MAC Advice for Facility Location Mechanism Design

Zohar Barak (TAU),

Anupam Gupta (NYU, Google Research),

Inbal Talgam-Cohen (TAU)

NeurIPS 2024

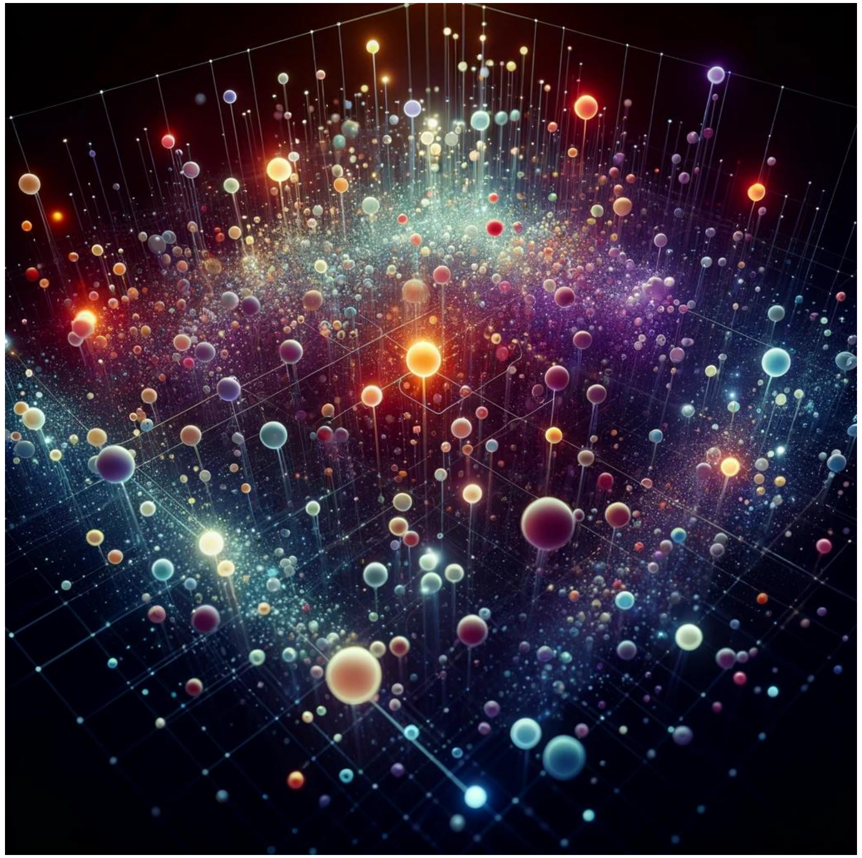

Algorithms with Predictions/Advice

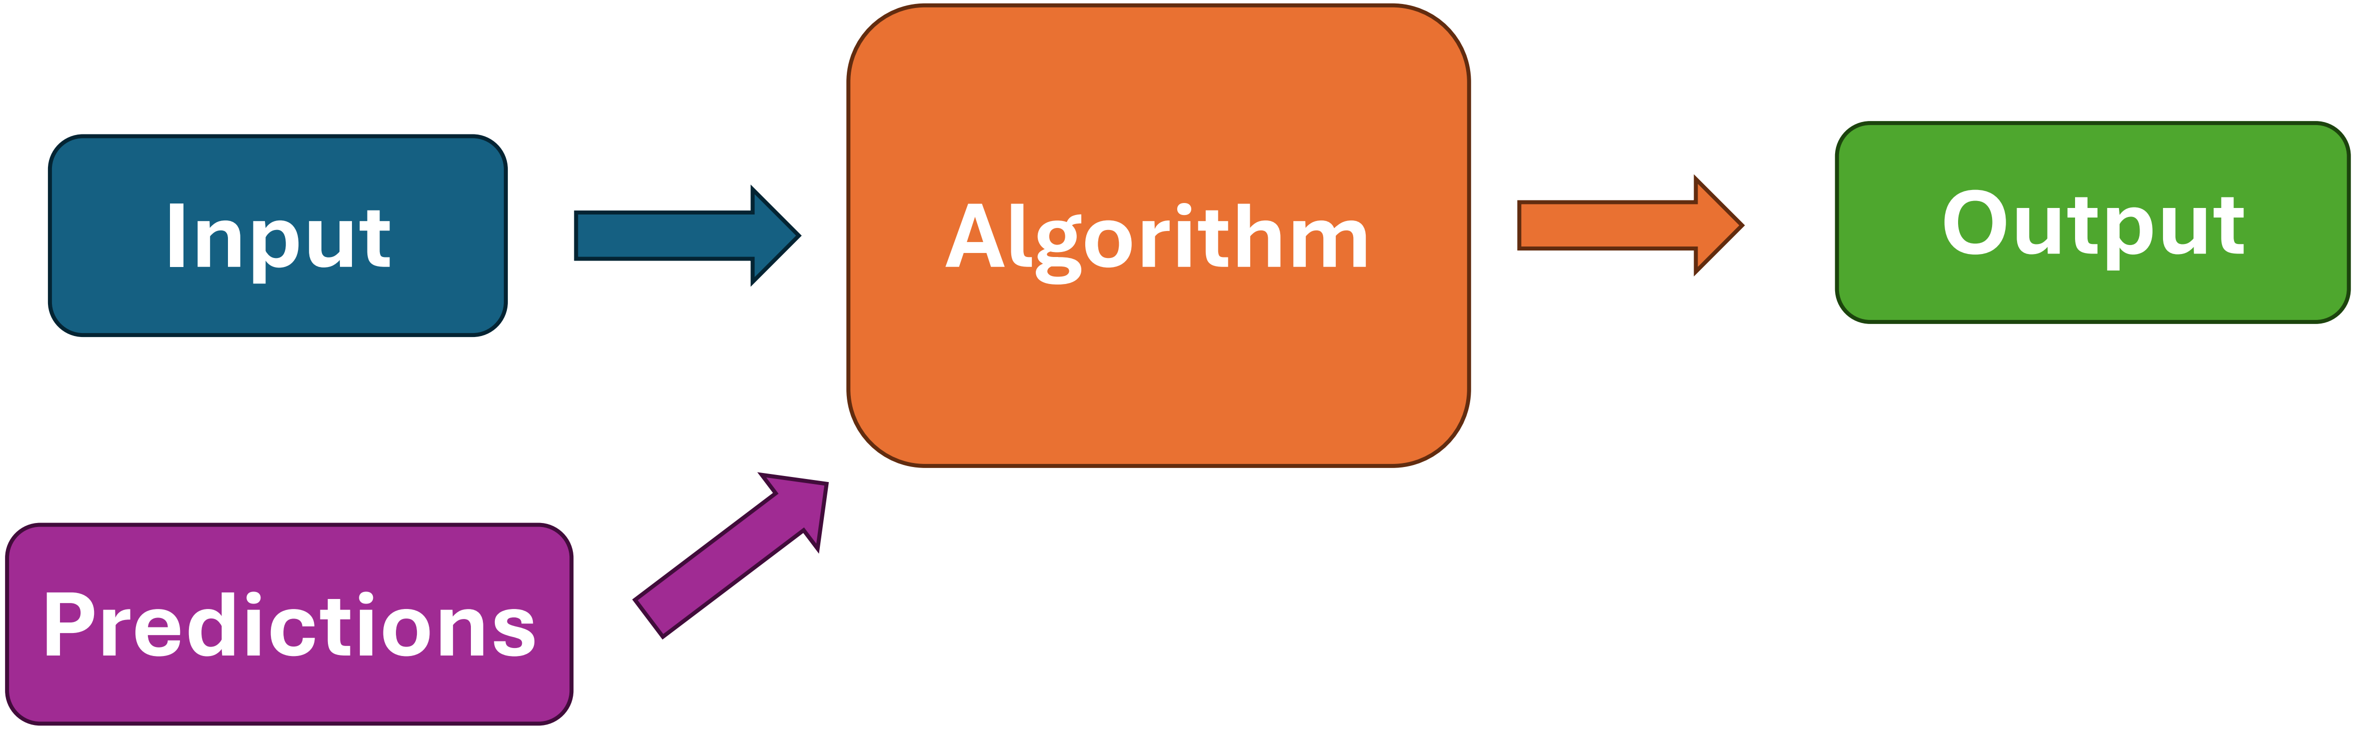

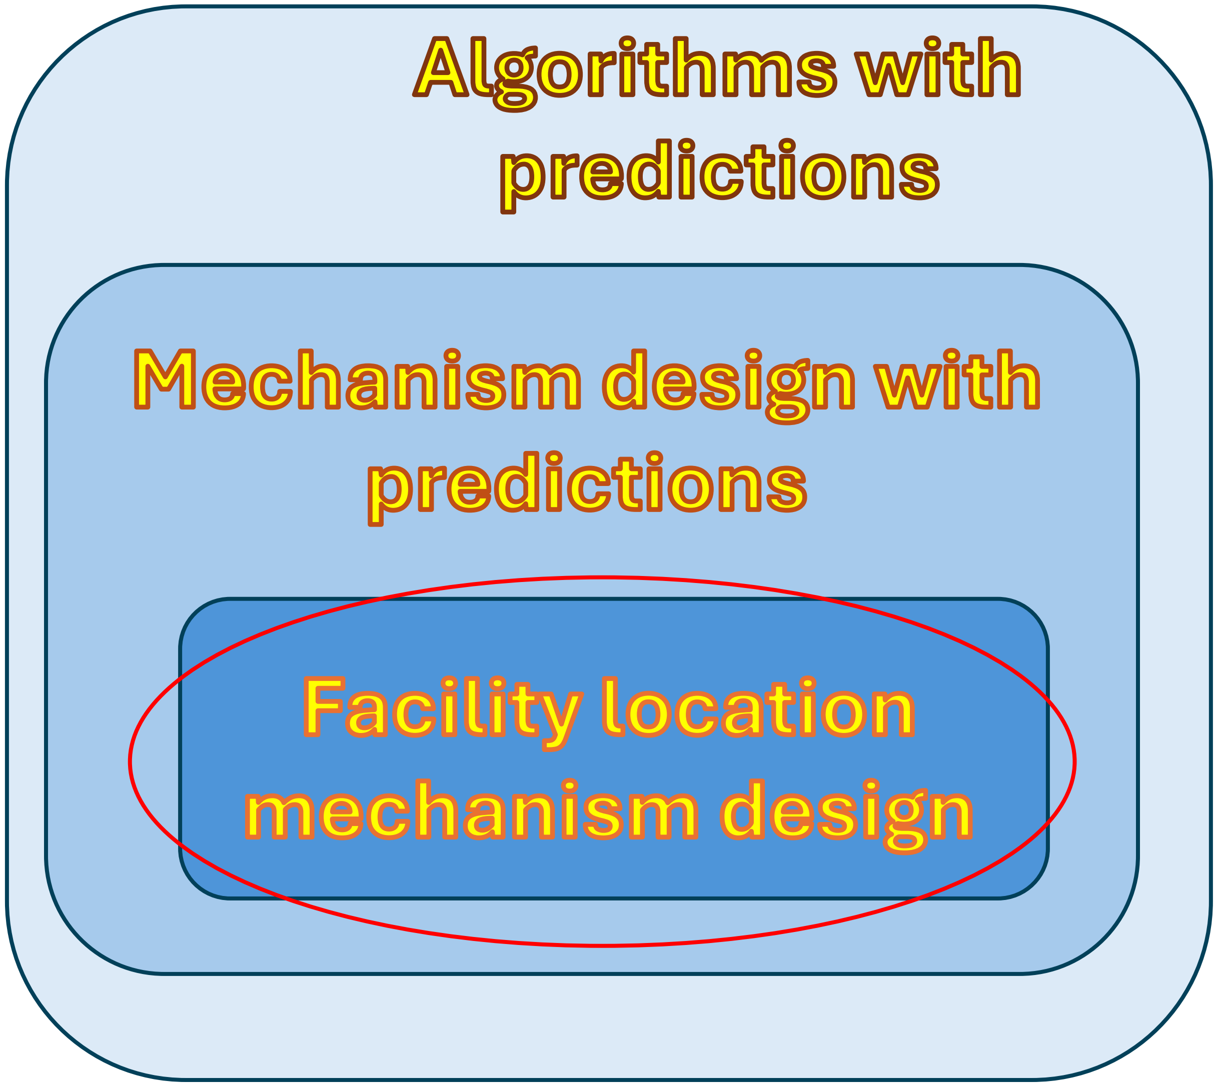

[Lykouris, Vassilvitskii 2020]

[Agrawal, Balkanski, Gkatzelis, Ou, Tan 2022]

[Xu and Lu 2022]

…

k Facility location mechanism design

Problem definition (no predictions)

k Facility location mechanism design

Input:
n **locations of strategic agents** in a metric space (each agent reports her location).

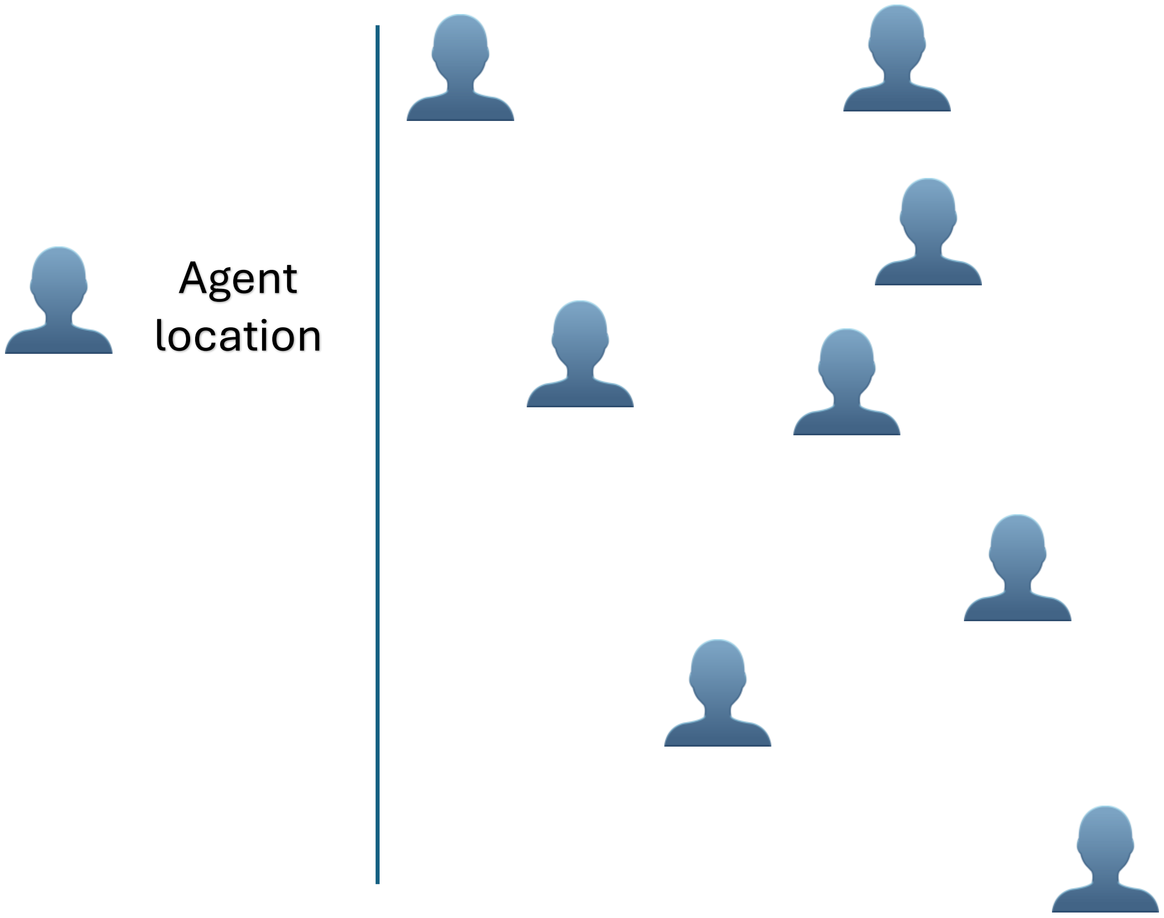

k Facility location mechanism design

<u>Input:</u>

n **locations of strategic agents** in a metric space (each agent reports her location).

<u>Goal:</u>

Return the **k facility locations** closest to the agents (minimize the **sum** of distances)

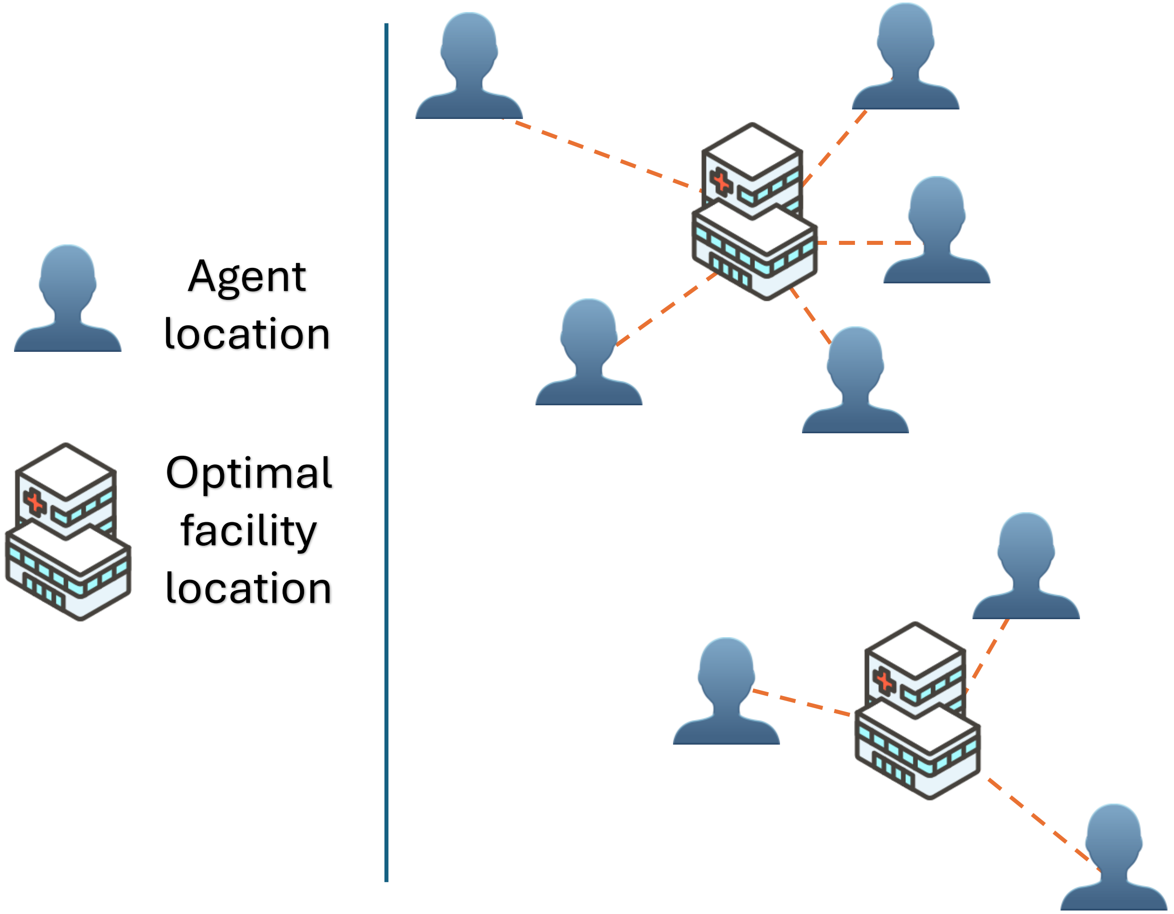

k Facility location mechanism design

Input:
n **locations of strategic agents** in a metric space (each agent reports her location).

Goal:
Return the **k facility locations** closest to the agents (minimize the **sum** of distances)

We need to design a **strategyproof mechanism** to do this.

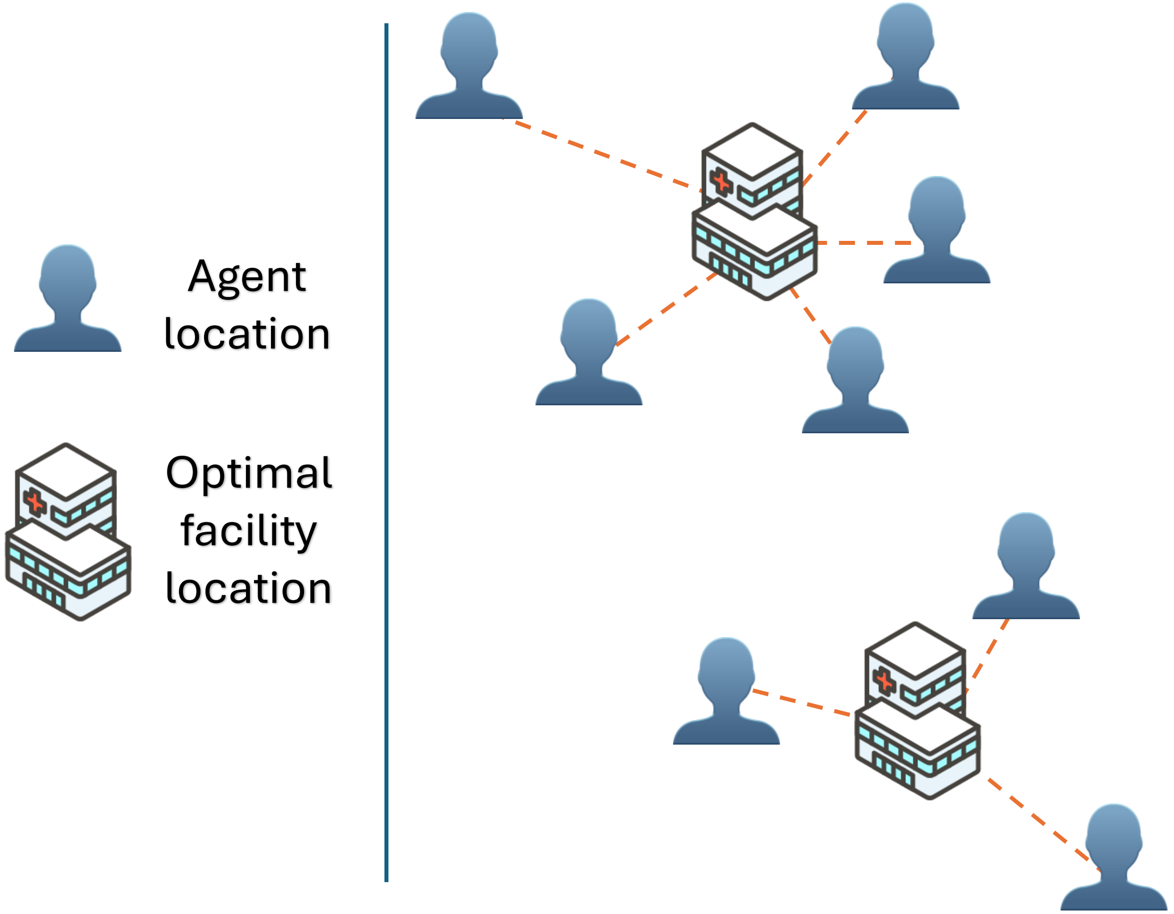

Setting with Predictions

In addition to the input, we also have predictions for each agent location.

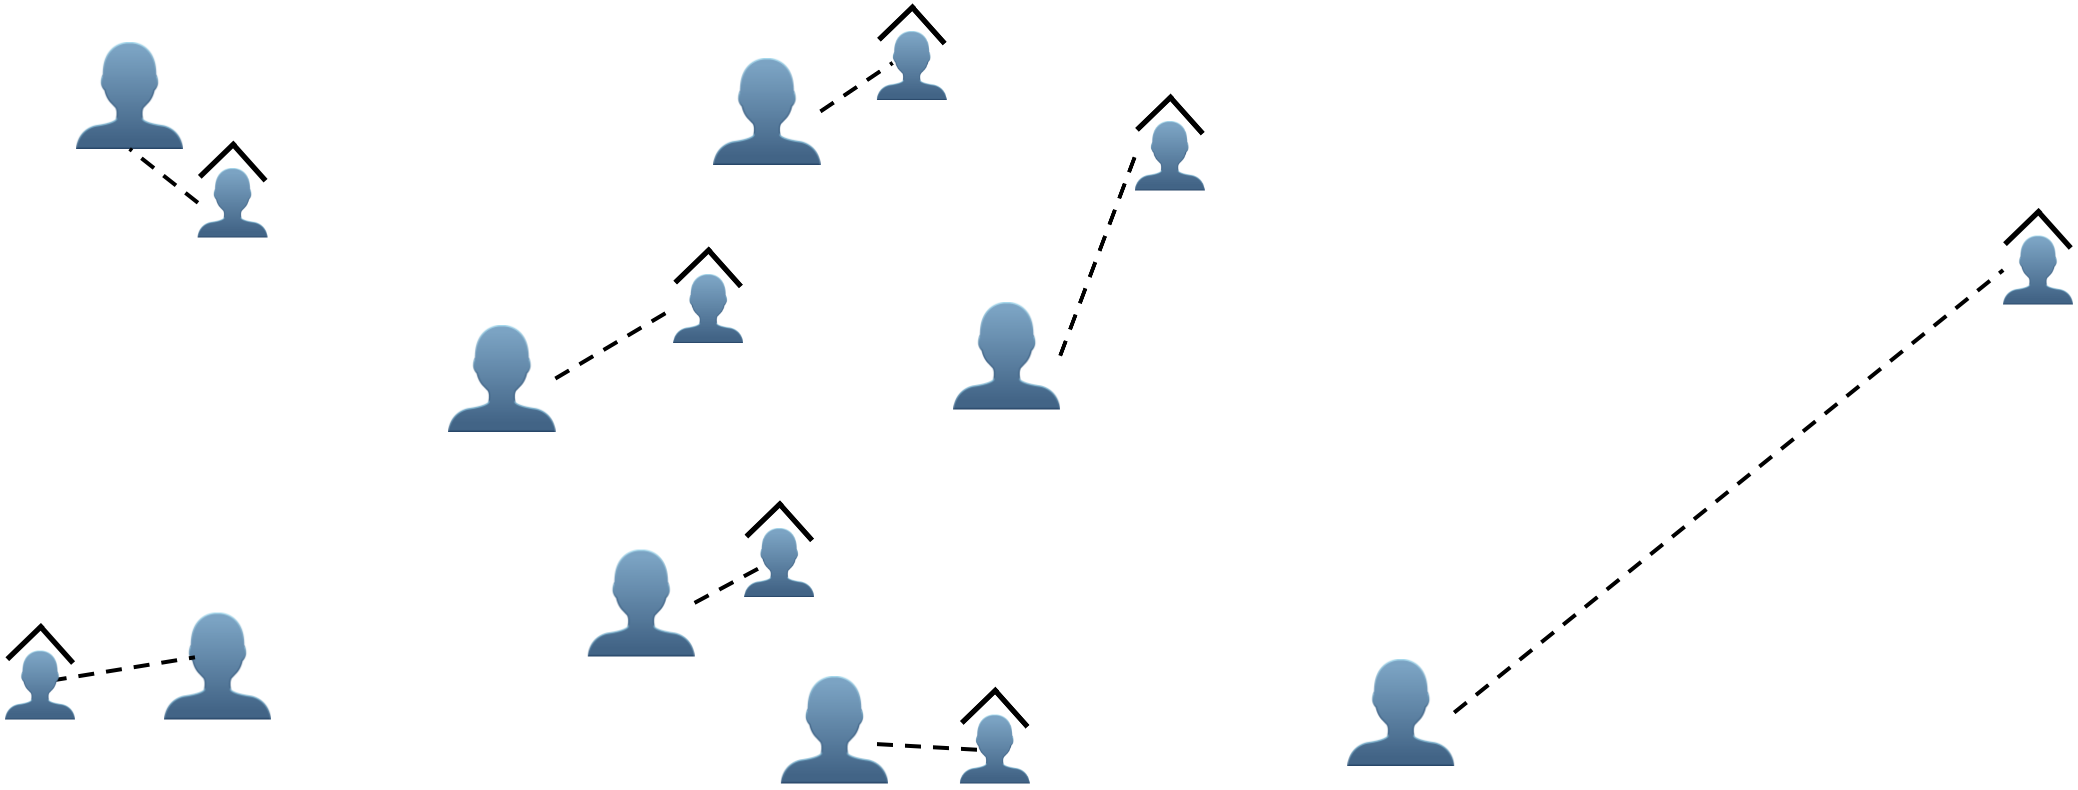

Determining the error model

If the predictions are "good" we want good performance.

What are "good predictions"?

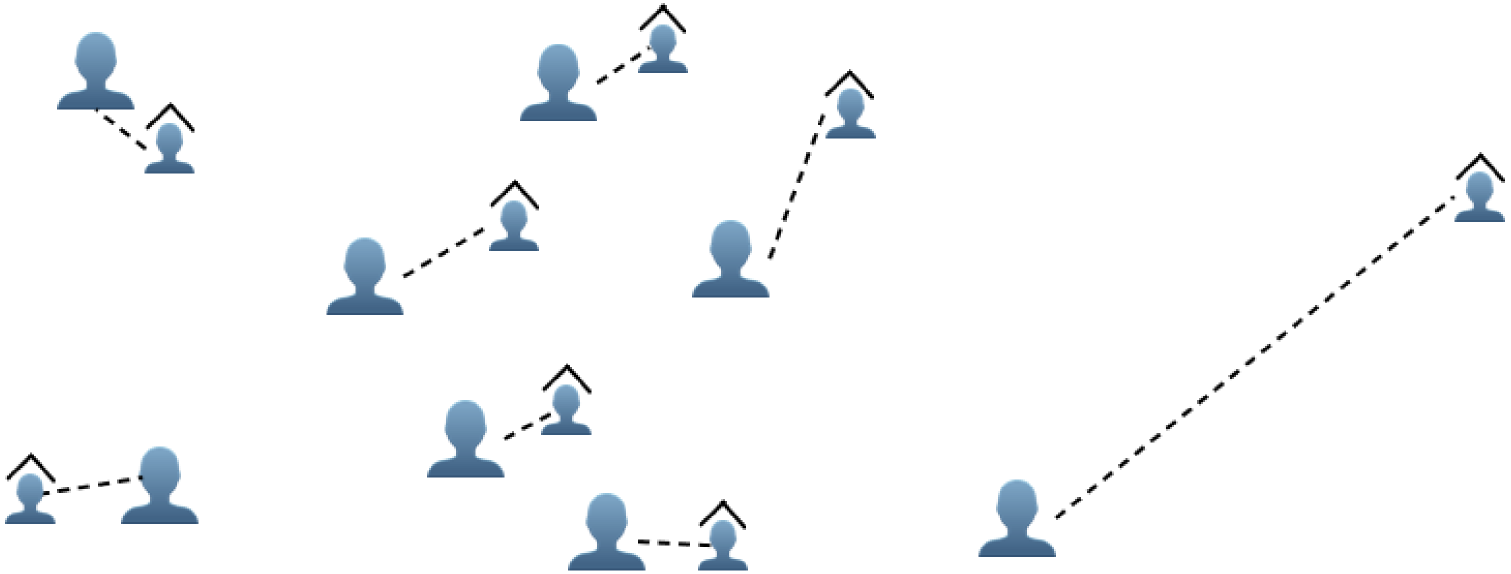

Standard "Worst case" error

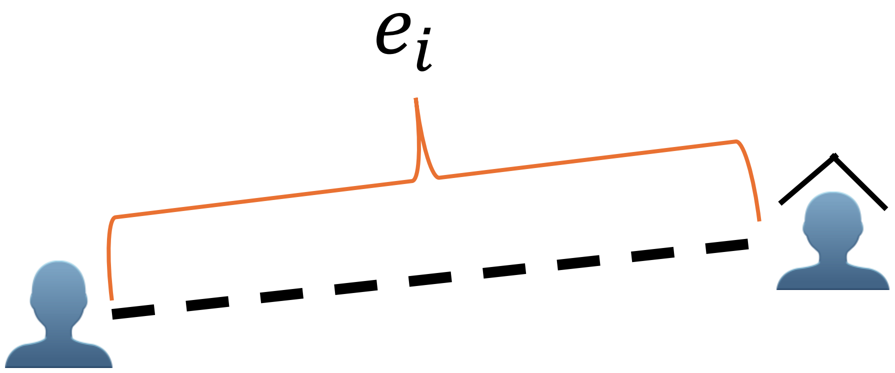

Error of each prediction i: e_i

Standard "Worst case" error η:

$$\eta := \sum_i e_i$$

$$or$$

$$\eta := \max_{\boldsymbol{i}} \boldsymbol{e_i}$$

Our model: MAC (ε, δ) predictions

MAC = Mostly Approximately Correct:

Most (At least $1-\delta$ fraction) of the predictions are **approximately correct** (up to an ε additive error).

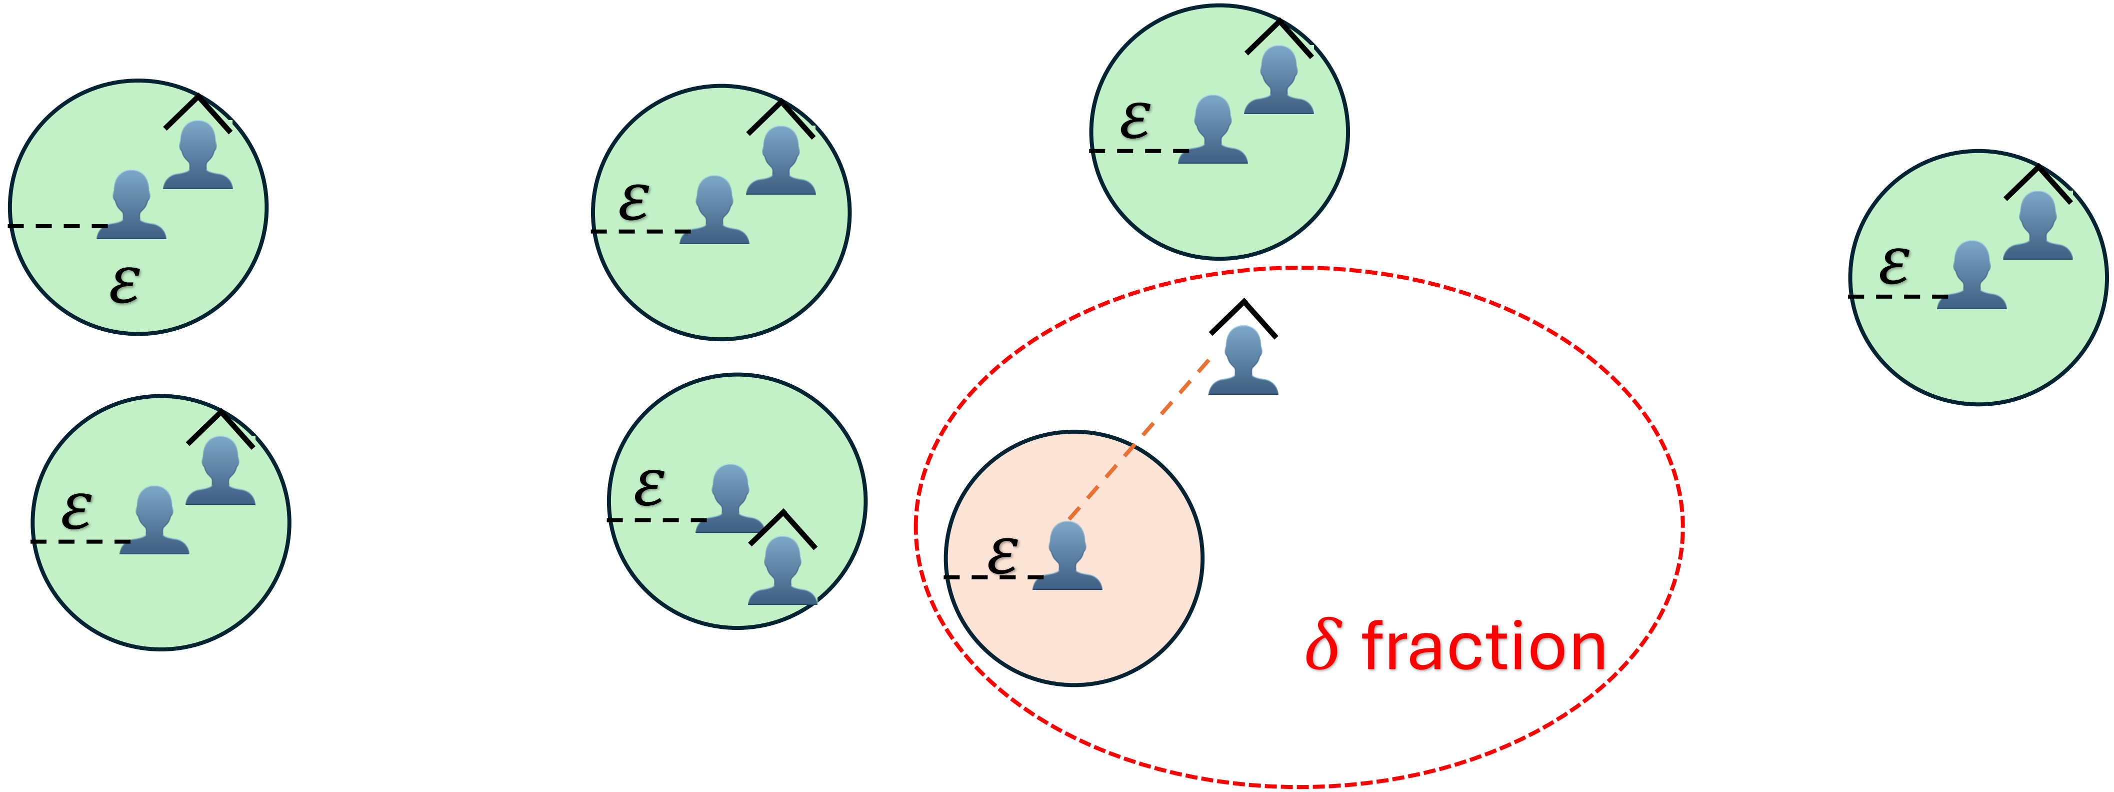

Results*

Deterministic mechanism design

Problem	Best known "no predictions" approximation ratio	MAC predictions approximation ratio
Single facility in $\mathbb{R}^d$	$\sqrt{d}$ [Meir et al. 2019]	$1 + O(\delta)$
β-balanced k facilities in $\mathbb{R}^d$	Linear ($O(n)$) [Aziz et al. 2020]	A constant depending on δ, β, k

Results*

Deterministic mechanism design

Problem	Best known "no predictions" approximation ratio	MAC predictions approximation ratio
Single facility in $\mathbb{R}^d$	$\sqrt{d}$ [Meir et al. 2019]	$1 + O(\delta)$
β-balanced k facilities in $\mathbb{R}^d$	Linear ($O(n)$) [Aziz et al. 2020]	A constant depending on δ, β, k

Random mechanism design

Problem	Best known "no predictions" approximation ratio	MAC predictions Approximation ratio
2 facility locations on a line	4 [Lu et al. 2010]	3.6 + $O(\delta)$

*ε is omitted from the results, as ε introduces another small additive term.

Results*

Deterministic mechanism design

Problem	Best known "no predictions" approximation ratio	MAC predictions approximation ratio
Single facility in $\mathbb{R}^d$	$\sqrt{d}$ [Meir et al. 2019]	$1 + O(\delta)$
β-balanced k facilities in $\mathbb{R}^d$	Linear ($O(n)$) [Aziz et al. 2020]	A constant depending on δ, β, k

Random mechanism design

Problem	Best known "no predictions" approximation ratio	MAC predictions Approximation ratio
2 facility locations on a line	4 [Lu et al. 2010]	3.6 + $O(\delta)$

Much harder setting in the presence of outliers

*ε is omitted from the results, as ε introduces another small additive term.

Techniques of independent interest

- **Distance and approximation robustness**
  - **Quantitative versions** of the known notion of **“breakdown point”** in robust statistics.
- **Second facility location problem**
  - Separate interesting “no predictions” mechanism design problem

Future directions

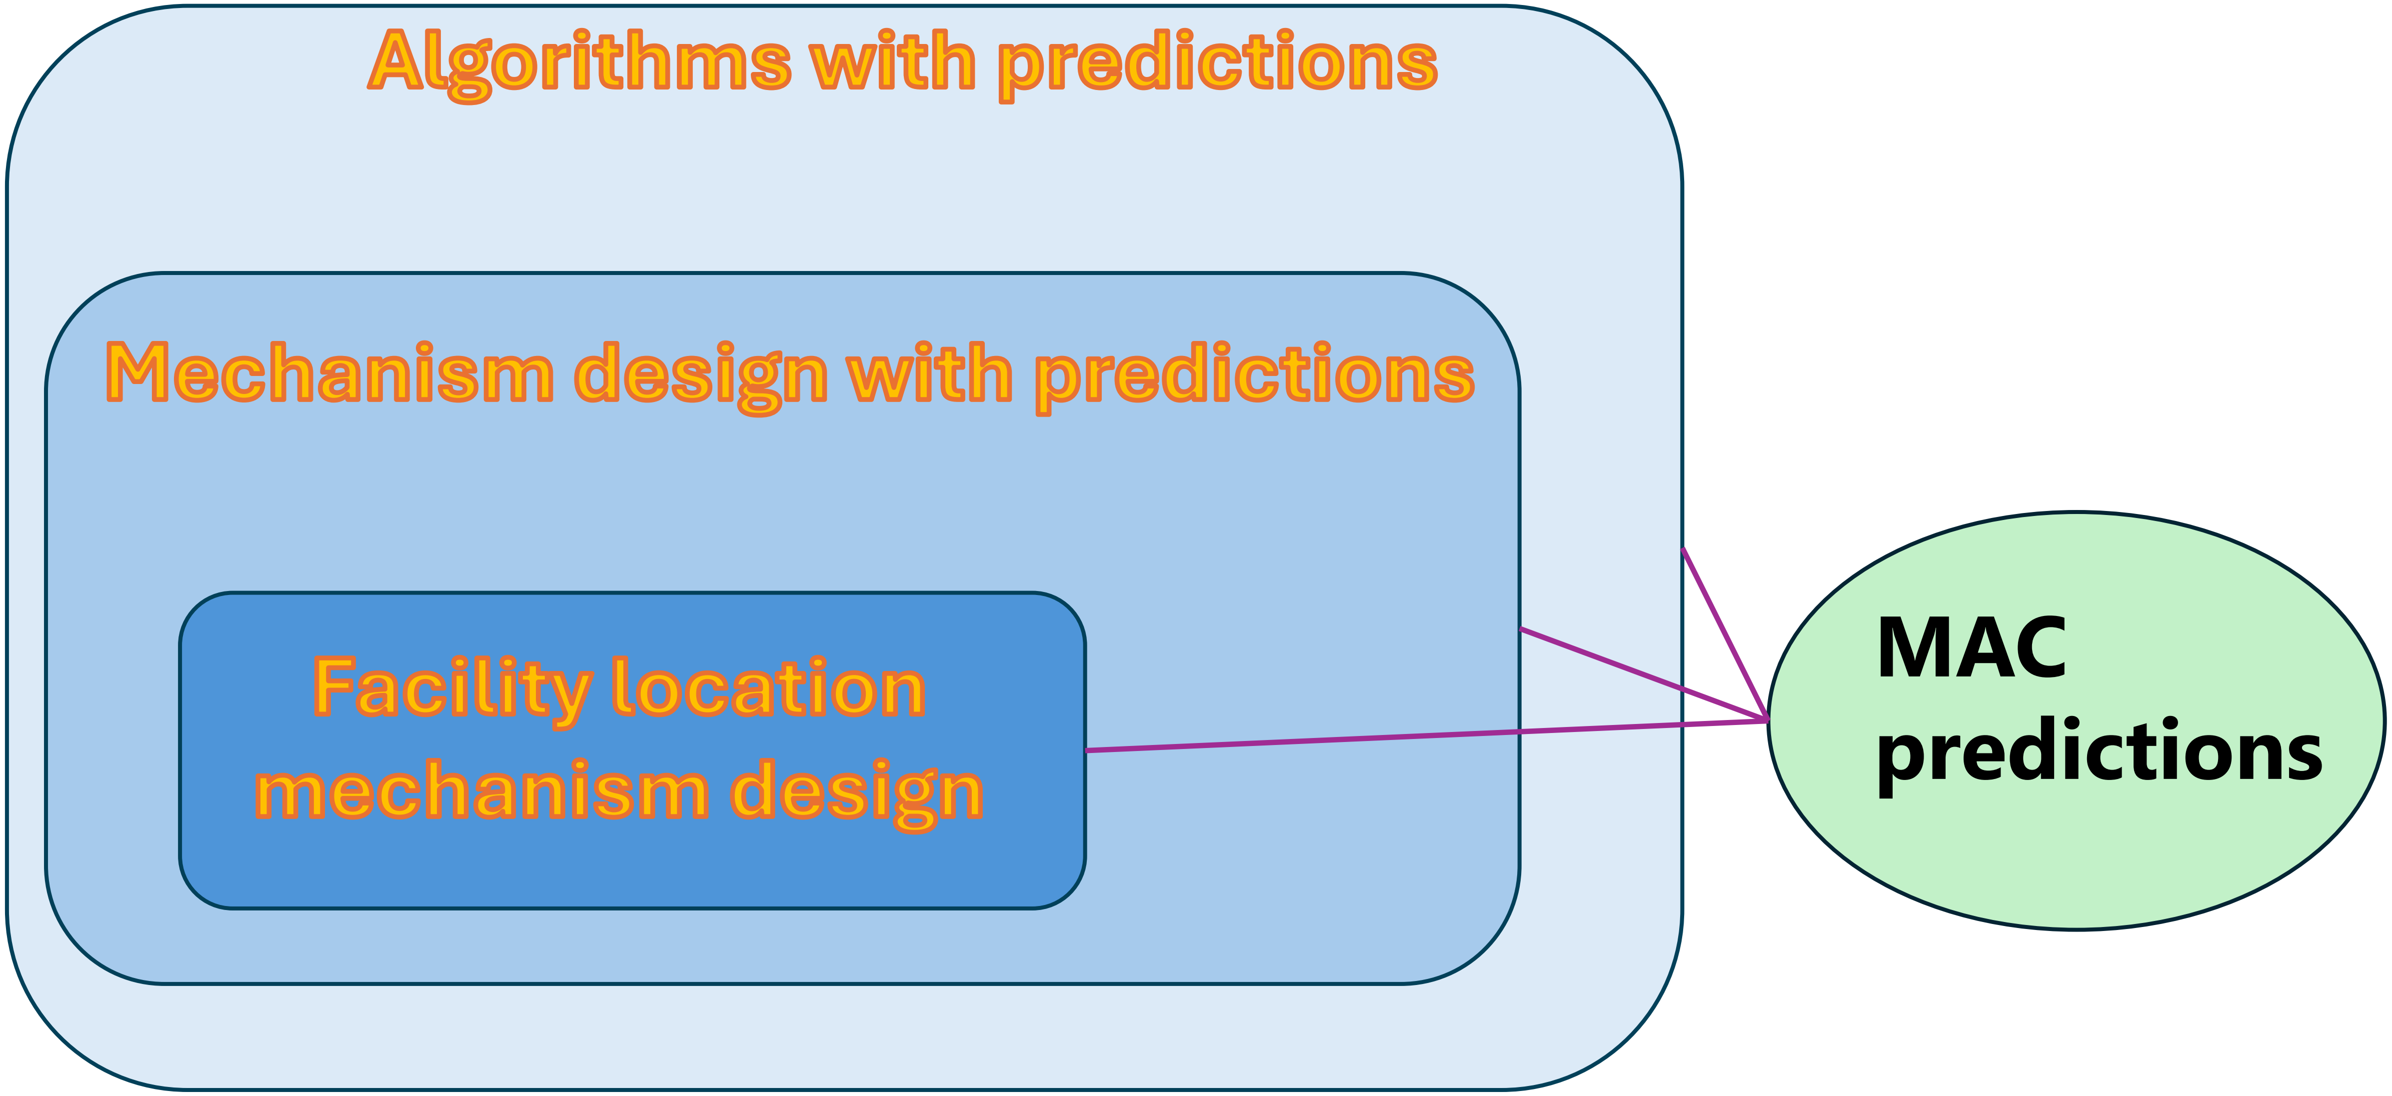

Thanks for listening