

Multi-agent Continual Coordination via Progressive Task Contextualization

Lei Yuan^{1,2}, Lihe Li¹, Ziqian Zhang¹, Fuxiang Zhang^{1,2}, Cong Guan¹, Yang Yu^{1,2,*}

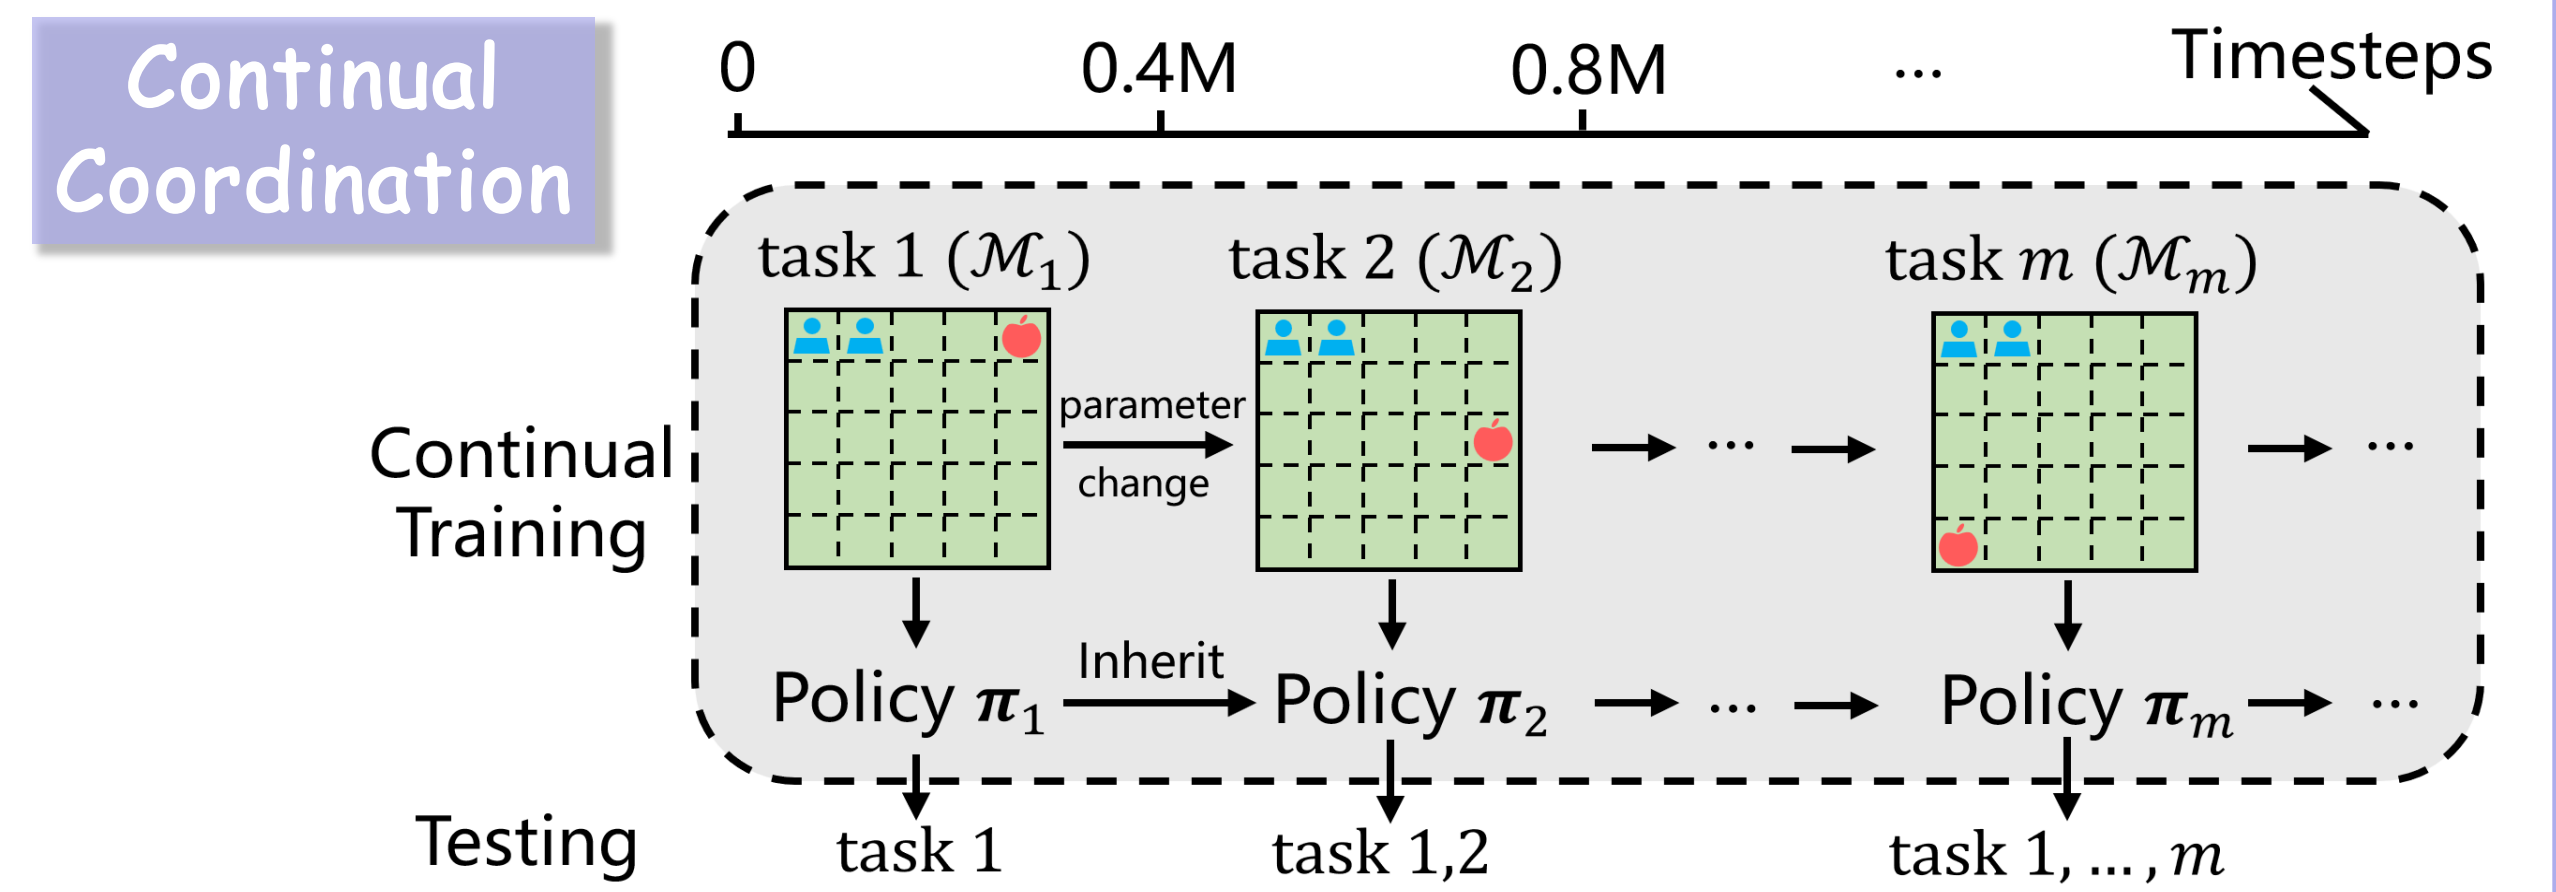
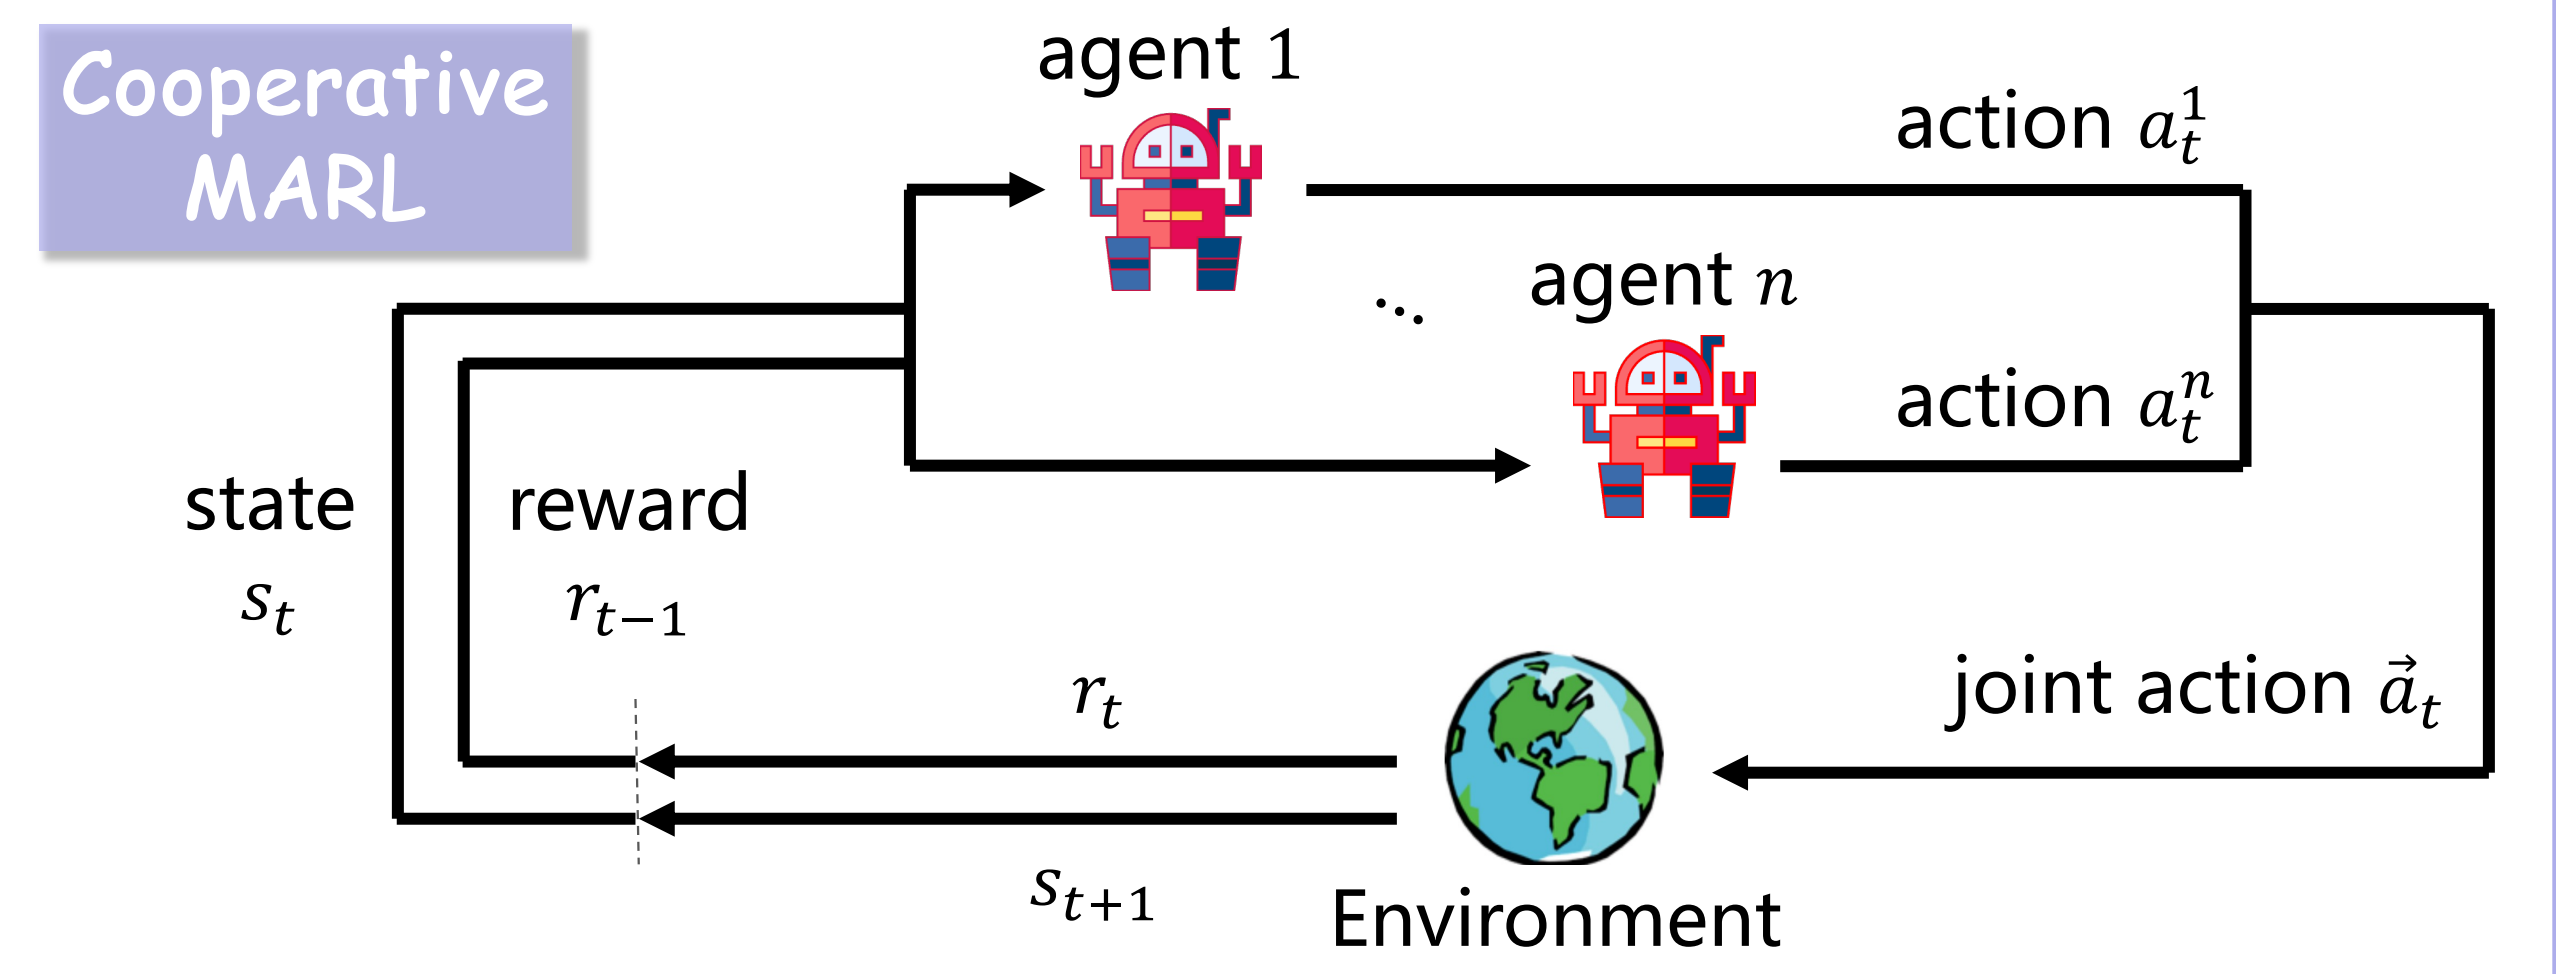
¹National Key Laboratory for Novel Software Technology, Nanjing University, ²Polixir.ai

Introduction

- Many real-world problems consist of multiple interactive agents, and can be modeled as **multi-agent systems (MAS)**.
- We can apply **reinforcement learning (RL)**, a machine learning paradigm that involves agents learning to make decisions, to train multi-agent systems to solve different tasks.
- If all agents in the multi-agent system share the same reward, this learning paradigm can be formulated as **cooperative multi-agent reinforcement learning (CMARL)**.

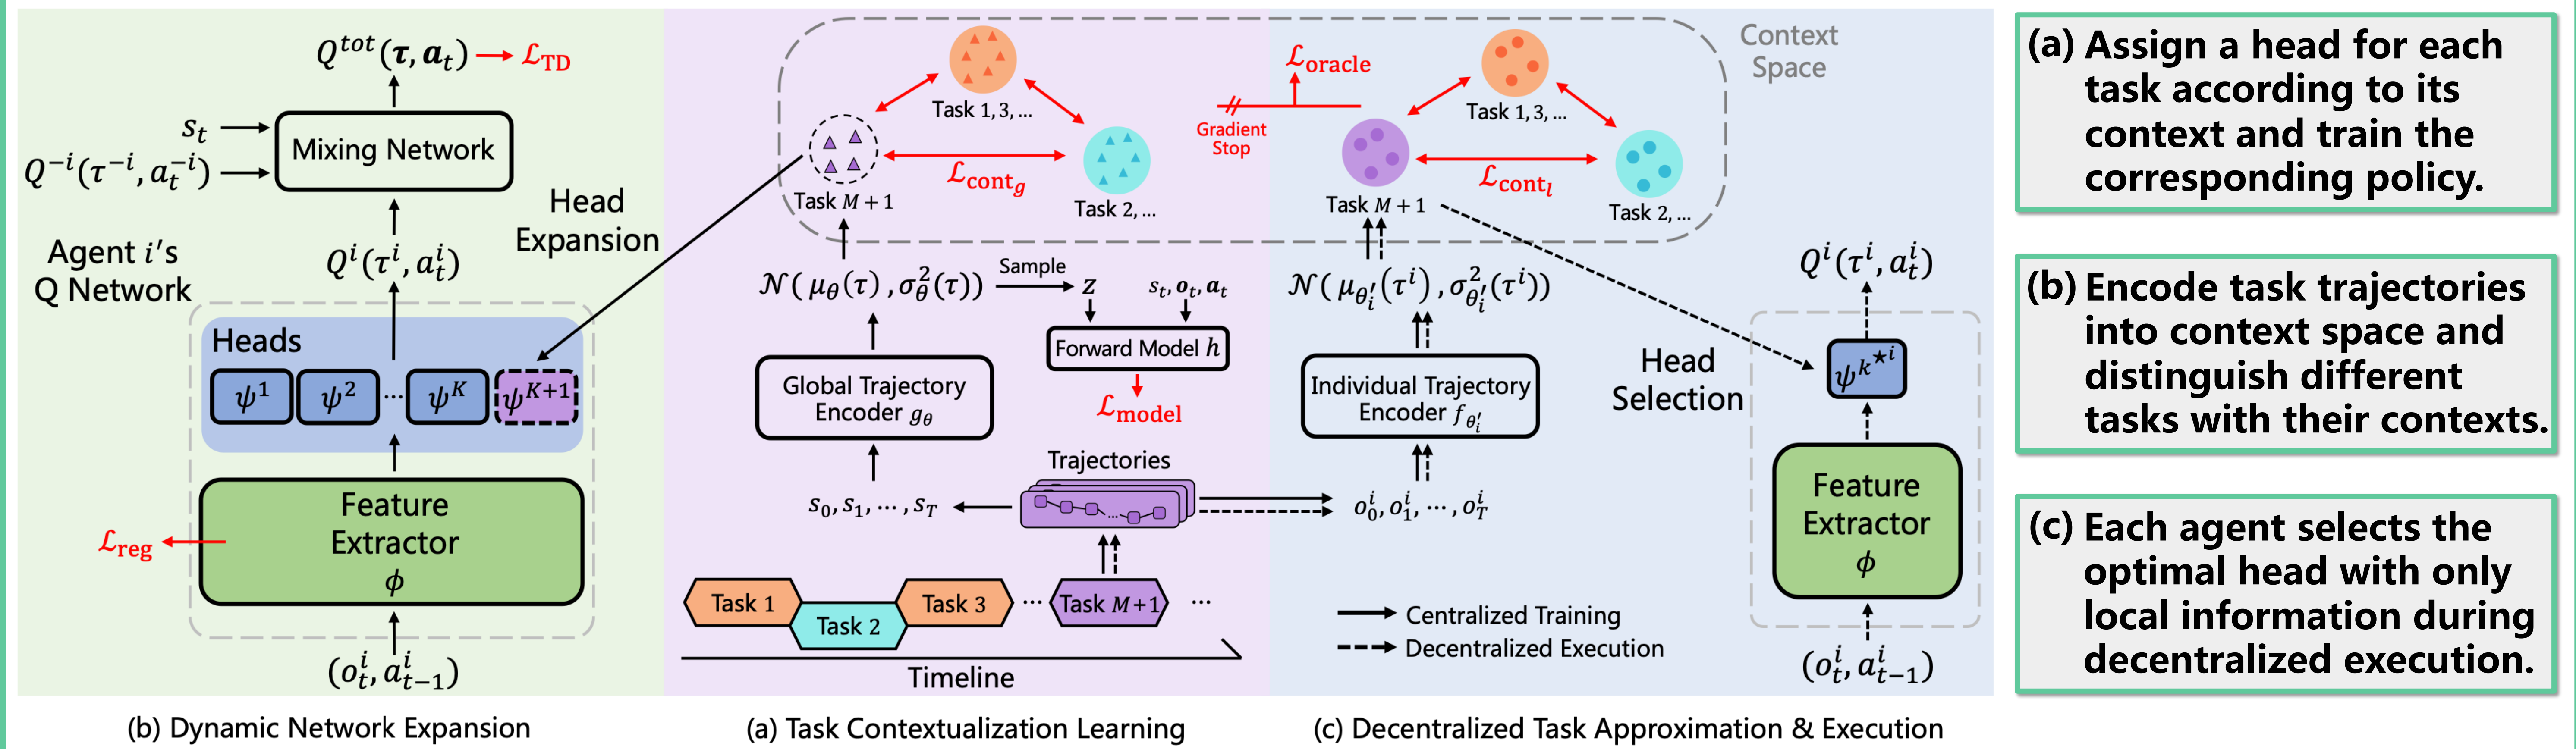
However,

- Mainstream CMARL methods are still restricted to being trained in one single task or multiple tasks simultaneously and thus **cannot learn a stream of tasks** in a continual manner.
- To solve this issue, we formalize **continual coordination** and learn a MAS in a stream of tasks via **progressive task contextualization**.



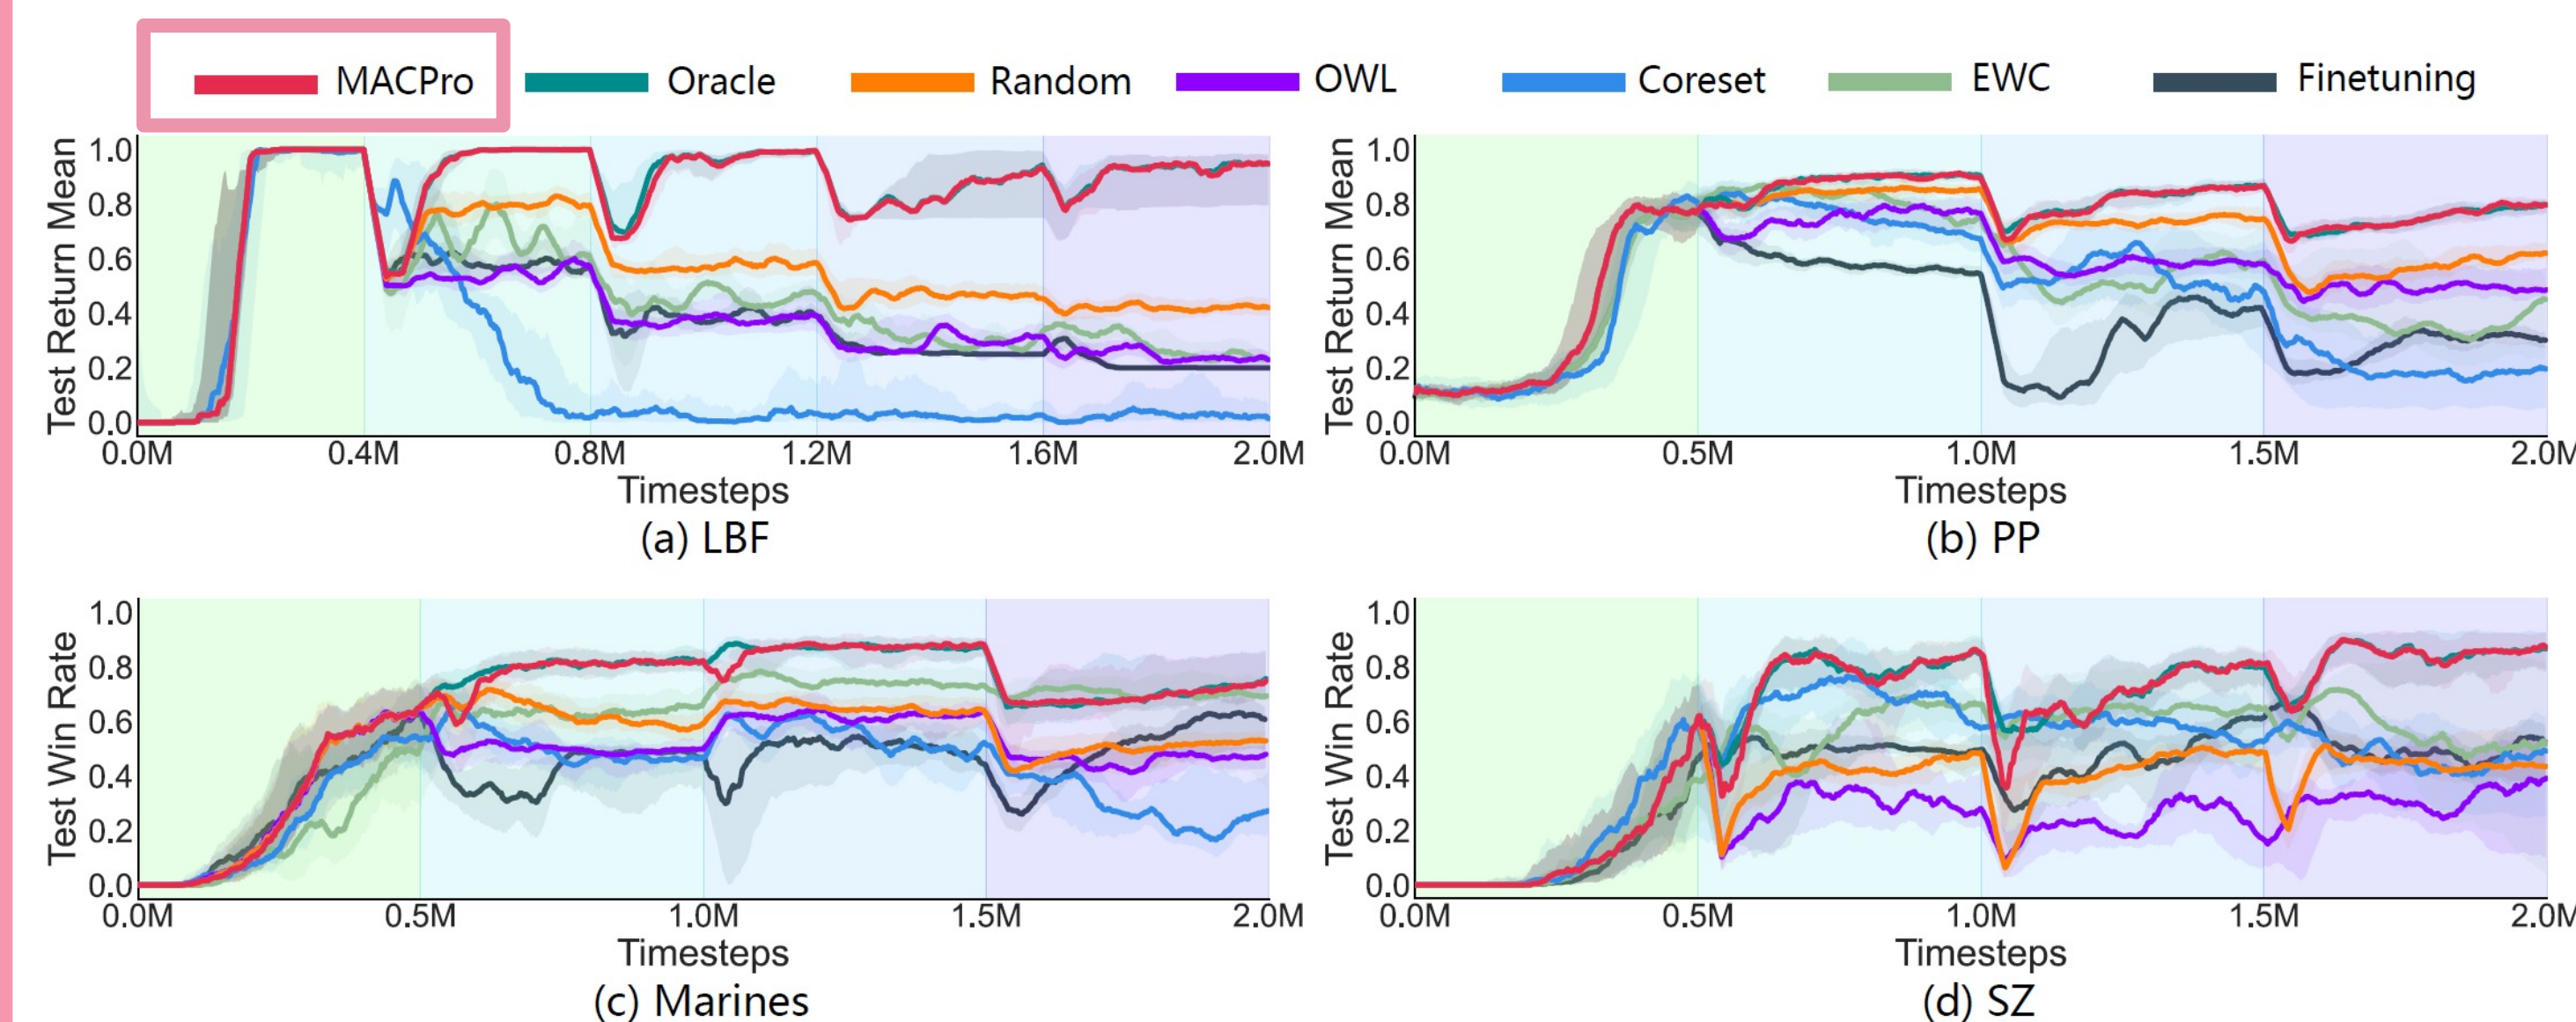
Method

- The overall workflow of our method MACPro



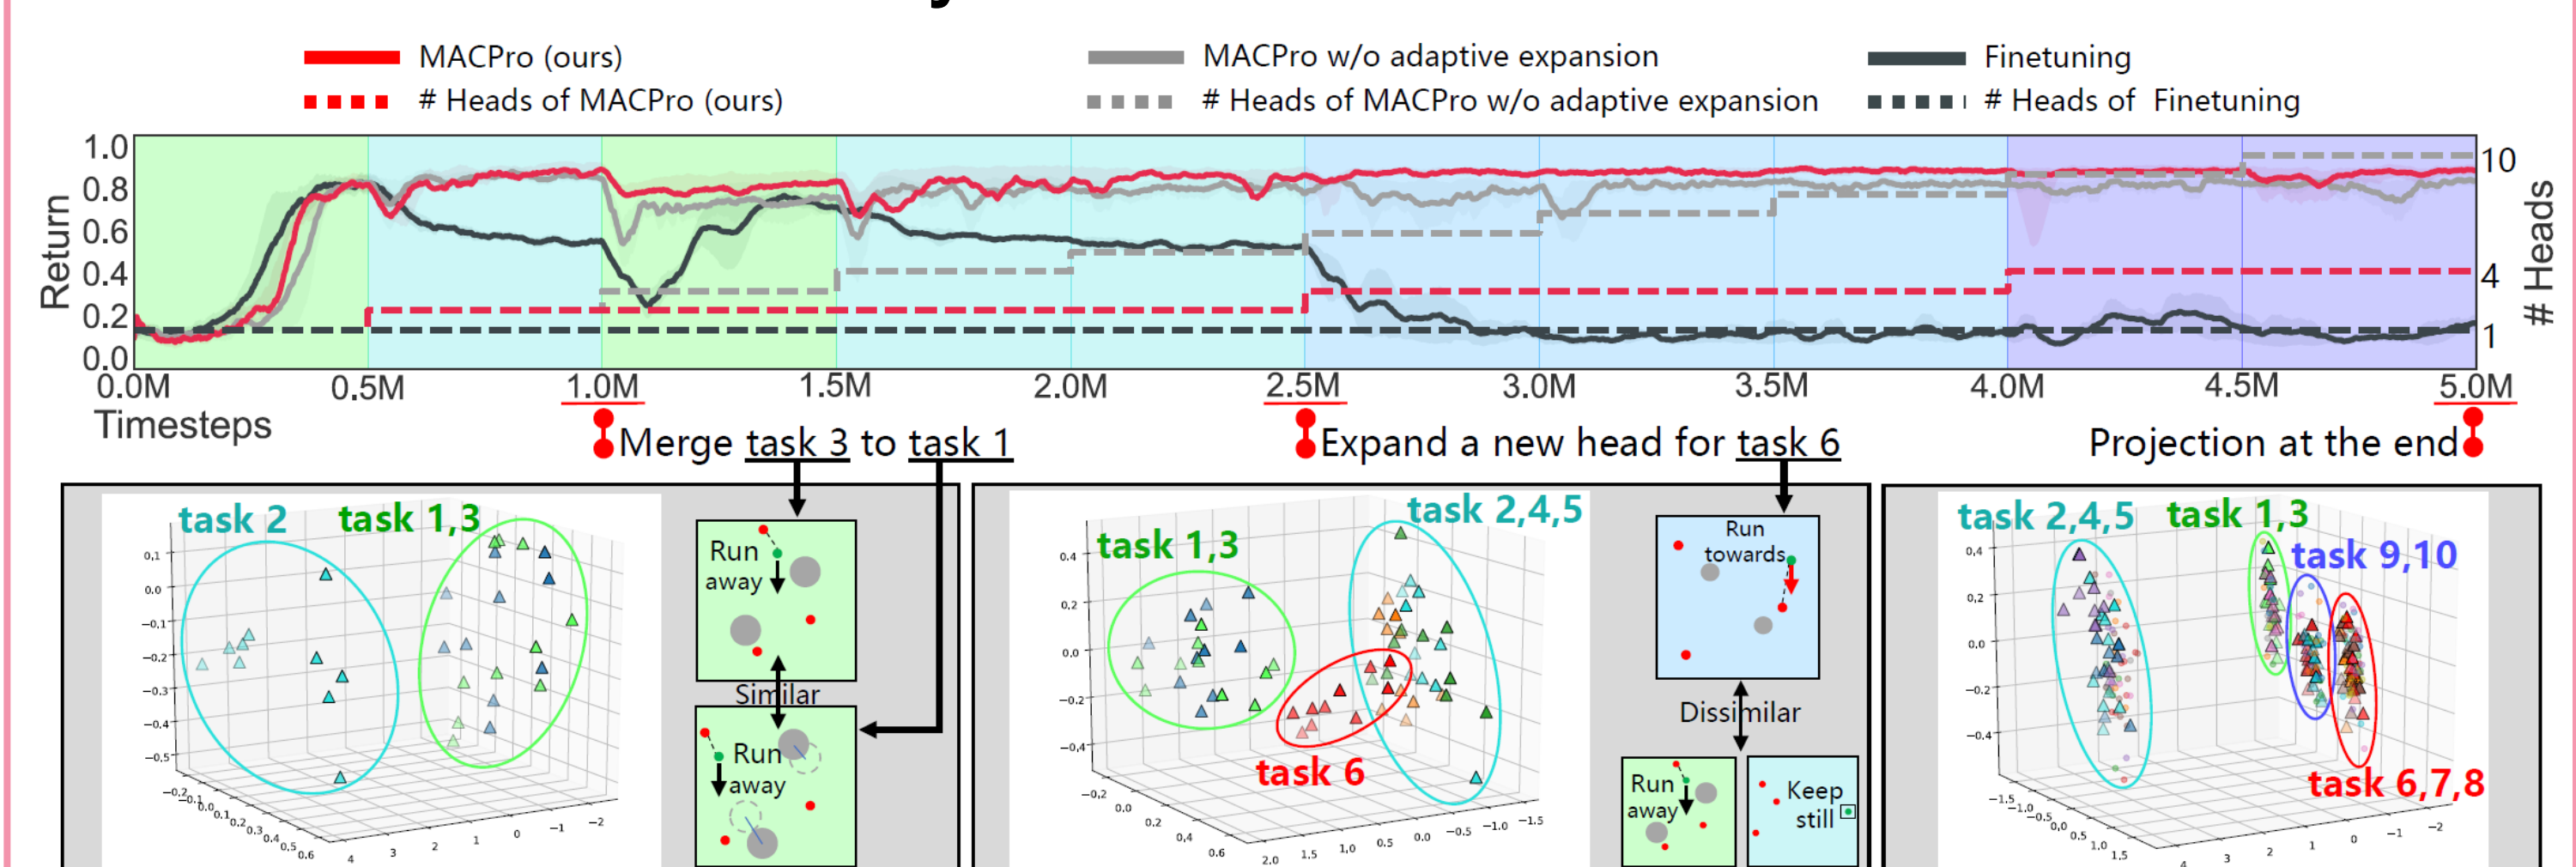
Experiments

- Performance on different benchmarks



- MACPro achieves the **best** continual learning performance on all multi-agent benchmarks.
- Swift knowledge transfer to new tasks (**plasticity**) while avoiding catastrophic forgetting (**stability**).

- Visualization Analysis



- Expand a new head only for **dissimilar** tasks, without sacrificing performance.
- Enjoy high **scalability** and learning **efficiency**.