Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location

Priyank Agrawal Columbia University

Eric Balkanski Columbia University

Vasilis Gkatzelis Drexel University

Tingting ou Columbia University

Xizhi Tan Drexel University

Learning-augmented mechanisms

- A serge of recent work on algorithm design with prediction has been proposed as an alternative to worst-case analysis.
  - **Consistency:** The worst-case performance over instances where the **prediction is accurate**
  - **Robustness:** The worst-case performance overall (even when the **prediction is arbitrarily bad**)
- We extend the framework to mechanism design problems where mechanisms is augmented with **predictions** regarding the private information
- **Objective**: Design learning-augmented mechanisms that achieve (near) optimal performance when the prediction is correct maintaining some non-trivial worst-case guarantees

Strategic facility location

- Each agent i has a **preferred location** $p_i \in \mathbb{R}^2$
- We need to choose a **single facility location** $f \in \mathbb{R}^2$
- The cost of agent i is the Euclidean distance from p_i to f
- Goal: minimizing the social costs

1. **Egalitarian** social cost: the **maximum cost** over all agents
2. **Utilitarian** social cost: the **average cost** over all agents

Mechanism design problem: Can we approximate these objectives when the preferred locations are **private**?

1. For egalitarian social cost, no strategyproof mechanism can achieve an approximation better than **2**. [Procaccia and Tennenholtz '09]
2. For utilitarian social cost, no strategyproof, anonymous mechanism can achieve an approximation better than $\sqrt{2}$. [Feigenbaum et al '17]

Our Results

Main Question: Can we design learning-augmented mechanisms to overcome pessimistic worst-case analysis results? Can learning-augmented mechanisms achieve good robustness and consistency trade-offs?

Problem	Lower bound w/o predictions	Our results		Optimality
		Consistency	**Robustness**	
Egalitarian in $\mathbb{R}$	2	1	2	Best of both worlds
Egalitarian in $\mathbb{R}^2$	2	1	$1+\sqrt{2}$	Optimal trade-off
Utilitarian in $\mathbb{R}$	1	-	-	-
Utilitarian in $\mathbb{R}^2$	$\sqrt{2}$	$\frac{\sqrt{2c^2+2}}{1+c}$	$\frac{\sqrt{2c^2+2}}{1-c}$	Optimal trade-off

Open Questions

- Learning augmented mechanism design for facility location in higher dimensions
- Other centralized mechanism design settings with predictions
- Decentralized mechanism design with predictions

Mechanisms with predictions and intuition

The mechanism has access to a **prediction regarding the optimal facility location $\hat{f}$**

For egalitarian social cost in $\mathbb{R}$:

- If the prediction is between the minimum and the maximum: **return prediction $\hat{f}$**

- Otherwise, return the point that is closest to the prediction

The mechanism is strategyproof, **1-consistent**, and **2-robust**; it achieves the best consistency and the best robustness at the same time!

For egalitarian social cost in $\mathbb{R}^2$, the **Minimum bounding box** mechanism returns the prediction if the prediction is inside the bounding box, and the point in the box that is closest to the prediction otherwise. The mechanism is **1-consistent** and $\mathbf{1+\sqrt{2}}$**-robust**

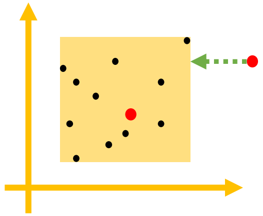

Any strategyproof mechanism that is better than **2**-consistent is no better than **1**+$\sqrt{\mathbf{2}}$-robust

For utilitarian social cost in $\mathbb{R}^2$, the **Coordinatewise Median with Predictions** mechanism with confidence $c \in (0,1]$ adds cn phantom points at the prediction $\hat{f}$ and outputs the point that is median on both axis. It achieves the **optimal trade-off** between consistency and robustness.

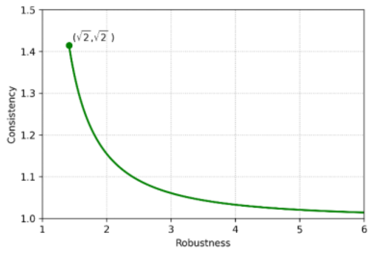

Acknowledgments: This project was supported in part by NSF grant CCF-2047907.