First-Order Methods for Wasserstein Distributionally Robust MDPs

Julien Grand-Clément, Christian Kroer

IEOR Department, Columbia University

What is this paper about?

- Solving distributionally robust Markov Decision Processes (MDP)
- Distributional robustness with Wasserstein distance
- Goal: improve convergence rate of Value Iteration algorithms

- Main idea: adapt First-Order Methods[1] (FOMs) for DR-MDP

- Challenges:
  - Adapt FOMs to a dynamically changing setting
  - Develop tractable proximal updates for interesting metrics/decision sets

[1] Mirror Descent, Mirror Prox, Primal-Dual Algorithms, etc.

Main complexity results

Consider an MDP with $S/A/N$ states/actions/observed kernels.

Complexity of our algorithm:

$$O\left(NA^{2.5}S^{3.5}\log(\epsilon^{-1})\epsilon^{-1.5}\right). \tag{1}$$

Note: complexity of classical Value Iteration:

$$O\left(N^{3.5}A^{3.5}S^{4.5}\log^2(\epsilon^{-1})\right). \tag{2}$$

$\Rightarrow$ improvement of $\Omega\left(N^{2.5}AS\right)$ ...

... at the price of a worst-dependence in ϵ^{-1}.

Numerical study

Comparing our algorithms with other methods:

- Three MDP instances (two real, one random).
- Other algorithms[2]: Value Iteration (VI), Gauss-Seidel VI, Accelerated VI, Anderson VI.
- We compare running times to return ϵ-optimal policy, when S, N grow larger.

[2] Puterman (1994), Geist and Scherrer (2018), Goyal and G.-C. (2019)

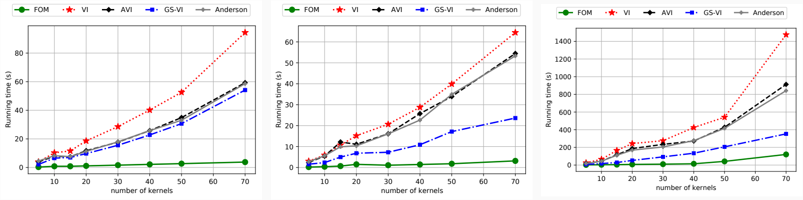

(a) Forest. **(b)** Machine. **(c)** Garnet.

Figure 1: Running times of various algorithms (vs. number of kernels N).

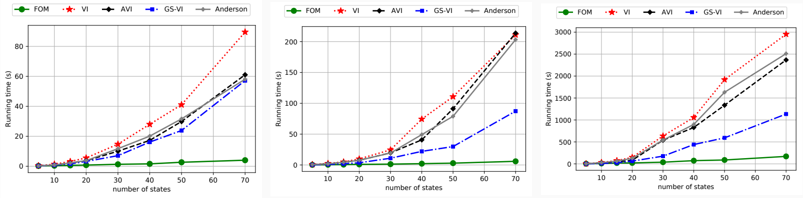

(a) Forest. **(b)** Machine. **(c)** Garnet.

Figure 2: Running times of various algorithms (vs. number of states S).

Conclusion

Our contributions.

- We present the first algorithm adapting FOM to *Distributionally Robust MDP.*
- In terms of state/action, we improve complexity to $O(NA^{2.5}S^{3.5})$ from previous $O(N^{3.5}A^{3.5}S^{4.5})$.
- Empirically, significant speedups on both random and structured MDP instances, even for small N, S, A.

In the paper:

- More details on Wasserstein setup,
- Details on convergence rate/complexity,
- Detailed simulation setup.