

MATE: Benchmarking Multi-Agent Reinforcement Learning in Distributed Target Coverage Control

Xuehai Pan Mickel Liu Fangwei Zhong
Yaodong Yang Song-Chun Zhu Yizhou Wang

June 9th, 2022

1 MATE: the Multi-Agent Tracking Environment

- Problem Settings
- Scalable
- Communicative Agents
- Heterogeneous & Asymmetric

2 Results

- Learning Cameras in Fully-Cooperative Game
- Learning Targets in Fully-Cooperative Game
- Zero-Sum Fully-Competitive Game

Outline

1 MATE: the Multi-Agent Tracking Environment

- Problem Settings
- Scalable
- Communicative Agents
- Heterogeneous & Asymmetric

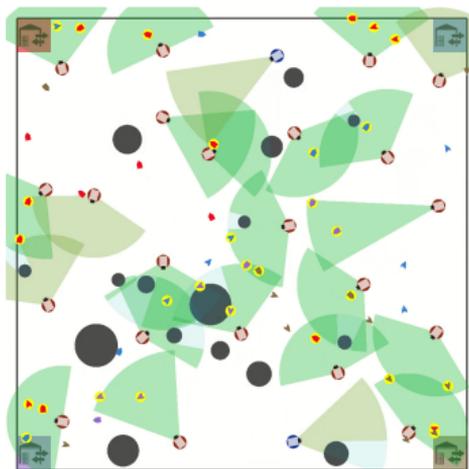
2 Results

- Learning Cameras in Fully-Cooperative Game
- Learning Targets in Fully-Cooperative Game
- Zero-Sum Fully-Competitive Game

MATE: the Multi-Agent Tracking Environment

An **asymmetric heterogeneous competitive-cooperative** multi-agent environment under **partial observations**.

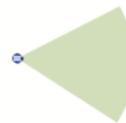
- **Cameras:** collaboratively maximize coverage on “targets” while minimizing repeated detection of the same “target”.
- **Targets (Vehicles):** transport cargoes between warehouses while minimizing “cameras” surveillance.



(a) 100C vs. 20T

(b) 24C vs. 48T

• Camera



• Target



Speed ++
Capacity +



Speed +
Capacity ++

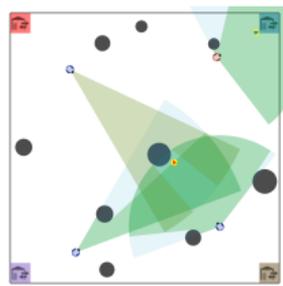
• Obstacle



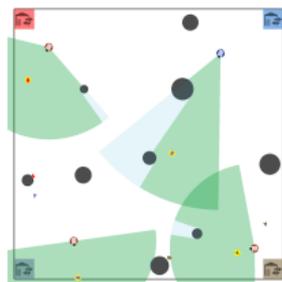
• Warehouse



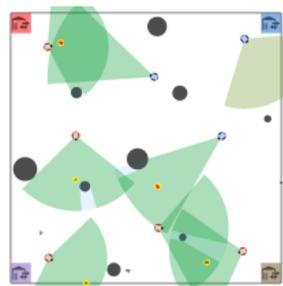
MATE: the Multi-Agent Tracking Environment



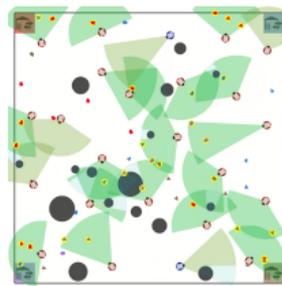
(a) 4C vs. 2T



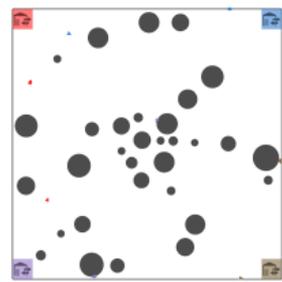
(b) 4C vs. 8T



(c) 8C vs. 8T



(d) 24C vs. 48T



(e) Navigation

Five snapshots of The Multi-Agent Tracking Environment at different scales. Note that here we abbreviate “camera” as “C”, “target” as “T”.

Problem Settings

The two sides of the opposing team constitute a **zero-sum** game.

- **Target Team** (vehicles):
 - There are several vehicles in charge of transporting cargo, with two types, one with **low speed and large capacity**, and the other with **high speed and small capacity**.
 - The vehicle will receive a reward and new cargoes when it transports its payload to the destination.
- **Camera Team** (sensor network):
 - The surveillance sensor network contains several **proactive** cameras distributed in the scene.
 - Each camera can independently adjust its orientation and perception range (the viewing angle and radius of the sector area).
- **Warehouses**: distributed in the environment, and each warehouse has several (random number) cargoes to be transferred to the others.
- **Obstacles**: randomly distributed in the scene. The camera cannot observe the targets behind the obstacle, and the target (vehicle) cannot enter the obstacle.

“Camera” Agent

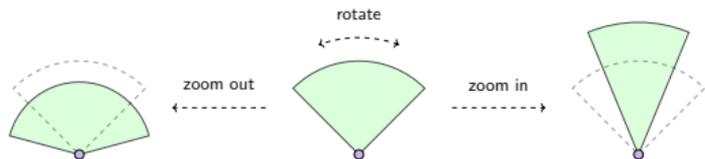
Camera: in-place zoomable directional sensor (the purple colored entity at center)

Observation:

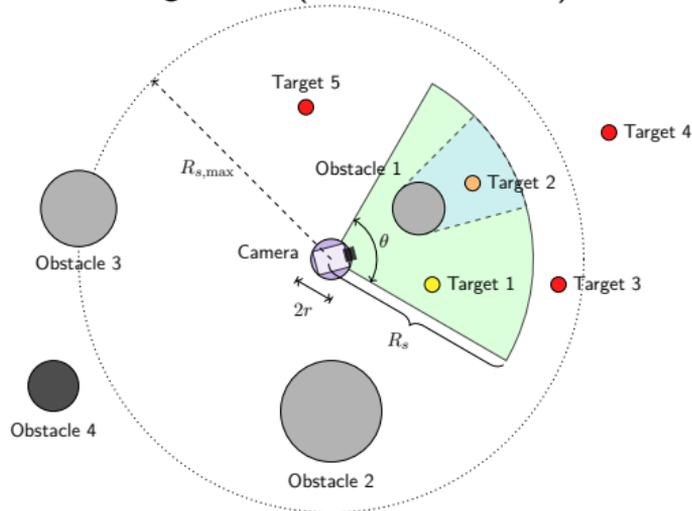
- the location of the vehicles (targets) and whether they are carrying cargo in the perception zone (the green colored area / blind spot (the cyan colored area))
- the location and size of obstacles within its maximum sensing radius (the dotted circle).

Action:

- rotate the orientation
- zoom in and out of the viewing angle.



Camera Action: rotation and zooming



Camera Observation (within the green sector)

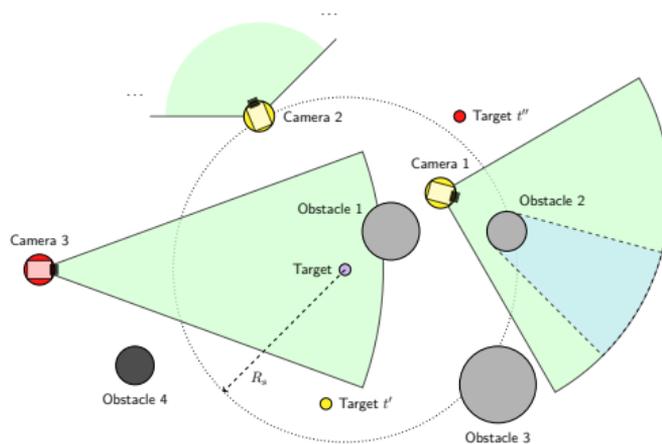
“Target” Agent

Target: movable vehicle equipped with omnidirectional sensors

Observation (in dotted circle):

- the camera location and pose (direction and opening angle)
- the location and size of obstacles
- the location of other vehicles (targets) and whether they are carrying cargo

Action: displacement relative to the current position ($\Delta x, \Delta y$)



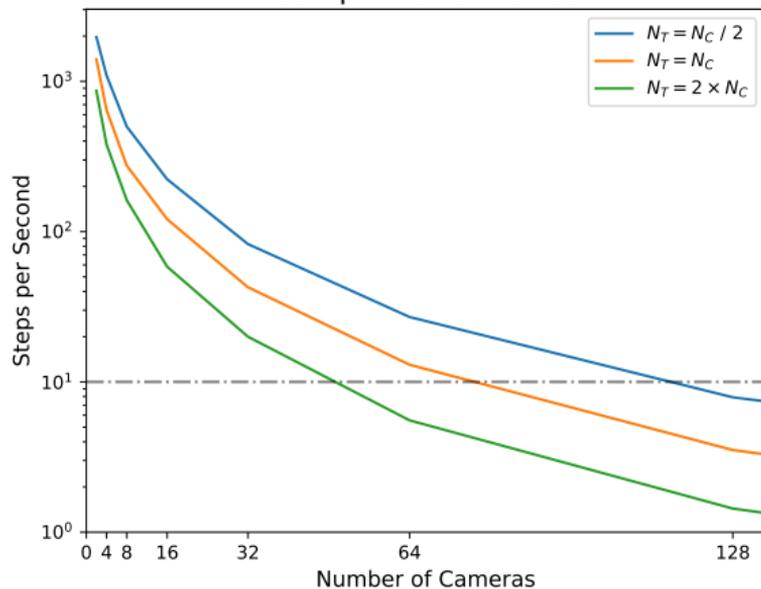
Target Observation (within dotted circle)

Core Features

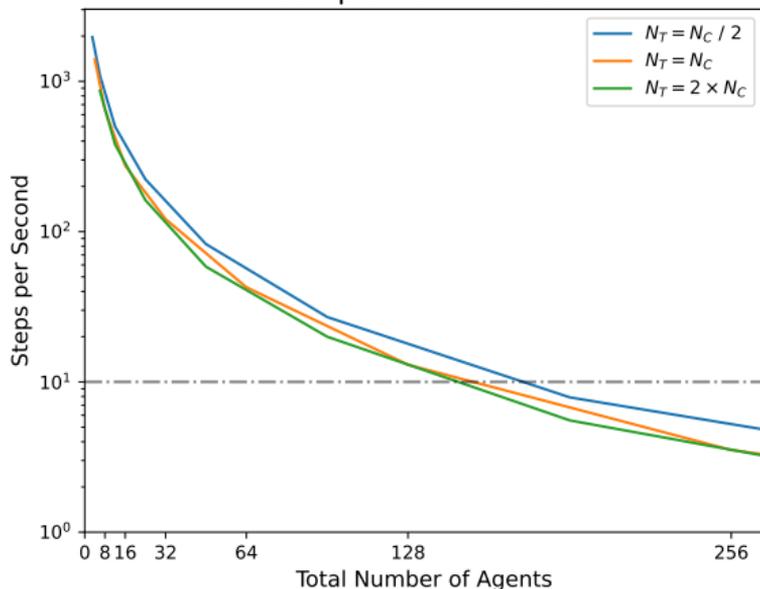
- **Sample Efficient and Easy to Use:**
 - 300+ FPS
 - easy to setup and out-of-box usable
- **Scalable**
 - up to 100+ agents
- **Communicative**
 - explicit communication channels
 - peer-to-peer and broadcast
- **Heterogeneous and Asymmetric**
- **Various Scenarios and Curricula Configurations**
 - built-in agents
 - built-in wrappers

Up to hundreds of agents interact with the environment simultaneously.

Sample Performance



Sample Performance



Communicative Agents

- Obtain extra information (partial observation)
- Sharing information and knowledge
- Reduce unnecessary exploration

More Effective Cooperation

Two types:

- peer-to-peer
- broadcast

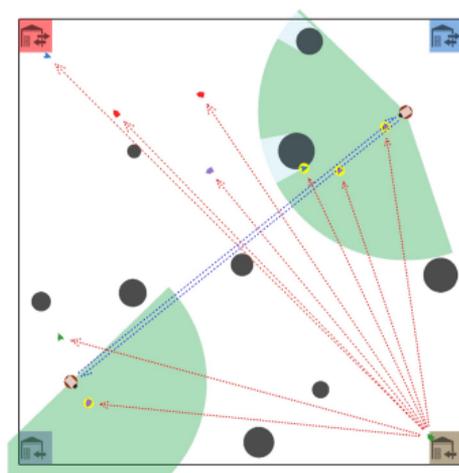
Multi-rounds:

- request
- response

⇒ negotiate

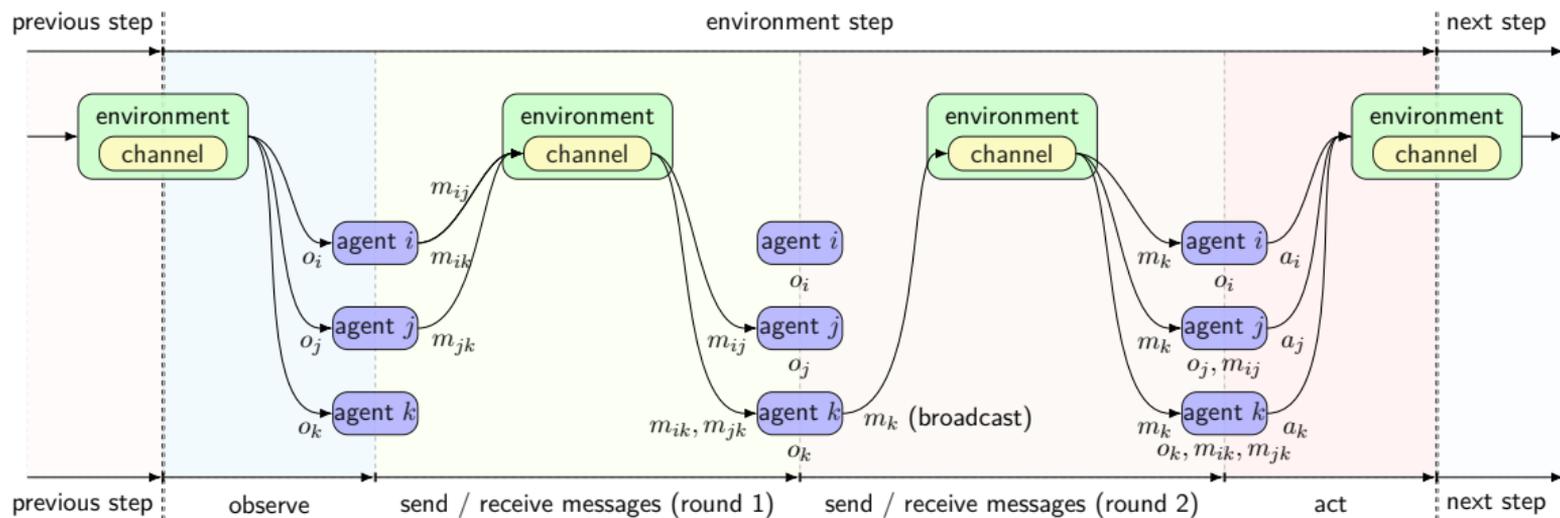
Realistic:

- noise
- delay
- range restriction
- ...



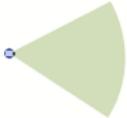
- Cameras' intra-team ----->
- Targets' intra-team ----->

Communicative Agents



Multi-round Communication Mechanism

Communicative Agents

- Camera 
- Target 
- Obstacle 
- Warehouse 

Rule-based Greedy Cameras vs. Rule-based Greedy Targets

Heterogeneous & Asymmetric

Heterogeneous:

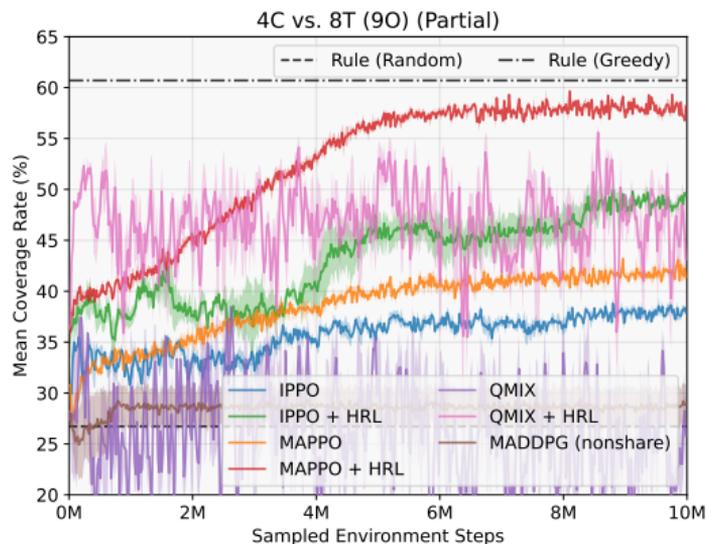
- **Inter-team:** the two types of agents have completely different dynamics and tasks.
- **Intra-team:** the “target” agents can use two different vehicles...
 - high speed and small capacity
 - low speed and large capacity

Asymmetric:

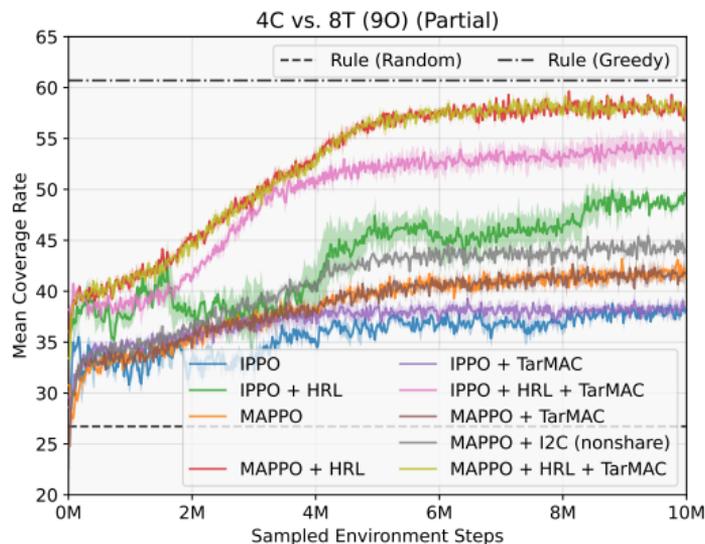
- The two sides' agents are heterogeneous.
- The game balance is affected by the capabilities of both agents.
- The size of both teams in the game is variable.

- 1 MATE: the Multi-Agent Tracking Environment
 - Problem Settings
 - Scalable
 - Communicative Agents
 - Heterogeneous & Asymmetric
- 2 Results
 - Learning Cameras in Fully-Cooperative Game
 - Learning Targets in Fully-Cooperative Game
 - Zero-Sum Fully-Competitive Game

Learning Cameras in Fully-Cooperative Game

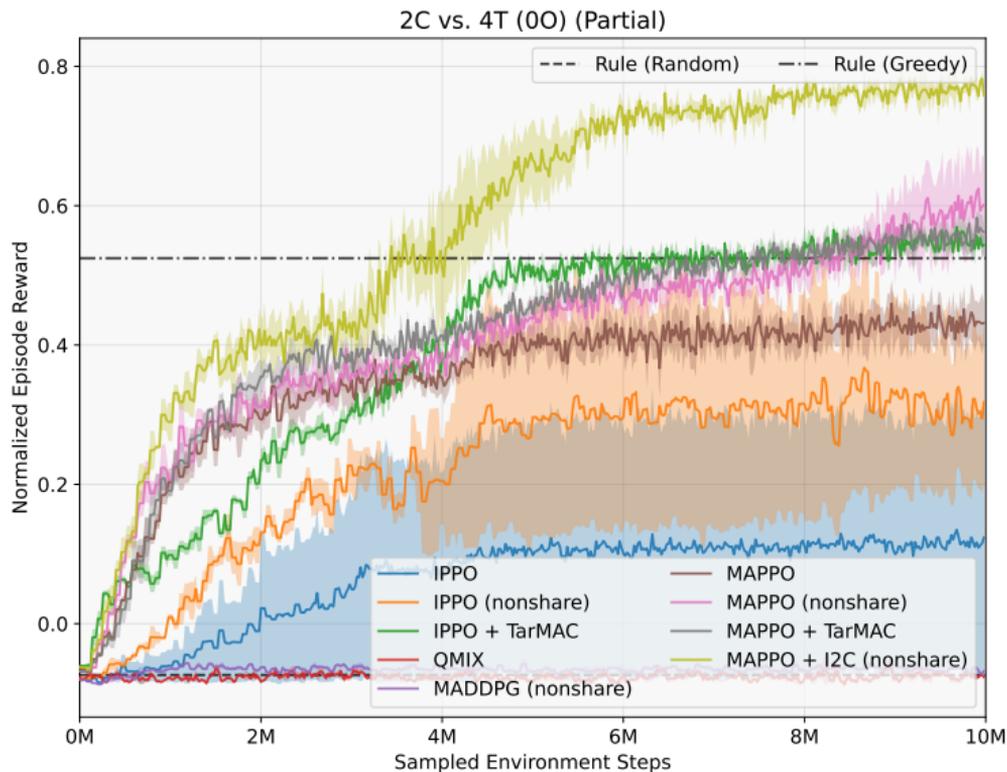


(a) MARL baselines



(b) MARL with multi-agent communication

Learning Targets in Fully-Cooperative Game



Learning Targets in Fully-Cooperative Game

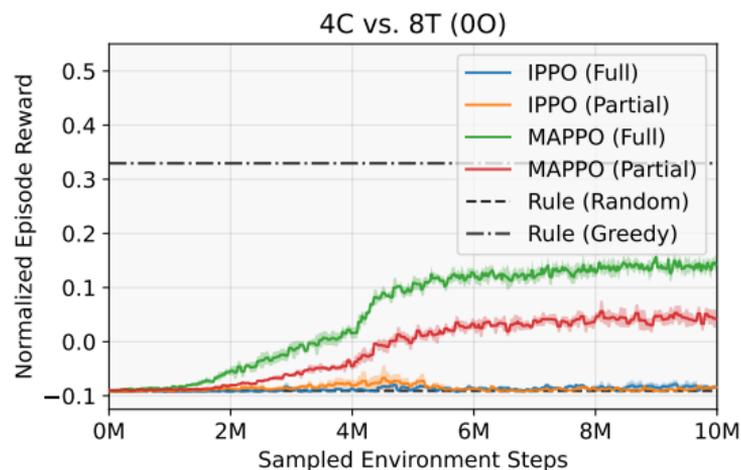
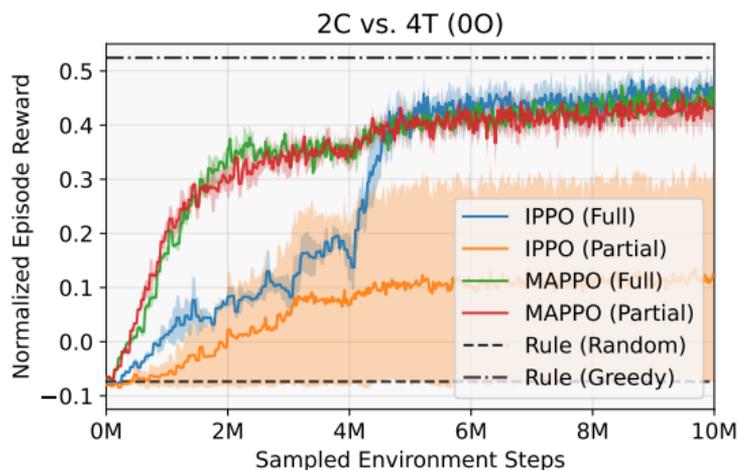
Emergence of Roles via Multi-Agent Communication



北京大学
PEKING UNIVERSITY

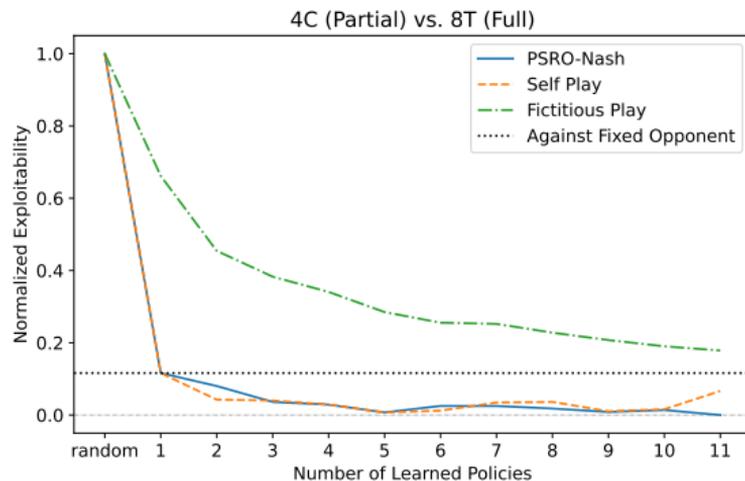
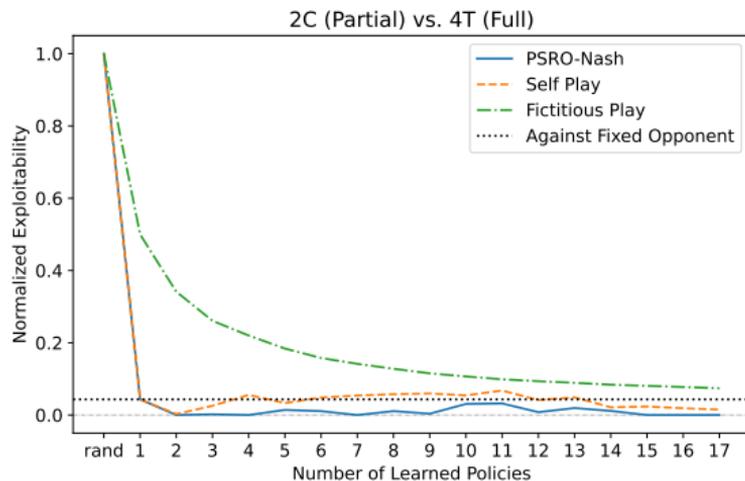
Learning Targets in Fully-Cooperative Game

Progressive Difficulty in Configurable Environments



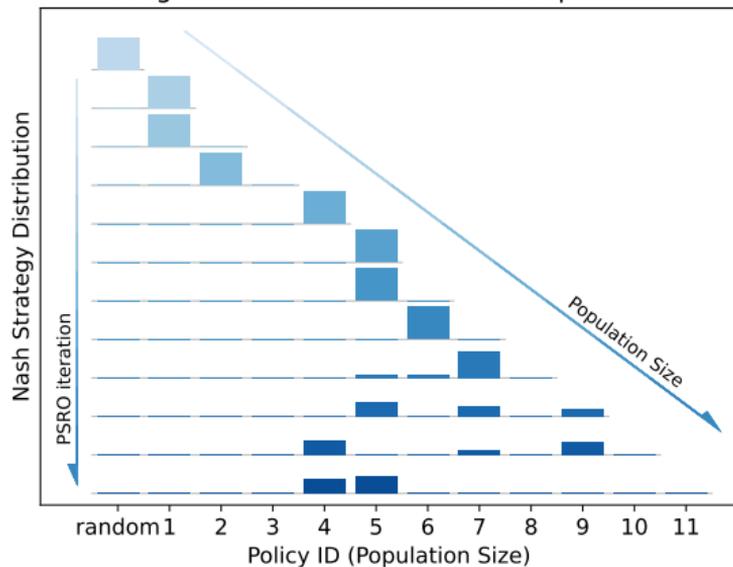
Zero-Sum Fully-Competitive Game

$$\text{exploitability}(\Pi^R, \Pi^C) = \frac{1}{2} \sum_{i \in \{R, C\}} [U^i(\mathbf{BR}(\Pi^{-i}), \Pi^{-i}) - U^i(\Pi^i, \Pi^{-i})],$$



Zero-Sum Fully-Competitive Game

Progression of Nash of Camera Population



Progression of Nash of Target Population

