

Hands-On Exercises in Multi-Agent Reinforcement Learning (MARL)

Introduction

These exercises are designed to help you explore key concepts in Multi-Agent Reinforcement Learning (MARL). You will work with a simulated environment where multiple agents interact, observe, and learn in both cooperative and competitive settings. Through these exercises, you will analyze the effects of rewards, policies, and interactions among agents.

Start by downloading the file `grid_world.py`, which contains the environment setup and basic functionalities.

Exercise 1: Independent Agents

Create agents that can move in a grid-world environment and interact with each other. Each agent should have a simple policy for movement and interaction. Its goal is to reach a target cell while avoiding obstacles.

Exercise 2: Cooperative Task - Reaching a Shared Goal

Objective: Implement a cooperative task where both agents need to reach a common goal.

Instructions:

1. Modify the environment by placing a single goal cell that both agents must reach simultaneously to receive a reward.
2. Adjust the reward structure:
 - +5 reward if both agents reach the goal at the same time.
 - -0.1 penalty for each movement.
 - -1 penalty for any collision.
3. Implement a coordination policy, allowing agents to consider each other's actions (e.g., communicate intended moves).

Exercise 3: Competitive Task - Competing for Resources

Objective: Introduce competition by placing a limited resource in the environment that both agents attempt to capture.

Instructions:

1. Place a single resource cell in the environment, rewarding only the agent that reaches it first.
2. Update the reward structure:
 - +3 reward for capturing the resource.
 - -0.1 penalty for each movement.
 - -1 penalty for collisions.
3. Allow agents to "observe" the other agent's position and adjust their strategies based on the competing presence.