\|x - h_\theta(s)\big\|^2$$

$$\ell^{\mathrm{s}}\big(x, h_\theta(s)\big) = F_\theta(s,x) - F_\theta\big(s, h_\theta(s)\big)$$

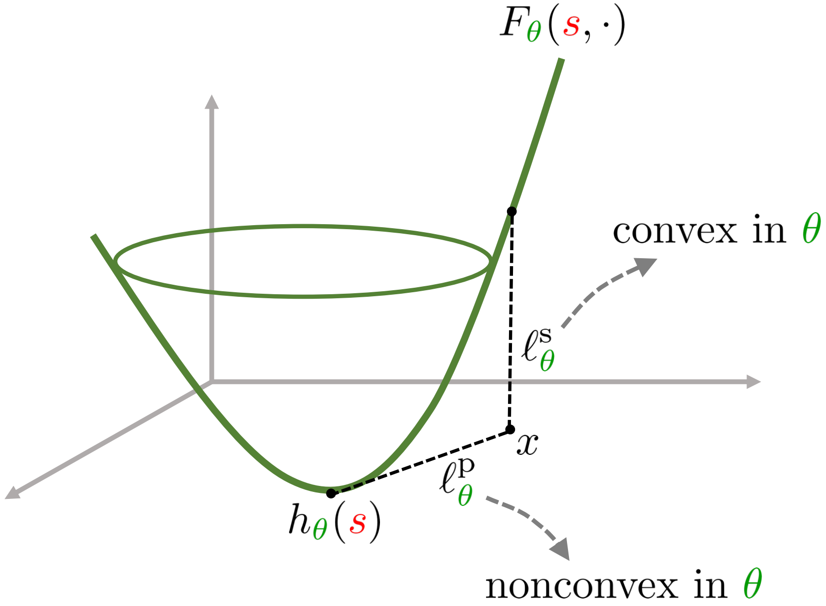

Inverse Optimization

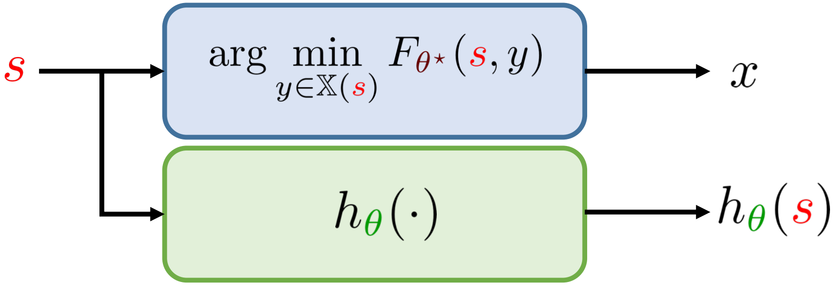

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\big(x, h_\theta(s)\big) = \big\|x - h_\theta(s)\big\|^2$$

$$\ell^{\mathrm{s}}\big(x, h_\theta(s)\big) = F_\theta(s, x) - F_\theta\big(s, h_\theta(s)\big)$$

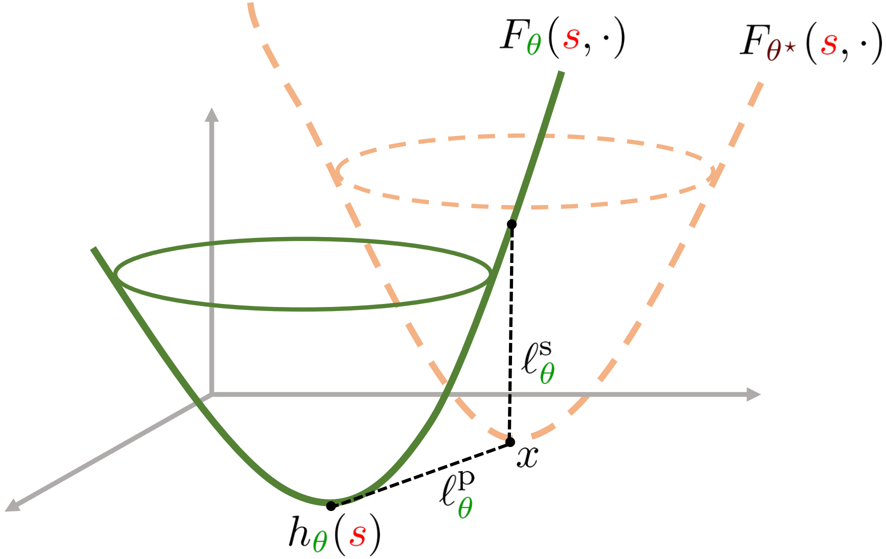

Inverse Optimization

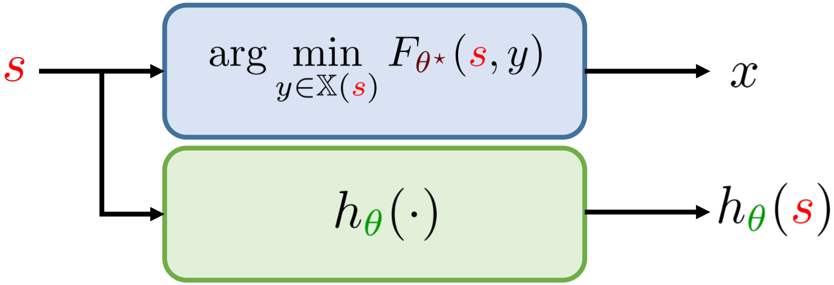

$$h_\theta(s) = \arg\min_{y \in \mathbb{X}(s)} F_\theta(s, y)$$

$$\ell^{\mathrm{p}}\big(x, h_\theta(s)\big) = \big\|x - h_\theta(s)\big\|^2$$

$$\ell^{\mathrm{s}}\big(x, h_\theta(s)\big) = F_\theta(s, x) - F_\theta\big(s, h_\theta(s)\big)$$

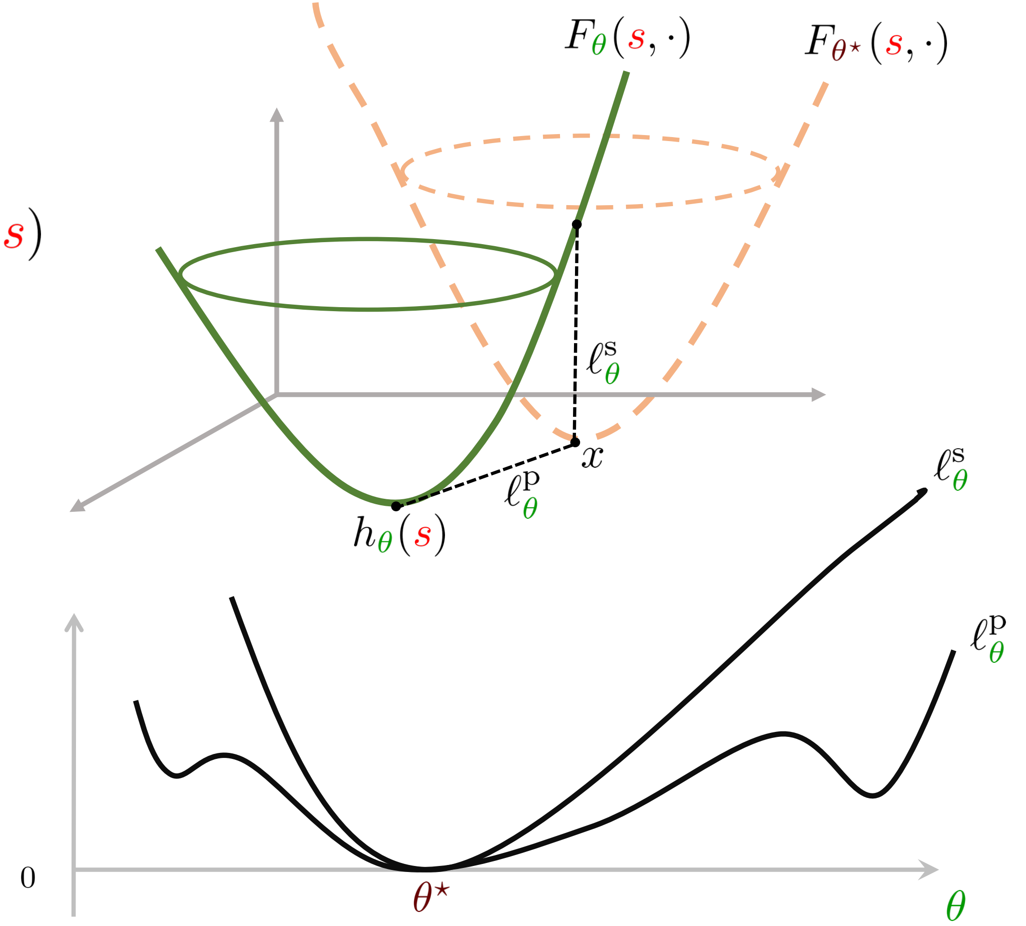

Inverse Optimization

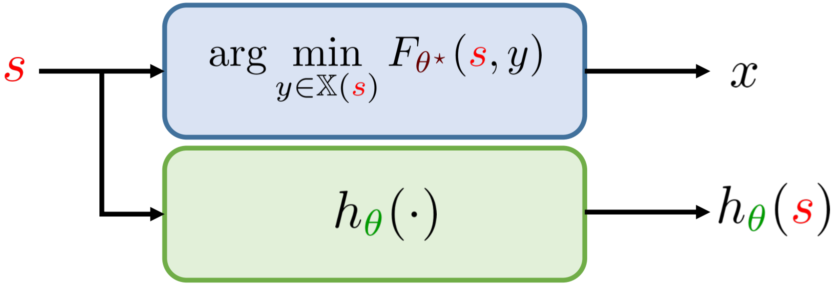

$$h_\theta(s) = \arg\min_{y\in\mathbb{X}(s)} F_\theta(s,y)$$

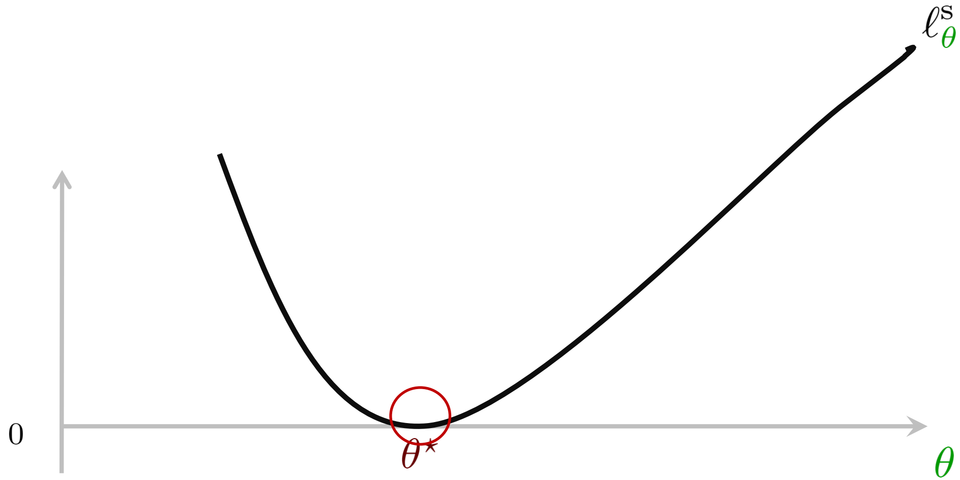

Inverse Optimization

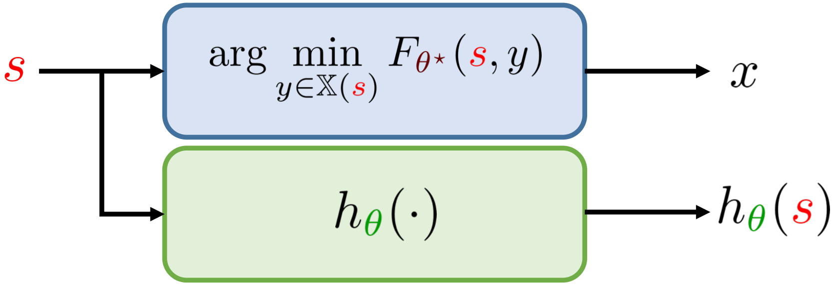

$$h_\theta(s) = \arg\min_{y\in\mathbb{X}(s)} F_\theta(s,y)$$

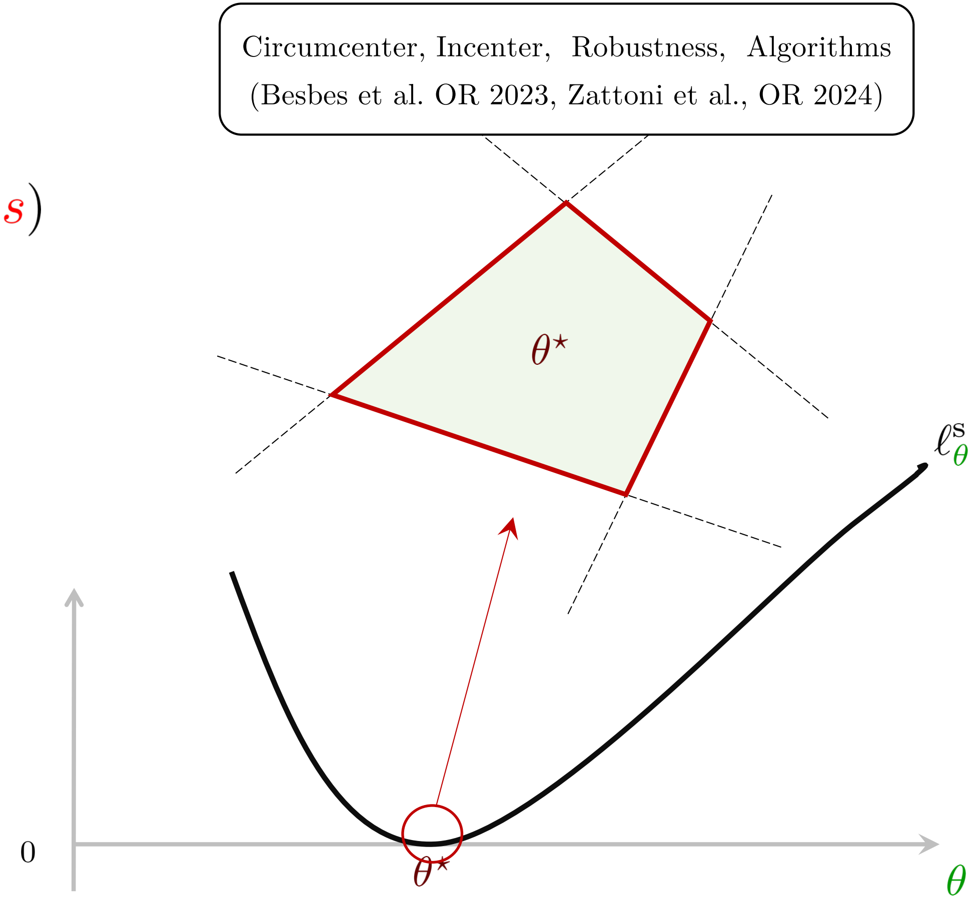

Inverse Optimization

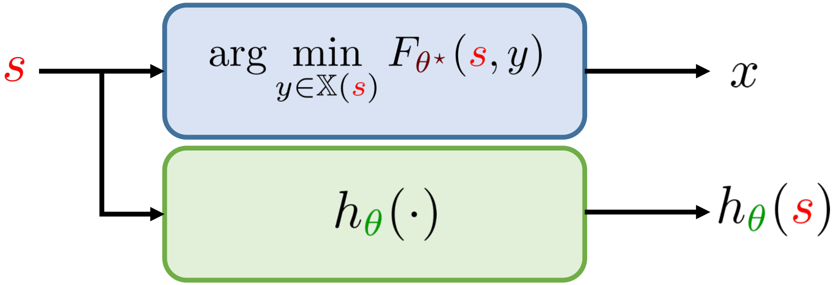

$$h_\phi(s) = \arg\min_{y\in\mathbb{X}(s)} \langle \underbrace{\phi(s)}_{}, y\rangle$$

Non-parametric learning

(Long et al., NeurIPS 2024)

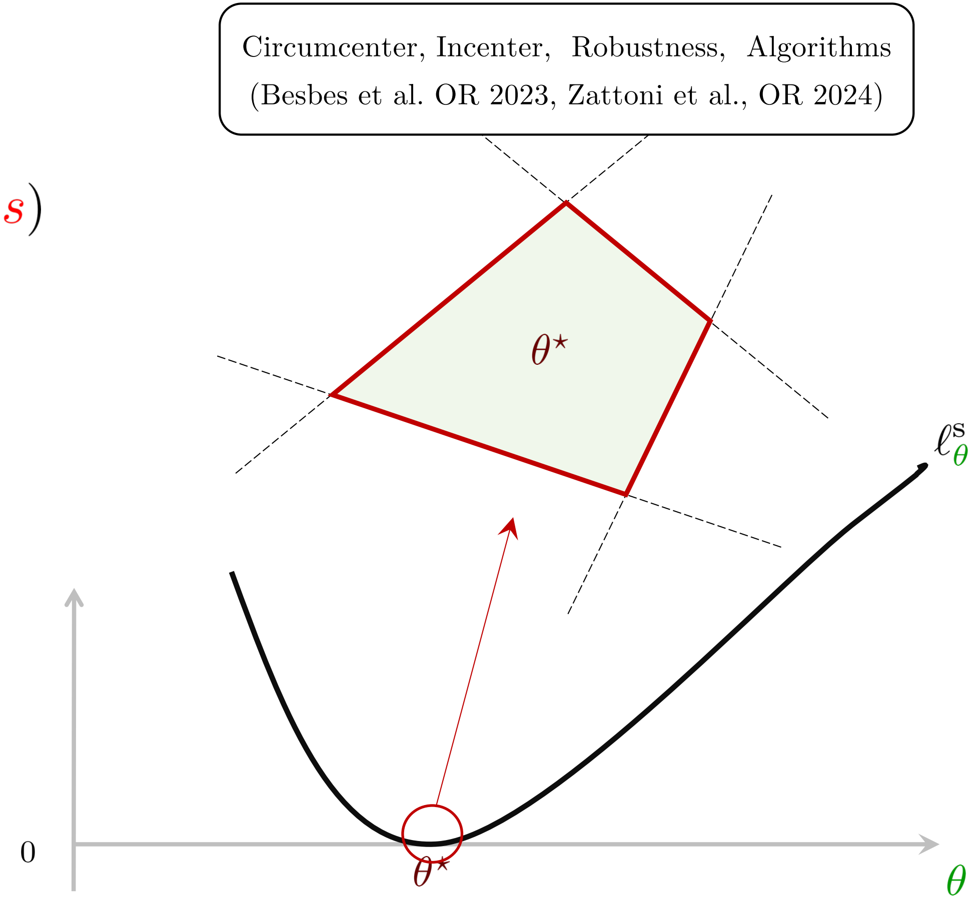

Inverse Optimization

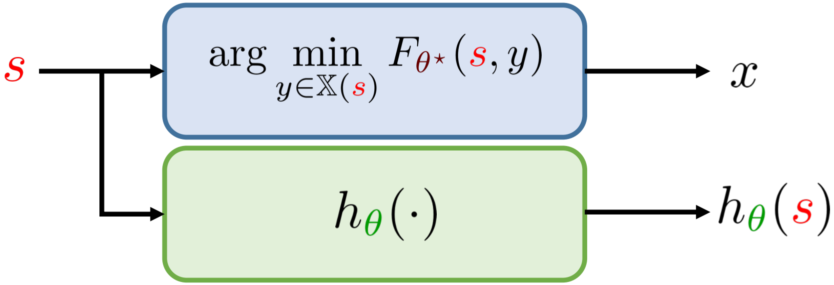

$$h_\phi(s) = \arg\min_{y\in\mathbb{X}(s)} \langle \underbrace{\phi(s)}, y\rangle$$

Non-parametric learning

(Long et al., NeurIPS 2024)

$$h_\theta(s) = \arg\min_{\underbrace{y\in\mathbb{X}_\theta(s)}} \langle \phi(s), y\rangle$$

Constraints learning

(Ke et al., ICML 2025)

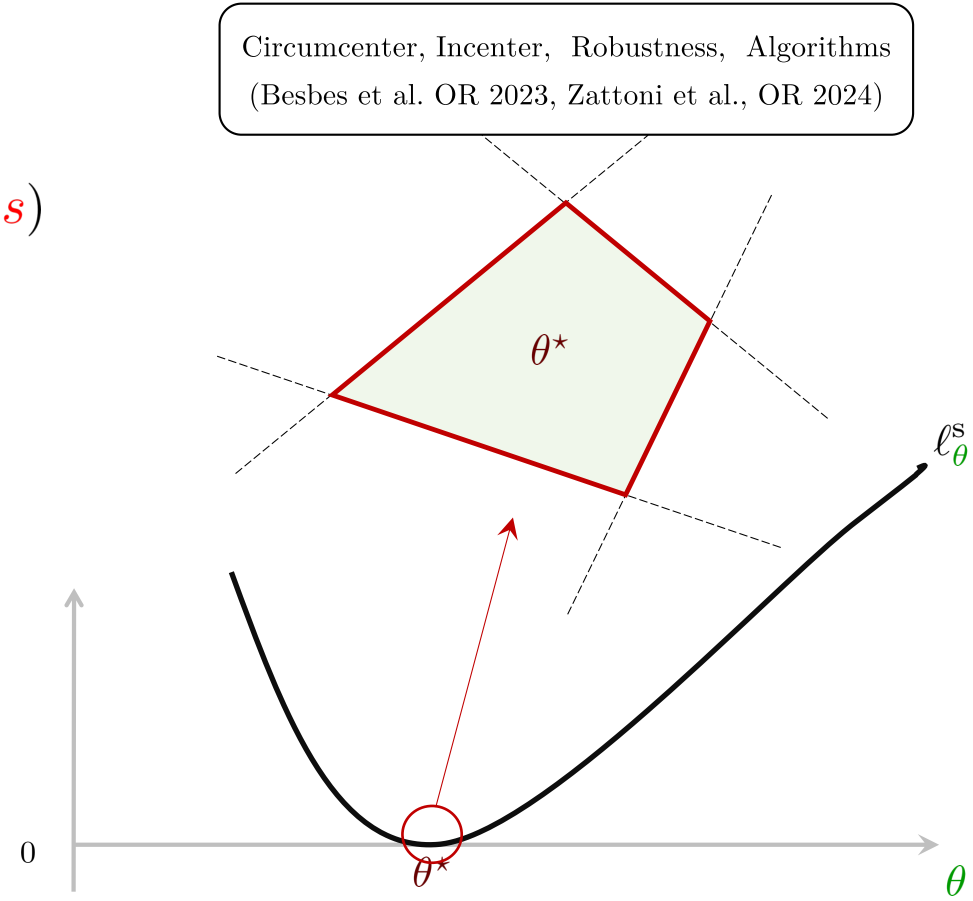

Outline

- Data-driven decision-making

- Inverse optimization

- Applications

  - A competition for Neural Networks !?

Model Predictive Control (MPC)

Lab helicopter

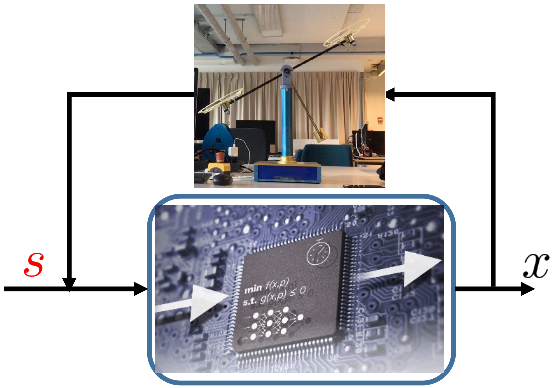

Model Predictive Control (MPC)

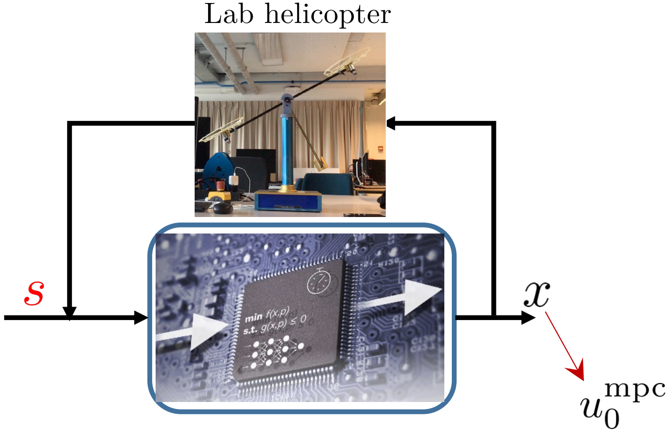

Model Predictive Control (MPC)

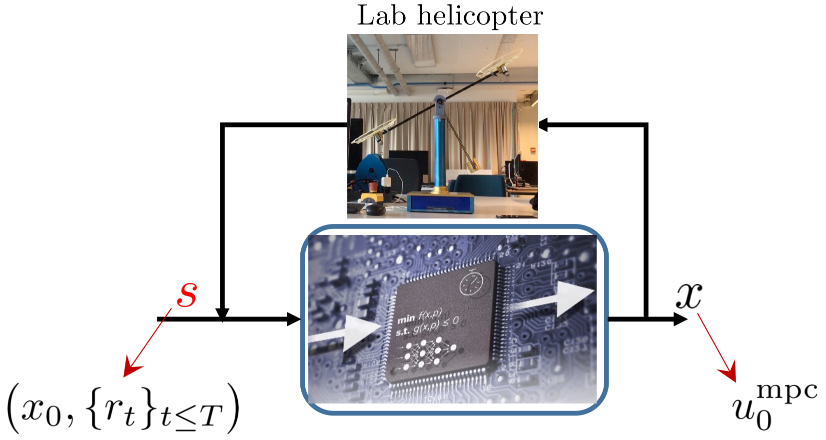

Model Predictive Control (MPC)

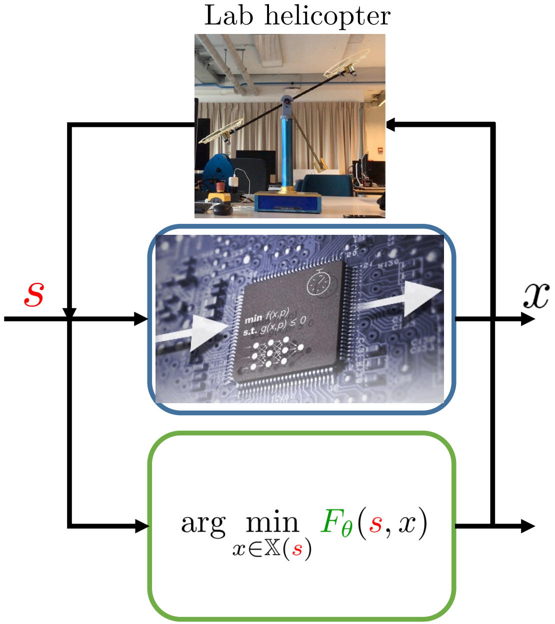

Model Predictive Control (MPC)

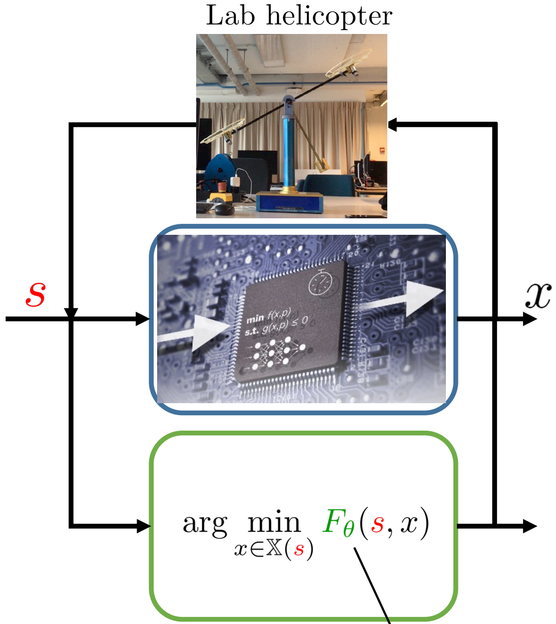

$$F_\theta(s,x) = x'\Theta_{xx}x + s'\Theta_{sx}x$$

<u>36</u> parameters

Model Predictive Control (MPC)

Lab helicopter

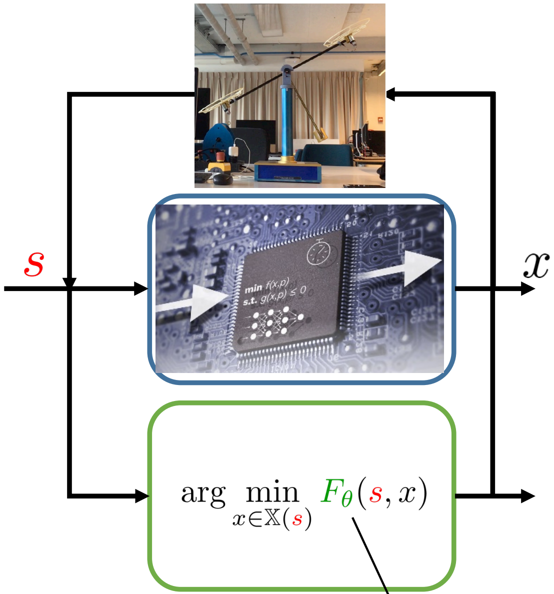

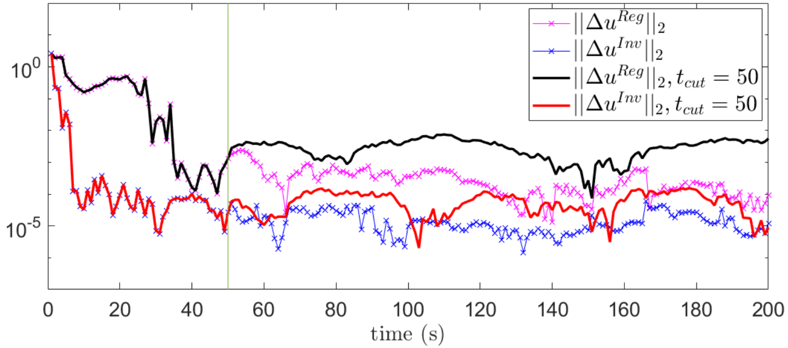

$$F_\theta(s,x) = x'\Theta_{xx}x + s'\Theta_{sx}x$$

<u>36</u> parameters

Model Predictive Control (MPC)

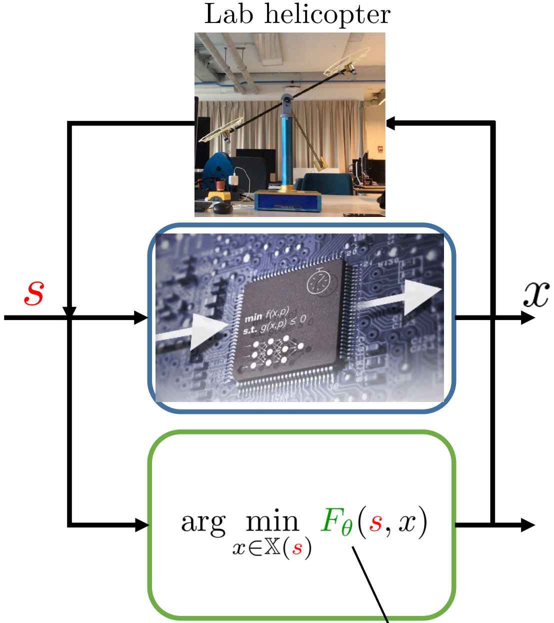

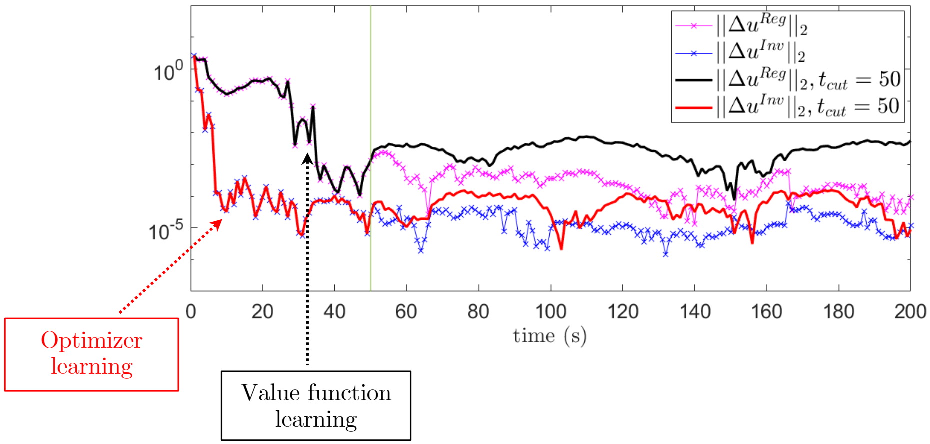

$$F_\theta(s, x) = x' \Theta_{xx} x + s' \Theta_{sx} x$$

36 parameters

MuJoCo Environments

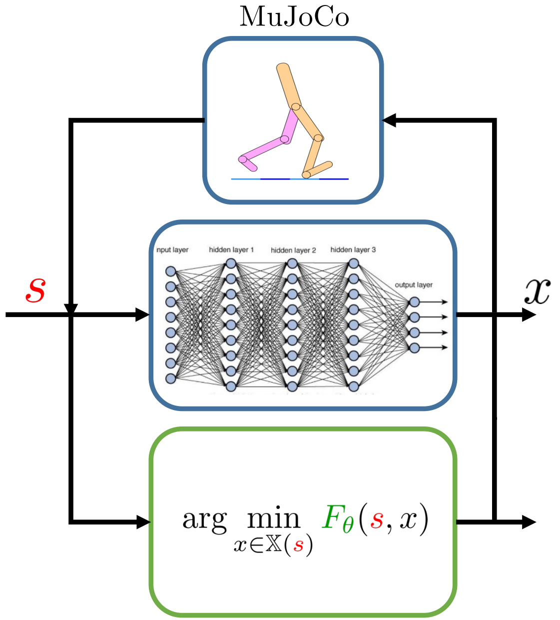

MuJoCo Environments

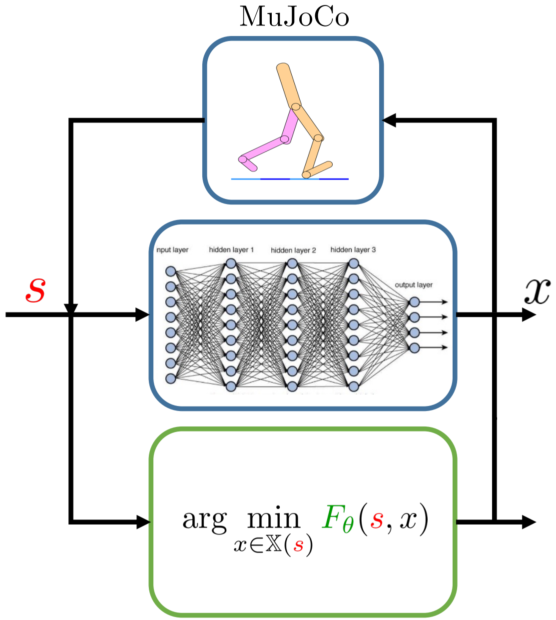

	# Parameters	Training dataset	Scores
Neural Network	2,489,949	1M	82.9
Inverse Optimization	840	5k	70.6

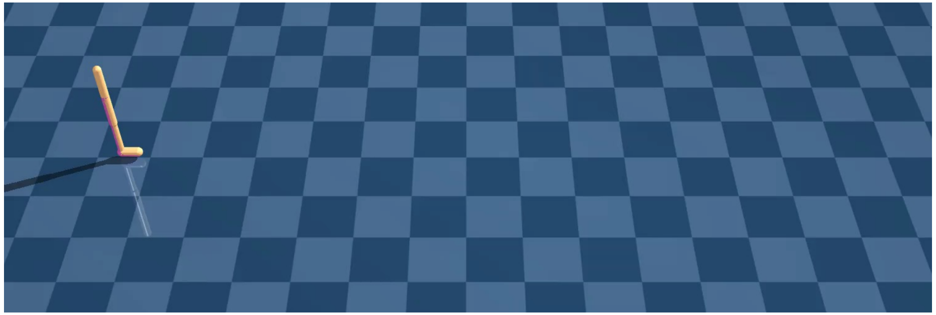

Last Mile Routing Challenge

demands
set of customers
time-windows
...

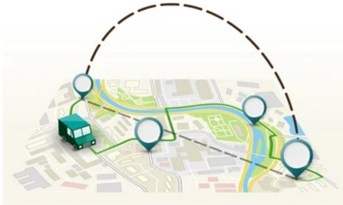

Last Mile Routing Challenge

demands
set of customers
time-windows
...

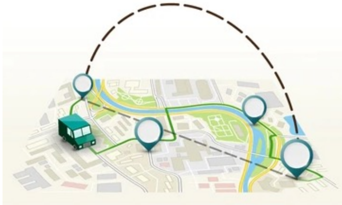

- "Tacit knowledge often contradicts optimized route plans"

Last Mile Routing Challenge

amazon

demands
set of customers
time-windows
...

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **6112** real-world routes from expert human drivers

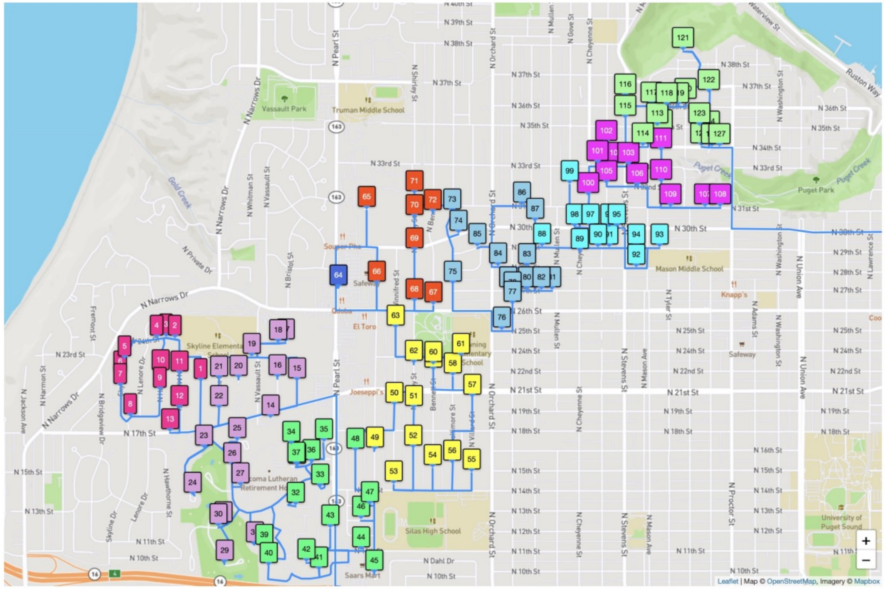

Last Mile Routing Challenge

amazon

demands
set of customers
time-windows
...

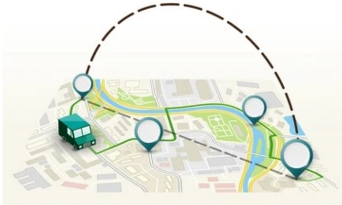

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **6112** real-world routes from expert human drivers
- **Goal**: learn how to route like expert human drivers.

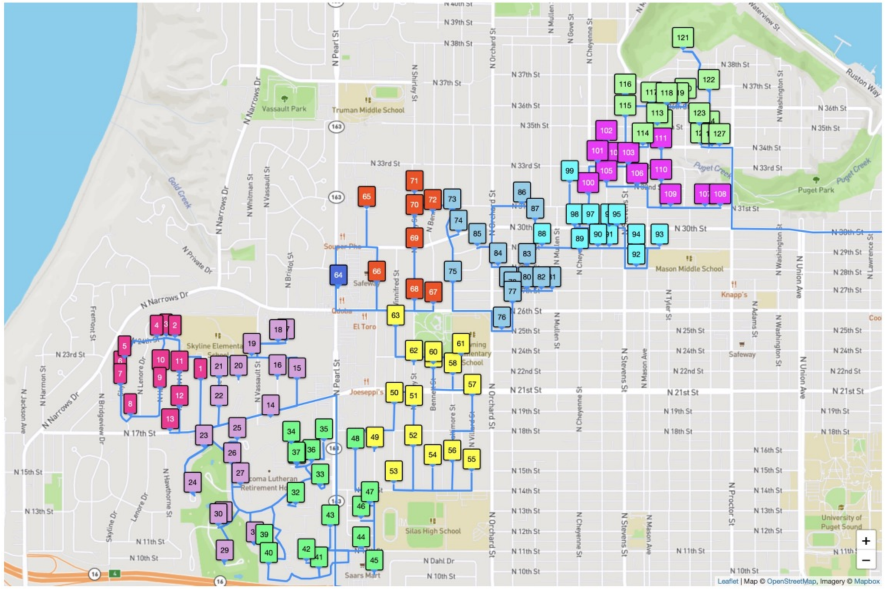

Last Mile Routing Challenge

demands
set of customers
time-windows
...

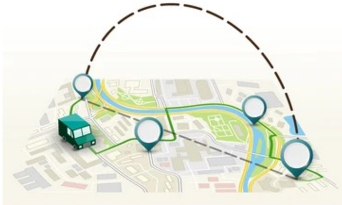

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **6112** real-world routes from expert human drivers
- **Goal**: learn how to route like expert human drivers.
- **$175,000** in prizes!

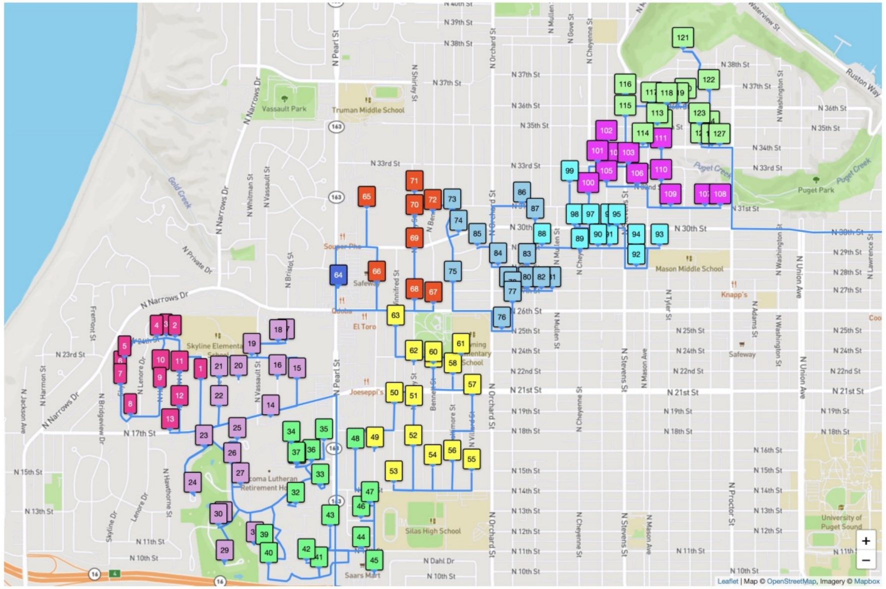

Last Mile Routing Challenge

amazon

demands
set of customers
time-windows
...

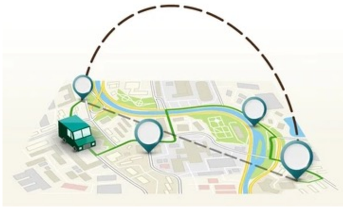

$$\arg\min_{x \in \mathbb{X}(s)} F_\theta(s, x)$$

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **<u>6112</u>** real-world routes from expert human drivers
- **<u>Goal</u>**: learn how to route like expert human drivers.
- **<u>$175,000</u>** in prizes!

Last Mile Routing Challenge

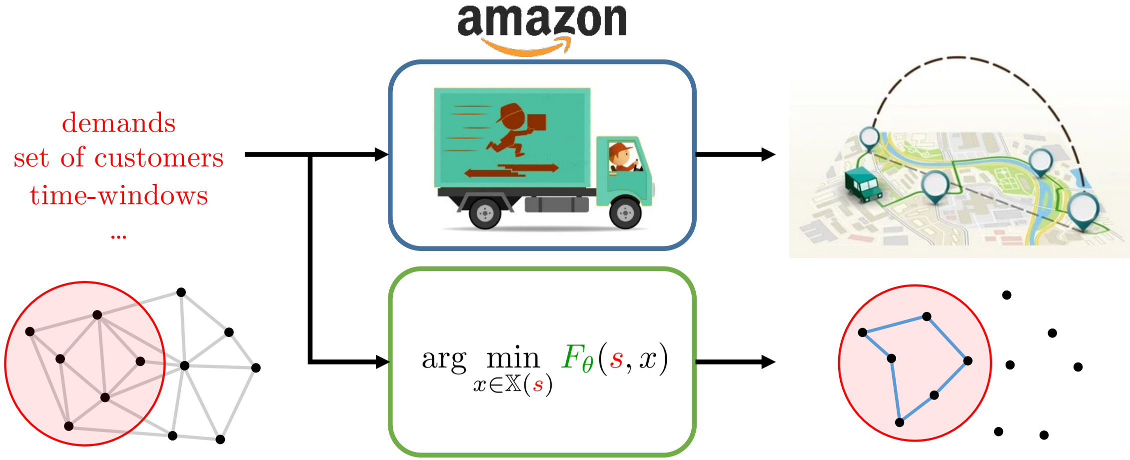

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **6112** real-world routes from expert human drivers
- **Goal**: learn how to route like expert human drivers.
- **$175,000** in prizes!

Last Mile Routing Challenge

demands
set of customers
time-windows
...

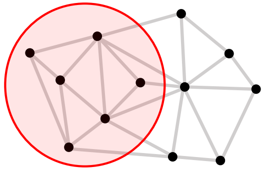

$$\arg \min_{x \in \mathbb{X}(s)} F_\theta(s, x)$$

traveling salesperson problem

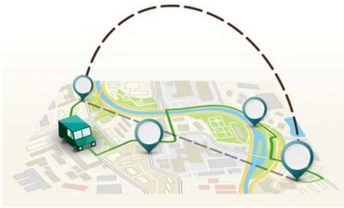

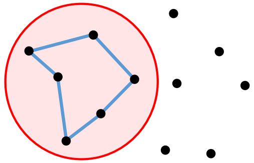

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **6112** real-world routes from expert human drivers
- **Goal**: learn how to route like expert human drivers.
- **$175,000** in prizes!

Last Mile Routing Challenge

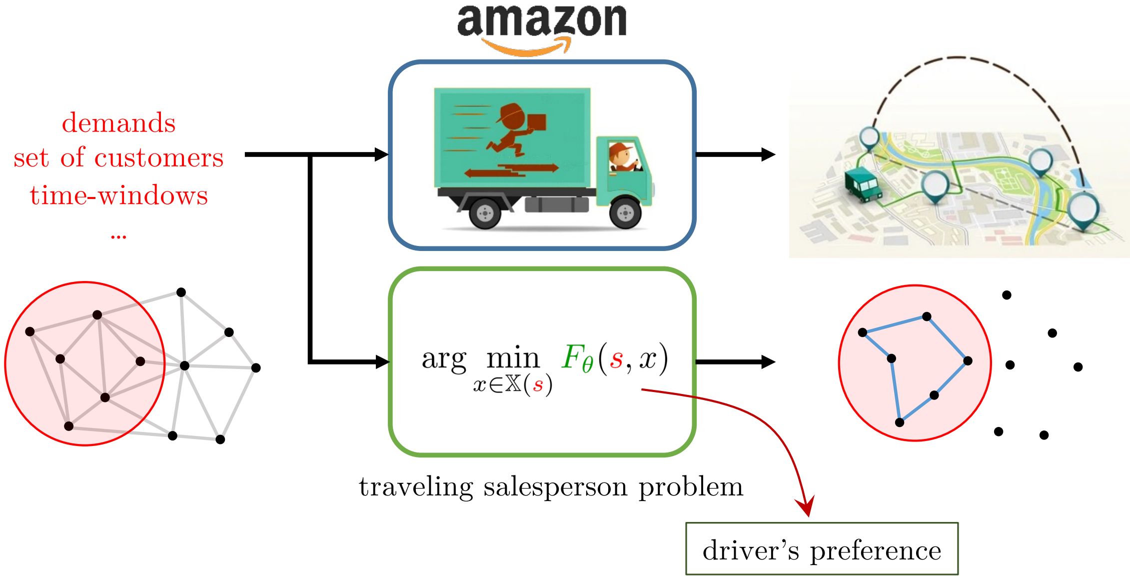

- "Tacit knowledge often contradicts optimized route plans"
- Dataset with **6112** real-world routes from expert human drivers
- **Goal**: learn how to route like expert human drivers.
- **$175,000** in prizes!

Last Mile Routing Challenge

amazon

demands
set of customers
time-windows
...

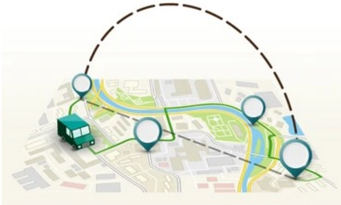

$$\arg \min_{x \in \mathbb{X}(s)} F_\theta(s, x)$$

traveling salesperson problem

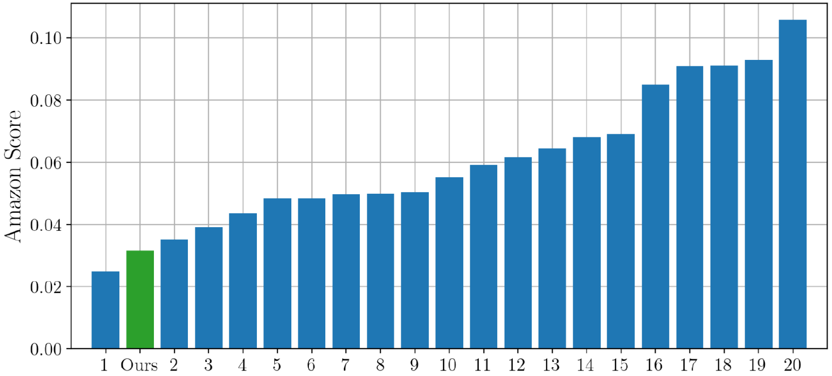

- “Tacit knowledge often contradicts optimized route plans”
- Dataset with **6112** real-world routes from expert human drivers
- **Goal**: learn how to route like expert human drivers.
- **$175,000** in prizes!

References

(available at https://mohajerinesfahani.github.io/)

Sub-optimality Loss & DRO	• **PME**, Shafieezadeh-Abadeh, Hanasusanto, and Kuhn, "Data-driven Inverse Optimization with Imperfect Information", *Mathematical Programming (MP)*, 2018
Incenter, robustness, algorithms	• Zattoni Scroccaro, Atasoy, **PME**, "Learning in Inverse Optimization: Incenter Cost, Augmented Suboptimality Loss, and Algorithms", *Operations Research (OR)*, 2024
Routing problems	• Zattoni Scroccaro, van Beek, **PME**, Atasoy, "Inverse Optimization for Routing Problems", *Transportation Science (TS)*, 2024
Nonparametric approach	• Long, Ok, Zattoni Scroccaro, **PME**, "Scalable Kernel Inverse Optimization", *Conference on Neural Information Processing Systems (NeurIPS)*, 2024
Constraints learning	• Ke, PME, Georghiou, "Inverse Optimization via Learning Feasible Regions", *International Conference on Machine Learning (ICML)*, 2025
Code	• Zattoni Scroccaro, "InvOpt: Inverse Optimization with Python", https://github.com/pedroszattoni/invopt, 2024