Scalable Kernel Inverse Optimization

Youyuan Long

Tolga Ok

Pedro Zattoni Scroccaro

Peyman Mohajerin Esfahani

Delft Center for Systems and Control (DCSC), Delft University of Technology

TU Delft

NEURAL INFORMATION PROCESSING SYSTEMS

Inverse Optimization

Assumption:

The expert optimizes an unknown cost to make a decision.

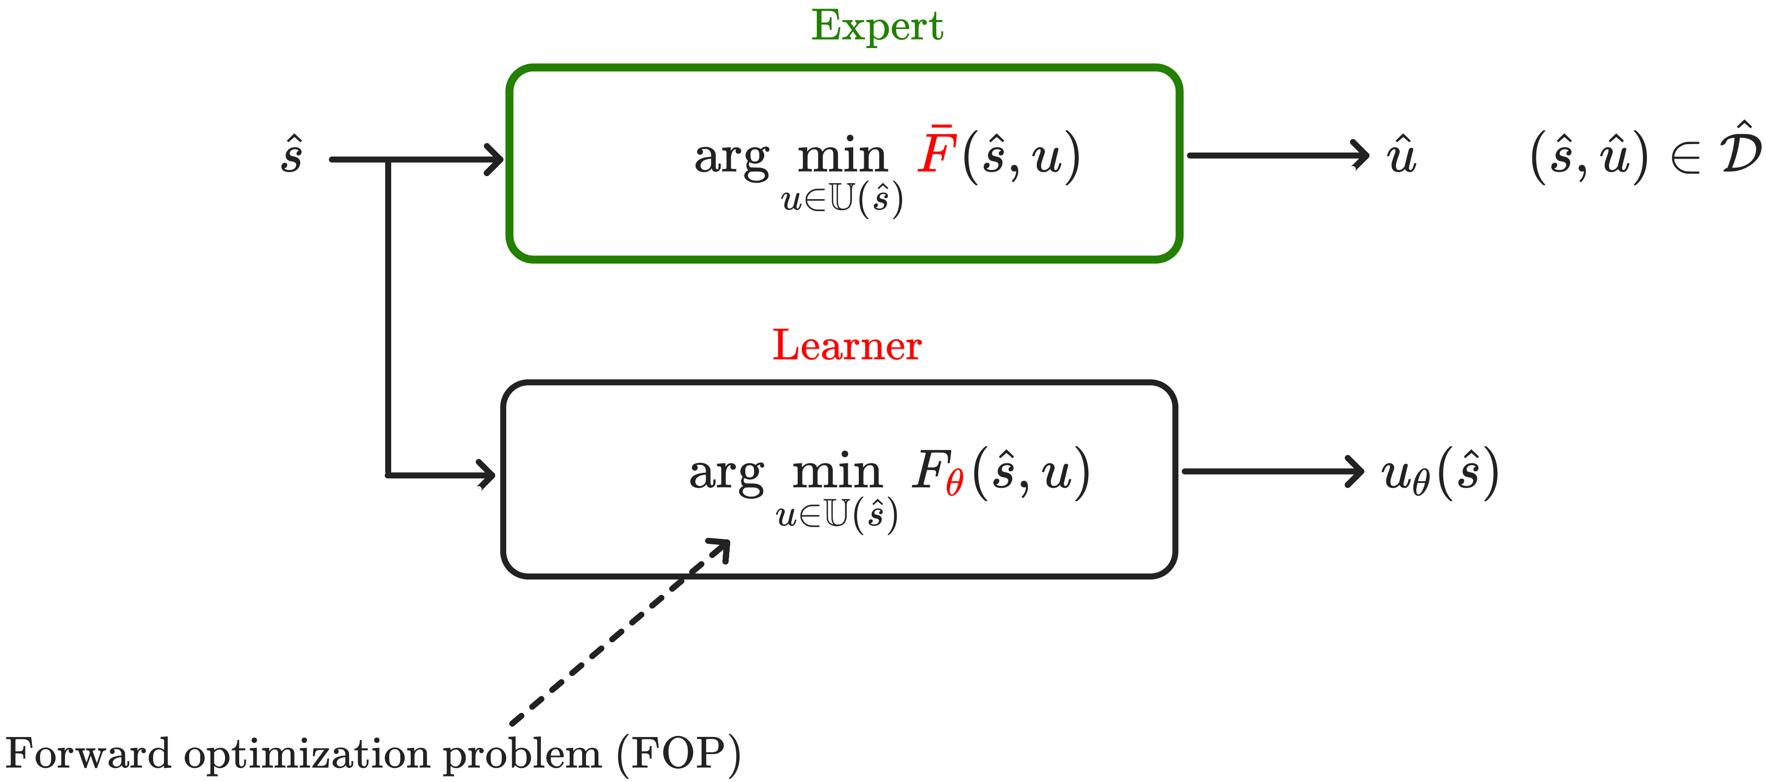

Inverse Optimization

Supervised Learning Perspective

1) Hyphothesis class of FOP: $\{F_\theta | \theta \in \Theta\}$

2) Loss function: $\ell(\hat{u}, u_\theta(\hat{s}))$

3) Optimization Algorithm: $\min_{\theta \in \Theta} \mathcal{L}(\theta)$

Inverse Optimization

Supervised Learning Perspective

1) Hyphothesis class of FOP: Quadratic F_θ

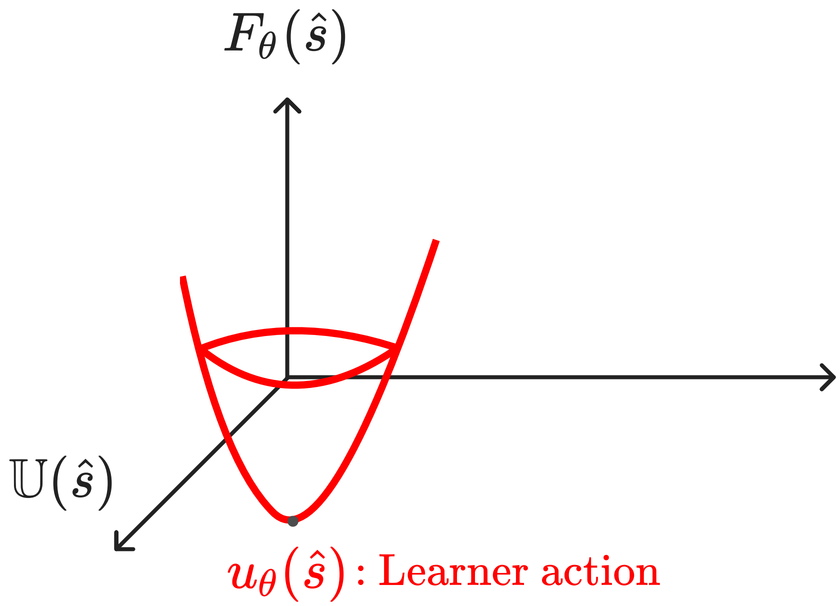

$$F_\theta(\hat{s}) = u^\top u + \phi(s)^\top \theta_{su} u$$

$$u_\theta(\hat{s}) = \arg \min_{u \in \mathbb{U}(\hat{s})} F_\theta(\hat{s}, u)$$

2) Loss function: $\ell(\hat{u}, u_\theta(\hat{s}))$

3) Optimization: $\min_{\theta \in \Theta} k\mathcal{R}(\theta) + \frac{1}{N} \sum_{i=1}^{N} \ell(\hat{u}, u_\theta(\hat{s}))$

Inverse Optimization

Supervised Learning Perspective

2) Loss function: Suboptimality loss[1]

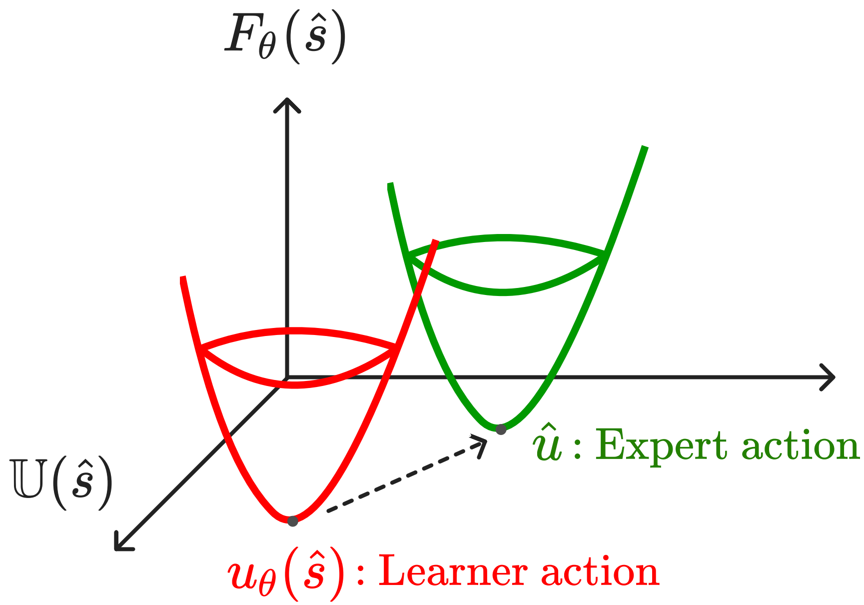

$$u_\theta(\hat{s}) = \arg\min_{u \in \mathbb{U}(\hat{s})} F_\theta(\hat{s}, u)$$

$$\ell(\hat{u}_i, \hat{s}_i, \theta) = \min_{u_i \in \mathbb{U}(\hat{s}_i)} \{F_\theta(\hat{s}_i, \hat{u}_i) - F_\theta(\hat{s}_i, u_i)\}$$ Convex !

[1] Mohajerin Esfahani P, Shafieezadeh-Abadeh S, Hanasusanto G A, et al. Data-driven inverse optimization with imperfect information[J]. Mathematical Programming, 2018, 167: 191-234.

Inverse Optimization

Supervised Learning Perspective

3) Optimization Algorithm: Semidefinite Program (SDP) solver

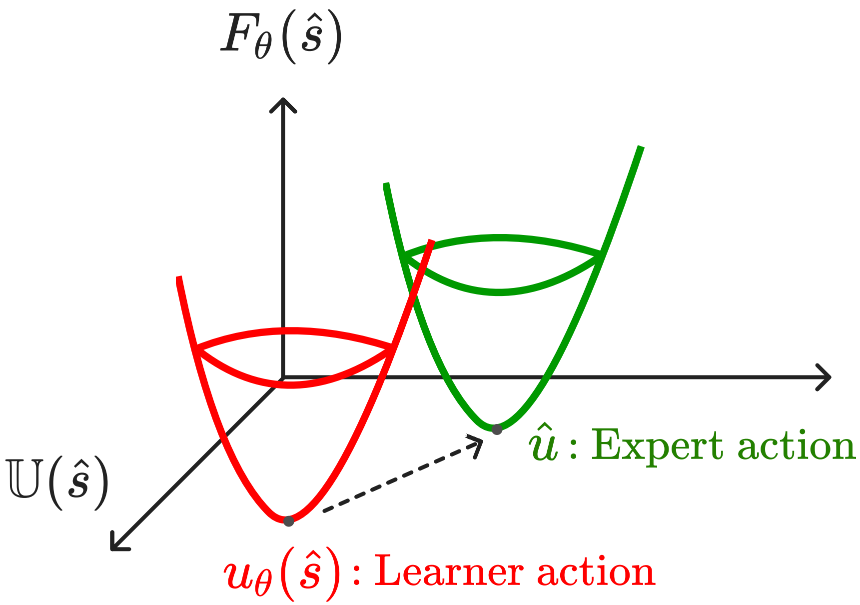

$$\min_{\theta \in \Theta} k\|\theta\|^2 + \frac{1}{N}\sum_{i=1}^{N} \ell(\hat{u}, \hat{s})$$

Primal Problem

LMI reformulation[1]

$$\begin{aligned} \text{minimize} \quad & x^\top c \\ \text{s.t.} \quad & G(x) \succcurlyeq 0 \\ & Ax = 0 \end{aligned}$$

[1] Akhtar S A, Kolarijani A S, Esfahani P M. Learning for control: An inverse optimization approach. IEEE Control Systems Letters, 2021, 6: 187-192.

Kernel Inverse Optimization

Can we generalize the function class of F? Yes!

Kernel Inverse Optimization

Can we generalize the function class of F? Yes!

$$\mathcal{L} : \theta \to \mathbb{R}$$

$$\min_{\theta \in \Theta} k\|\theta\|^2 + \frac{1}{N}\sum_{i=1}^{N} \ell(\hat{u}, \hat{s}, \theta)$$

Θ: Parameter space

$$\mathcal{L} : \mathcal{H} \to \mathbb{R}$$

$$\min_{f \in \mathcal{H}} k\|f\|_{\mathcal{H}}^2 + \frac{1}{N}\sum_{i=1}^{N} \ell(\hat{u}, \hat{s}, f)$$

$\mathcal{H}$: Reproducing Kernel Hilbert Space

Kernel Inverse Optimization

The resulting Problem is Convex!

$$\mathcal{L} : \theta \to \mathbb{R}$$

$$\min_{\theta \in \Theta} k\|\theta\|^2 + \frac{1}{N}\sum_{i=1}^{N} \ell(\hat{u}, \hat{s}, \theta)$$

Θ: Parameter space

$$\mathcal{L} : \mathcal{H} \to \mathbb{R}$$

$$\min_{f \in \mathcal{H}} k\|f\|_{\mathcal{H}}^2 + \frac{1}{N}\sum_{i=1}^{N} \ell(\hat{u}, \hat{s}, f)$$

$\mathcal{H}$: Reproducing Kernel Hilbert Space

$$\arg\min_{u \in \mathbb{U}(\hat{s})} u^\top u + \phi(s)^\top \theta_{su} u$$

Quadratic FOP

$$\arg\min_{u \in \mathbb{U}(\hat{s})} u^\top u + \sum_{i=1}^{N} \kappa(\hat{s}, \hat{s}_i)\alpha_i^\top u$$

Kernelized FOP

Kernel Inverse Optimization

Performance Evaluation

MuJoCo[1] Benchmark on D4RL[2] dataset!

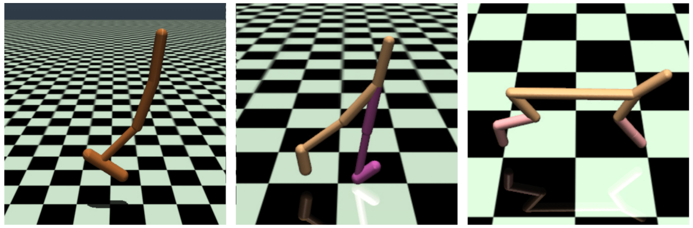

Hopper Walker2d HalfCheetah

Kernelized FOP → KIO (5k)

Quadratic FOP → IO (5k)

Medium quality data collecting agent → Teacher

Task	KIO (5k)	IO (5k)	TD3	CQL	Teacher
Hopper	**50.2**	20.6	30.0	29.0	47.2
Walker	**74.6**	0.0	11.4	6.6	68.1
HalfChetah	**39.0**	-3.1	**36.6**	**36.1**	40.7

[1] E. Todorov, T. Erez and Y. Tassa, "MuJoCo: A physics engine for model-based control," 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura-Algarve, Portugal, 2012

[2] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, & Sergey Levine (2020). D4RL: Datasets for Deep Data-Driven Reinforcement Learning. CoRR, abs/2004.07219.

Scalable Kernel Inverse Optimization

Sequential Selection Optimization (SSO)

The memory complexity is $\mathcal{O}(N^2)$.

$$\mathrm{FOP}(\hat{s}; \alpha_i) = \arg\min_{u \in \mathbb{U}(\hat{s})} u^\top u + \sum_{i=1}^{N} \kappa(\hat{s}, \hat{s}_i) \alpha_i{}^\top u$$

SSO: Coordinate descent[1] based algorithm

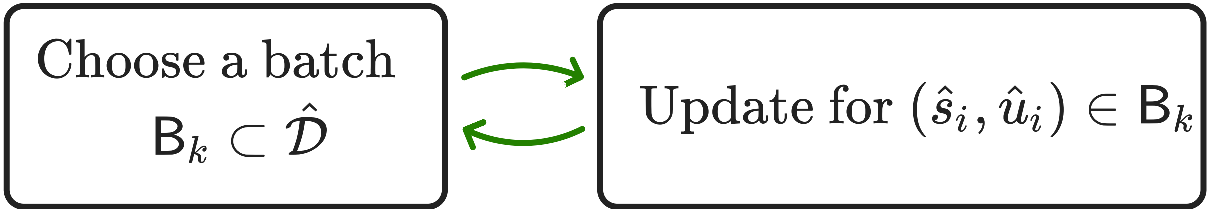

`kkt_violator(i)`

[1] Nesterov, Y. (2012). Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems. SIAM Journal on Optimization, 22(2), 341-362.

Scalable Kernel Inverse Optimization

Sequential Selection Optimization (SSO)

Converges to the same global optimal!

$$\mathrm{FOP}(\hat{s};\alpha_i) = \arg\min_{u\in\mathbb{U}(\hat{s})} u^\top u + \sum_{i=1}^{N} \kappa(\hat{s},\hat{s}_i){\alpha_i}^\top u$$

Task	KIO	KIO-SSO
Hopper	50.2	51.8
Walker	74.6	74.9
HalfChetah	39.0	39.7

SSO: Coordinate descent[1] based algorithm

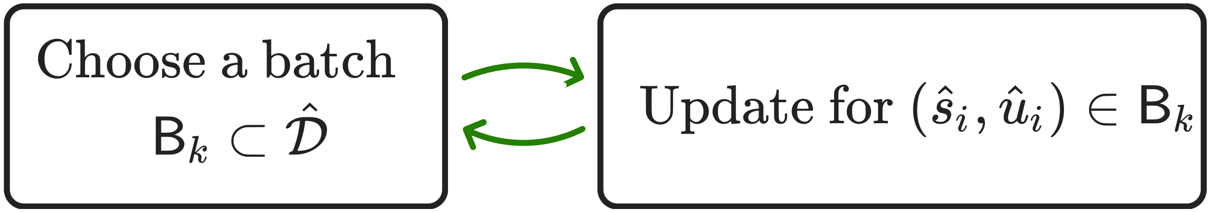

`kkt_violator(i)`

[1] Nesterov, Y. (2012). Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems. SIAM Journal on Optimization, 22(2), 341-362.

Scalable Kernel Inverse Optimization

Conclusion

Convex Problem

Generalization capabilities with relatively low-data

Possible extensions via sophisticated kernels

The cost of evaluating the FOP is $\mathcal{O}(N)$

Lack of convergence guarantees for SSO with **`kkt_violator`**