Dynamic Noises of Multi-Agent Environments Can Improve Generalization: Agent-based Models meets Reinforcement Learning

Mohamed Akrout, Amal Feriani, Bob McLeod

Introduction

Observations:

- Several real-world dynamics cannot be expressed using mathematical models
- RL agents are prone to exploiting idiosyncrasies of specific implementations of RL environments from which they learn infeasible behaviors in the real world
- Artificial noise injection is proposed to improve dynamics generalization in vision-based problems

Contributions:

- Advocate the use of agent-based models (ABM) to design RL environments
- Examine the effect of the intrinsic dynamic randomization of ABMs on the generalization

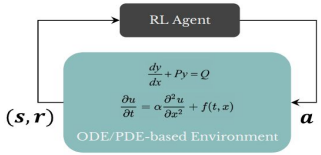

Average behavior

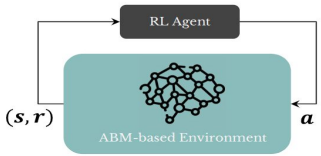

Microscopic behavior

Use Case: Epidemic Control

SIR Model:

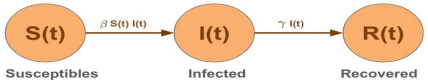

The finite state machine representation of the SIR model.

The dynamics of the SIR mode is governed by :

- β : the effective contact rate of the disease
- γ : the mean recovery rate

ABM-based SIR dynamics

At each time step:

- If an individual is infected, it transmit the disease with a probability β/N to its susceptible neighbors
- Each infected agent transmits to the recovery state after 1/γ days

Intrinsic Noise in ABM-based Dynamics

- The non-deterministic dynamics of ABMs yield an intrinsic noise in the SIR dynamics

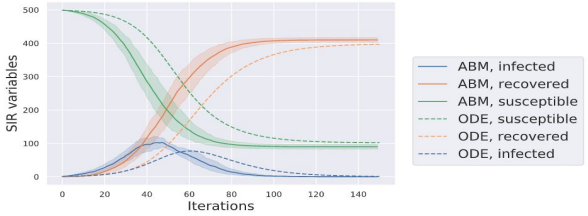

Results

Average Reward Improvement

	ODE	Randomized ODE	ABM
$\beta = 0.2$	0.77 ±0.12	-0.16 ± 0.15	**0.84 ± 0.14**
$\beta = 0.8$	0.78 ±0.8	-0.15 ± 0.13	**0.85 ± 0.11**

Better Generalization

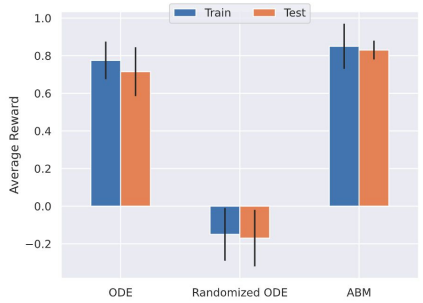

Takeaways

- ABM-based modelling improves the performance and the generalization of RL agents
- Artificial noise injection is harmful and deteriorates the agent performance
- Further experiments on different tasks should be conducted to confirm the paper's findings.