PADERBORN UNIVERSITY

Department of Electrical Engineering and Information Technology
Automatic Control Group
https://en.ei.uni-paderborn.de/de/rat/

Master Thesis

Hierarchical Optimization and Deep Multi-Agent Reinforcement Learning for Coverage Control with unknown Points of Interest

Background

Coverage control is a fundamental problem in multi-robot systems, where robots are equipped with potentially qualitatively different sensing units to cover a given area of interest over time to detect events, collect data, and make predictions. Notably, coverage control is essential for map building and area surveillance. Deep Multi-Agent Reinforcement Learning (DeepMARL) algorithms are powerful tools that can solve coverage control problems without knowledge about the environment and robot dynamics [2]. However, DeepMARL algorithms often scale poorly to many agents/robots, requiring long episodic training on a simulator [1]. For coverage control, this is even more severe when points of interest (POIs) or targets (see Figure 2) are apriori unknown and potentially time-varying, such that training is inherently based on incomplete information.

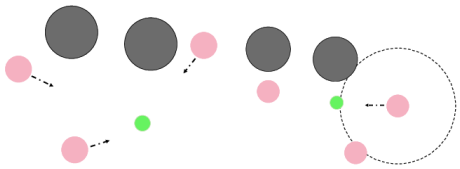

Figure 1: Particle environment as used by OpenAI in [1].

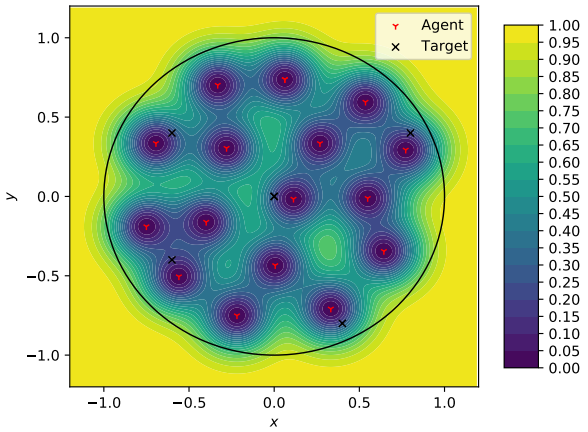

Figure 2: Illustrated coverage control setting. Contour colors visualize the average error probability of sensor reading at every point (x,y).

Thesis Goals

The goal is to tackle the joint POI detection and coverage control problem with a **hierarchical optimization and DeepMARL framework**. Distributed stochastic gradient methods are particularly suitable to propose potential POI and points with potentially good general coverage while guaranteeing fast convergence. A DeepMARL algorithm can, therefore, focus on the underlying control problem under unknown dynamics guided toward points proposed by the optimization iteration.

Intermediate Goals

- Literature review of DeepMARL algorithms and coverage control approaches.
- Familiarise with OpenAI gym multi-particle simulator [1] and extend it with an environment for sensor coverage control with unknown time-varying POIs.
- Implement a distributed stochastic gradient descent (DSGD) approach to propose target points.
- Implement a hierarchical framework where a DeepMARL algorithm should cover target points proposed by DSGD while satisfying constraints (e.g., obstacle avoidance).
- Test the proposed framework in conjunction with state-of-the-art DeepMARL algorithms.

Knowledge

Required:

- Good programming skills (Python).
- Machine learning (SGD), System Theory, and basic RL.

Not required, but a big +:

- Knowledge in Deep Reinforcement Learning.

References

[1] R. Lowe, Y. WU, A. Tamar, J. Harb, O. Pieter Abbeel, and I. Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

[2] S. Meng and Z. Kan. Deep reinforcement learning-based effective coverage control with connectivity constraints. *IEEE Control Systems Letters*, 6:283–288, 2021.

Contact: Adrian Redder, aredder@mail.upb.de, Room: P1.7.09.2