Introduction to Multi-Agent Deep Deterministic Policy Gradient

Authors : Xu Jie

Abstract : Multi-Agent Reinforcement Learning (MARL) is an increasingly important area in artificial intelligence, where multiple agents learn to make decisions and interact within a shared environment. One of the key challenges in MARL is the non-stationary dynamics that emerge from interactions between multiple agents, which can complicate the learning process. Multi-Agent Deep Deterministic Policy Gradient (MADDPG) is a prominent method that addresses this issue by introducing centralized training with decentralized execution. This paper provides an overview of MADDPG, highlighting its architecture, advantages, and its application in various multi-agent environments.

Keywords : multi-agent reinforcement learning (MARL), non-stationary dynamics, multi-agent systems, cooperative and competitive agents

Conference Title : ICSLP 2024 : International Conference on Speech and Language Processing **Conference Location :** San Francisco, United States

Conference Dates : November 04-05, 2024