





Project Type _____

- Master Thesis
- Bachelor Thesis
- Research Project

Supervisors _____

-  Niklas Freymuth
-  niklas.freymuth@kit.edu
-  Gerhard Neumann
-  gerhard.neumann@kit.edu

Difficulty _____

Algorithmic



Math



Application



Spatial Credit Assignment for Swarm Reinforcement Learning

Description

Swarm systems are groups of actors that act in a collaborative fashion. In nature, such systems appear in the form of bee swarms, ant colonies and migrating birds. In all these cases, the individual actors perform simple actions, but the swarm as a whole often exhibits surprisingly complex behavior.

The field of Swarm Reinforcement Learning (RL) [2] attempts to replicate this phenomenon; individually acting agents collaborate to maximize a shared, often global reward that would not be obtainable for any single agent. An example can be seen in Figure 1, where a number of actors (blue) chase and eventually catch an evader (red) that is much faster than any individual agent.




Figure 1: A swarm of agents (blue) collaborating to catch an evader (red).

The communication between agents is naturally modelled as a message-passing problem on graphs, where the nodes are the agents that communicate along the edges. As such, we realize the policy using Graph Neural Networks [1], which are neural networks that act on graphs as inputs.

The problem that arises from this setup is that the agents share a single (graph-wise) global reward that disregards the performance of the individual agents. As a result, the correlation between per-agent performance and reward becomes extremely noisy, and not all agents learn to perform the task equally well. In the above example, the agent on the left of the evader at $t = 121$ lags behind the other agents but is still rewarded due to a good global solution.

The goal of this project will be to develop an algorithm that allows for an individual reward for each agent. Limiting ourselves to geometric graphs, this can be realized via *spatial credit assignment*, i.e., by decomposing the reward based on the locality of the agents relative to each other and based on the objective of the task.

The result will be a framework that allows multiple agents to learn complex joint behavior while making sure that every individual agent is contributing to the underlying task in a meaningful way.

Tasks

- Literature Review: Get familiar with Graph Neural Networks, (Swarm) RL and the theory behind it
- Algorithm Design: Develop an algorithm for spatial credit assignment for swarm systems, assigning individual rewards to agents based on how well they contribute to the solution of a shared geometric task.
- Evaluation: Evaluate your algorithm on simulated swarm systems and compare it to existing approaches.

References

- [1] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- [2] Maximilian Hüttenrauch, Sosic Adrian, Gerhard Neumann, et al. Deep reinforcement learning for swarm systems. *Journal of Machine Learning Research*, 20(54):1–31, 2019.