

IMP-MARL: a Suite of Environments for Large-scale Infrastructure Management Planning via MARL

Pascal Leroy, Pablo G. Morato, Jonathan Pisane, Athanasios Kolios, Damien Ernst

Paper: <https://arxiv.org/abs/2306.11551>


Project page: https://github.com/moratodpg/imp_marl

Motivation

Multi-agent reinforcement learning (MARL)

- Cooperative settings
- Many available **open-source** methods
- Common benchmarks are **games** or **simulators**

Gap:
Only few **real-world** environments
Only few **large-scale** environments



SMAC, MaMuJoCo, MPE, Hanabi,...

Motivation

Infrastructure management planning (IMP)

- Impactful **real-world** application
- Inspections and maintenance planning, minimising system failure risk and maintenance costs
- Effective multi-component policies can be learned via **MARL**

Gap:

Most studies are **not open-sourced**

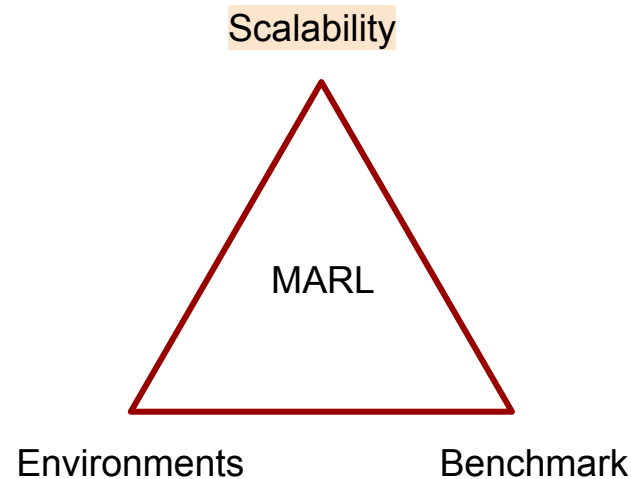
Most development **are not compared** against SOTA MARL algorithms



Wind turbines, bridges, water networks, ...

IMP-MARL: Main contributions

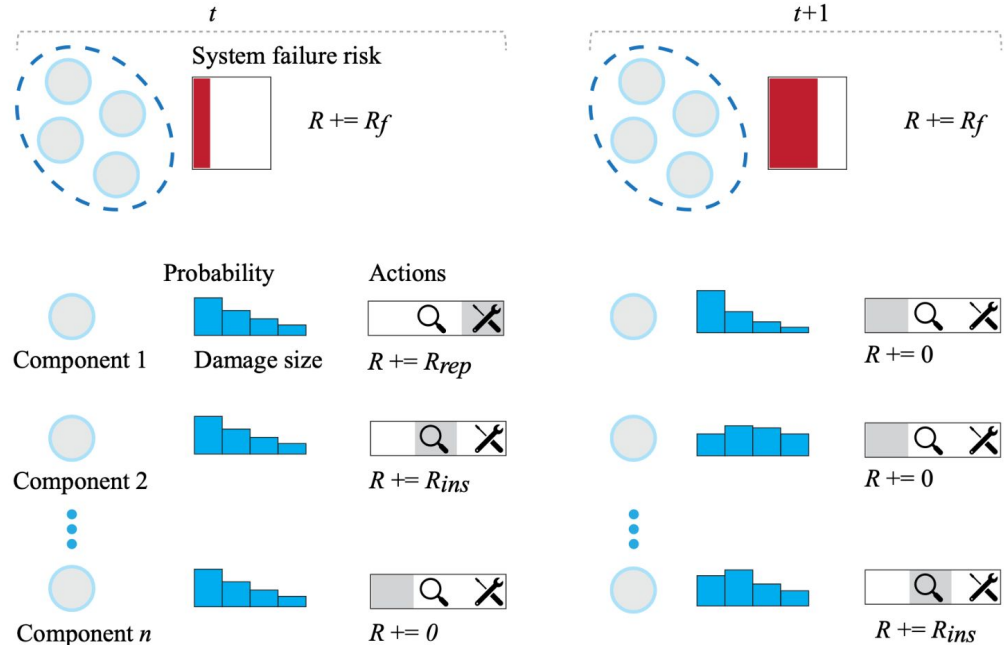
1. A novel **open-source** suite with **real-world** environments
“Up to 100 agents!”
2. Benchmark **SOTA** cooperative MARL methods
3. Highlight important challenges that must be resolved
“Are cooperative **MARL** methods **scalable?**”



Infrastructure Management Planning (IMP)


Managing n components:

- *Inspect or repair or do-nothing* based on components' damage probability?
- **System failure risk depends** on the **components' failure probability**
- **Goal:**
 - Minimise maintenance costs
 - Avoid system's failure
- **Challenge:**
 - Joint **action space** exponentially growing with n (number of agents)
 - Do-nothing action usually dominates (class imbalance)



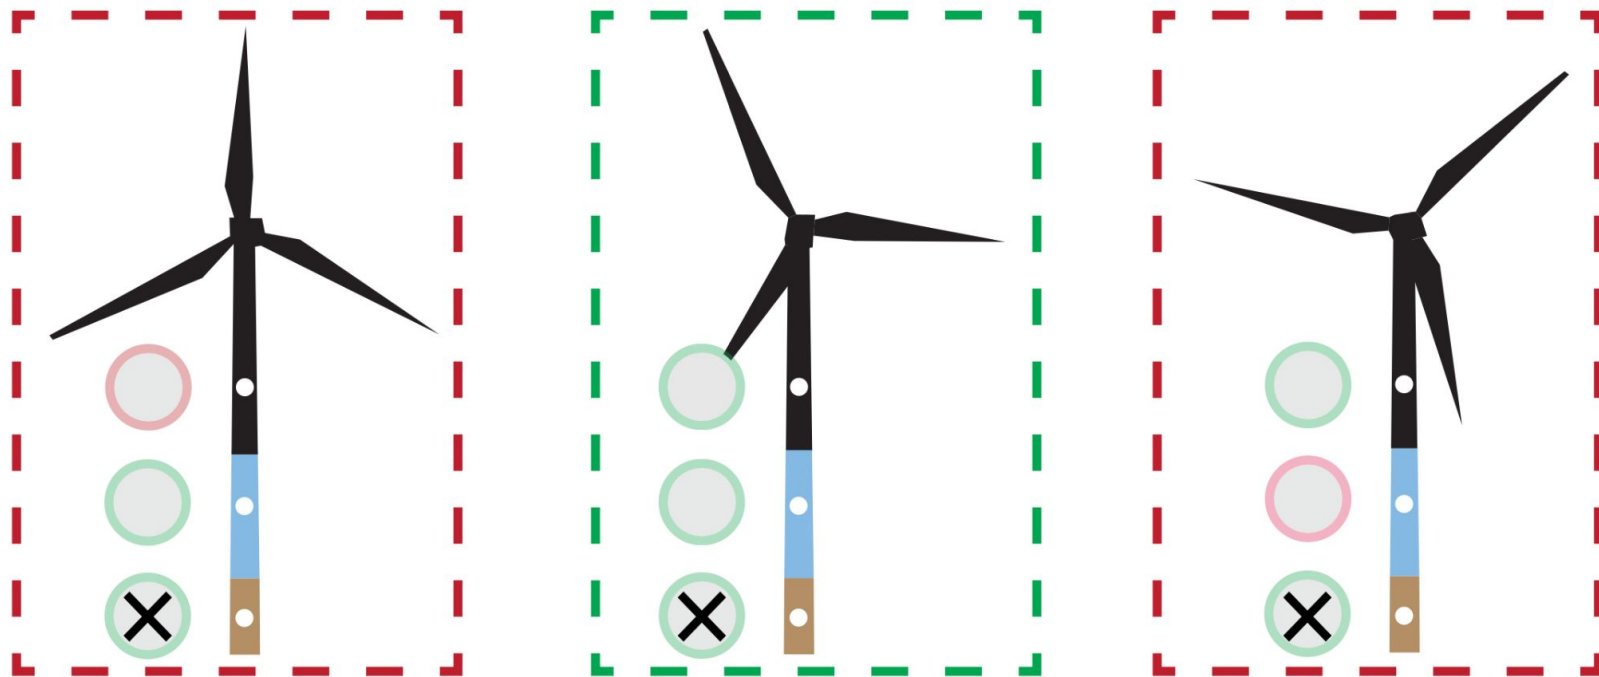
IMP-MARL: Environments

k-out-of-n systems: (4-out-of-5 system)



IMP-MARL: Environments

Offshore wind farm: 3 representative components per wind turbine



IMP-MARL benchmark

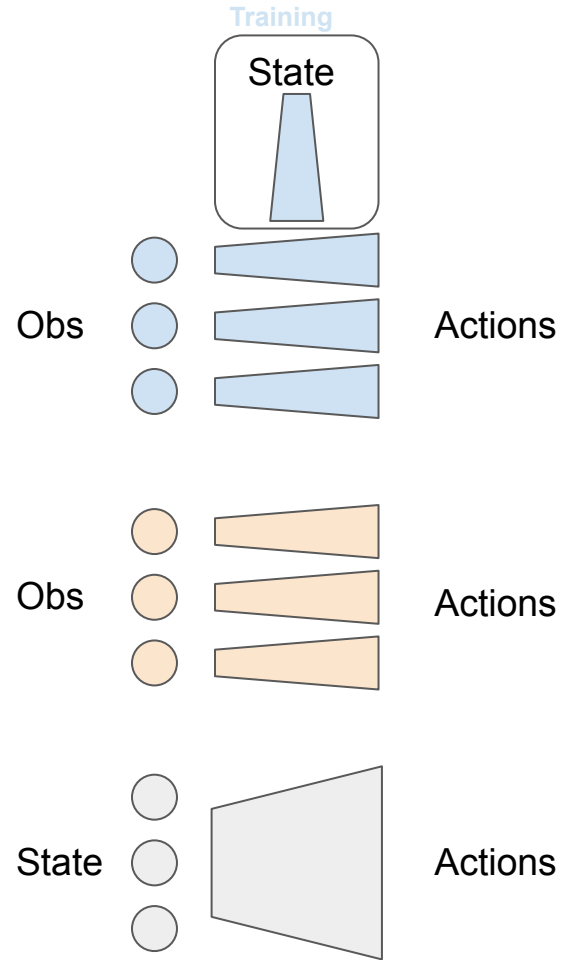
Goals:

- MARL vs heuristic ?
- Scalability ?

Methods

- Centralised training with decentralised execution: QMIX, QVMIX, QPLEX, COMA, FACMAC
- Decentralised: IQL (DQN for each agent)
- Centralised: DQN
- Heuristic: rule-based **baseline** from the reliability community

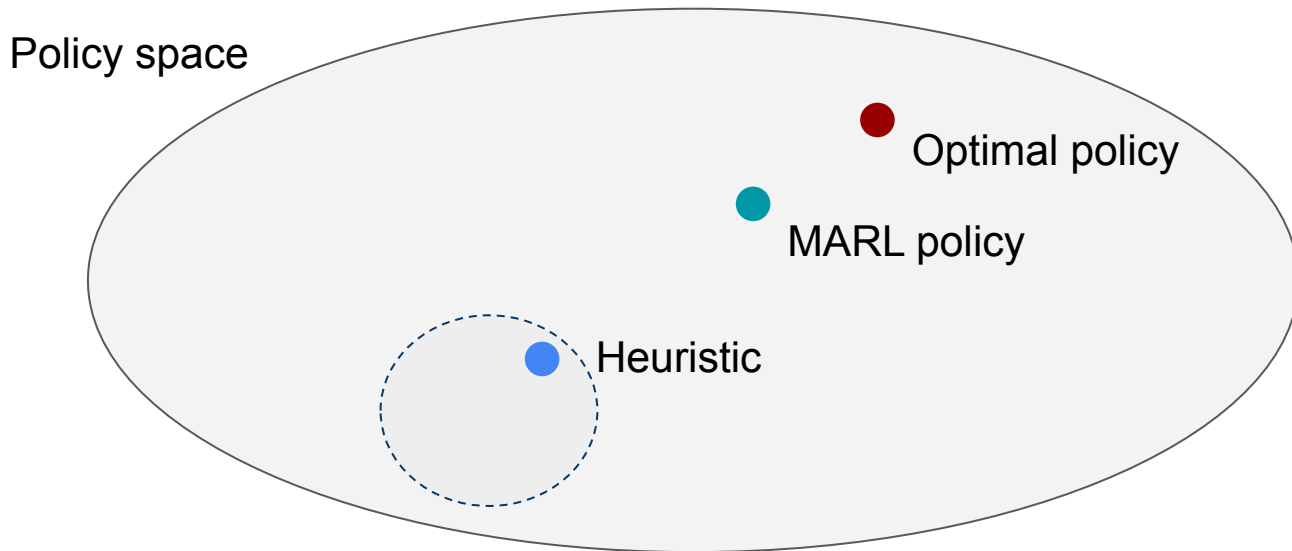
*Performance evaluated with respect to the heuristic



Benchmark results

MARL vs heuristic:


- **CTDE** methods generally **outperform heuristics**



Benchmark results

Scalability ?

- **Centralised RL** methods **do not scale well** with the number of agents
- IMP environments demand **cooperation** among agents: CTDE >> decentralised
- Remaining challenges
 - Correlated environments
 - Group campaign costs



Future work

Ultimate goal:

- **Scalable MARL** applied to real-world problems

What we did:

- Py**MARL** benchmark

What we have:

- **Compatibility** with CleanRL, TorchRL, BenchMARL, Epy Marl,...

What we need:

- Use **IMP-MARL** in your study
- **New IMP** environments
- More **challenges**
- Contribute to the repository!

