SYMPHONY

Synergistic Multi-agent Planning with Heterogeneous Language Model Assembly

Wei Zhu, Zhiwen Tang*, Kun Yue

School of Information Science and Engineering, Yunnan University

Yunnan Key Laboratory of Intelligent System and Computing

zhuwei@stu.ynu.edu.cn, {zhiwen.tang, kyue}@ynu.edu.cn

NeurIPS 2025

The Problem with Single-Agent Planning

Existing Approach

Most LLM-based planners use a **single agent** (one LLM) for MCTS.

- They rely on querying the **same model** repeatedly to simulate search branches.
- This assumes model stochasticity is enough for diverse exploration.

The Limitation

This single-agent paradigm leads to **low reasoning diversity**.

- Outputs often reflect the same "dominant reasoning pattern."
- This results in narrow search, suboptimal plans, and susceptibility to local optima.

Our Solution: SYMPHONY

The Key Idea

Move from a single agent to a **heterogeneous multi-agent pool**.

- Leverage the diverse reasoning patterns from different LLMs (e.g., Qwen, Llama, GPT).
- SYMPHONY: **S**ynergistic **M**ulti-agent **P**lanning with **H**eterogene**O**us la**N**guage model assembl**Y**.

How It Works

Different agents are assigned to generate actions at search nodes.

- This introduces **structural diversity** directly into the MCTS tree.
- It increases exploration, reduces model-specific bias, and finds complementary solution paths.

Framework Overview

Core Components

- **Heterogeneous Agent Pool:**
- **Agent Scheduling (UCB):**
- **Pool-wise Memory Sharing:**
- **Entropy-Modulated Node Evaluation(EMCS):**

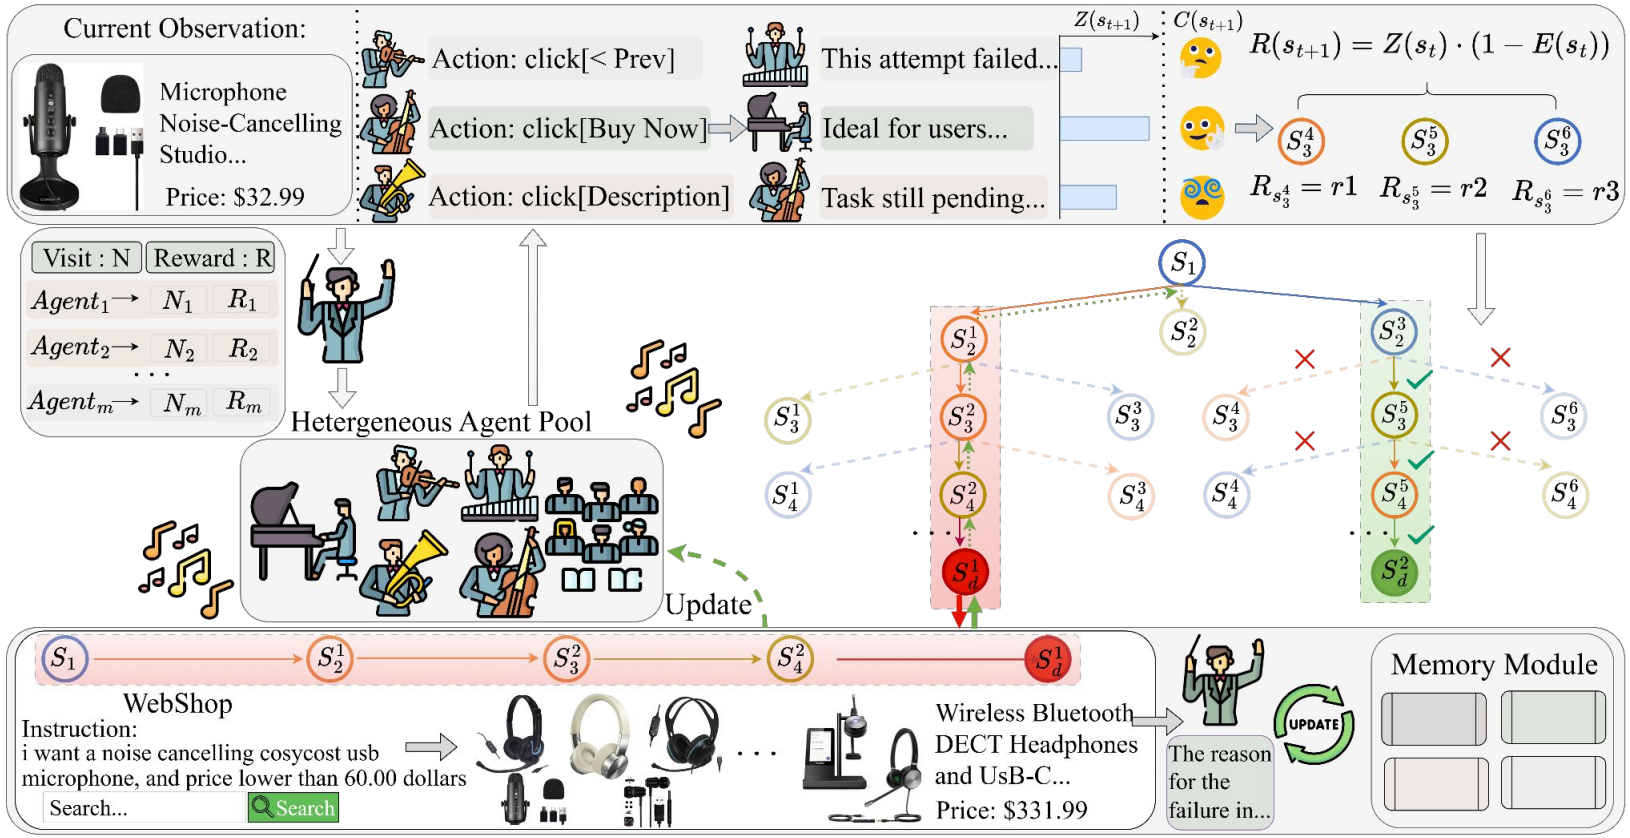

Experimental Results

SYMPHONY consistently outperforms SOTA baselines across diverse tasks.

Method	HotpotQA (EM) ↑	WebShop (SR) ↑	MBPP (Pass@1) ↑
ReAct	0.39	0.32	0.710
LATS	0.71	0.38	0.811
MASTER	0.76	-	0.910
SYMPHONY-S (Local)	0.59	0.56	0.927
SYMPHONY-L (API)	**0.79**	**0.72**	**0.965**

SYMPHONY-L (with API models) achieves new SOTA.

SYMPHONY-S (with local models) is highly competitive and efficient.

Analysis: Diversity is Key

Why Does It Work?

We analyzed **branch diversity** versus the agent pool composition on all three tasks.

- **Finding:** Increasing agent heterogeneity (from a single agent to the full trio) directly increases the proportion of "4-Unique" branches.
- This structural diversity strongly correlates with task performance.
- **Conclusion:** More reasoning diversity leads to better exploration and more robust planning.

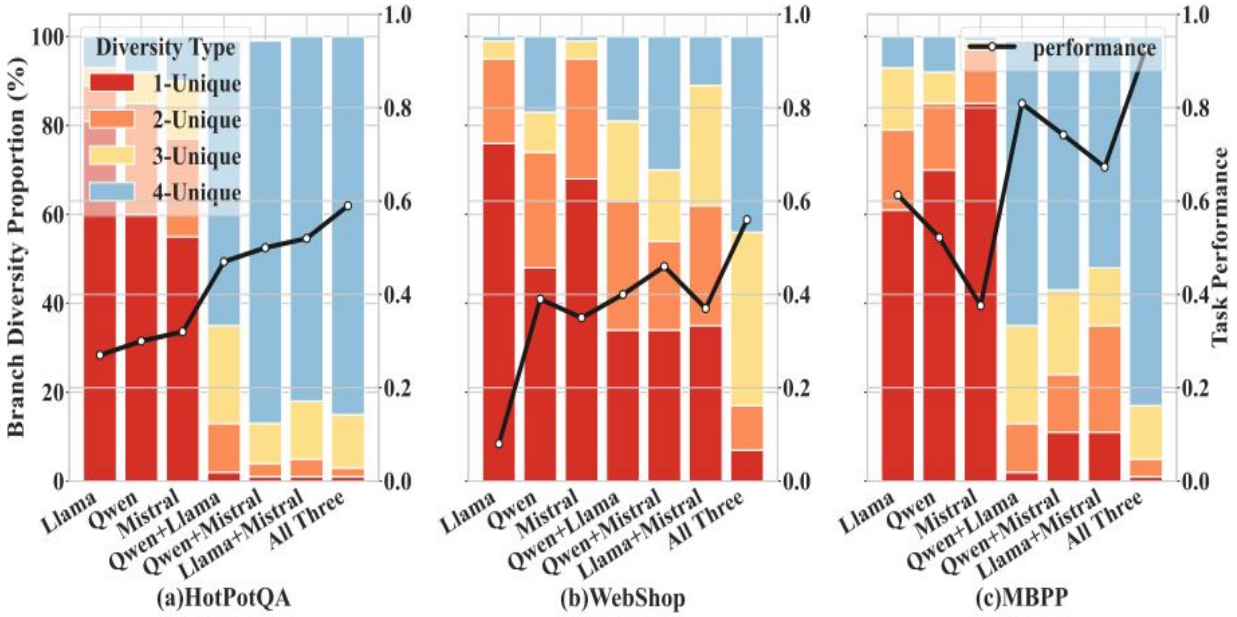

Efficiency and Cost-Effectiveness

Search Efficiency (HotpotQA)

SYMPHONY requires **far fewer node expansions** to find the solution.

- LATS (K=50): 66.65 nodes
- **SYMPHONY-L (K=10): 9.47 nodes**

Method	K	HotpotQA ↑	#Nodes ↓
ToT	10	0.34	33.97
RAP	10	0.44	31.53
LATS	10	0.44	28.42
ToT	50	0.49	84.05
RAP	50	0.54	70.60
LATS	50	0.61	66.65
SYMPHONY-S	**10**	**0.59**	**16.39**
SYMPHONY-L	**10**	**0.79**	**9.47**

Cost-Effectiveness (HotpotQA)

SYMPHONY-L intelligently schedules agents, reducing calls to expensive models.

- **GPT-4 used in only ~35% of calls**, vs. 100% for baselines.
- Achieves SOTA performance at a fraction of the API cost.

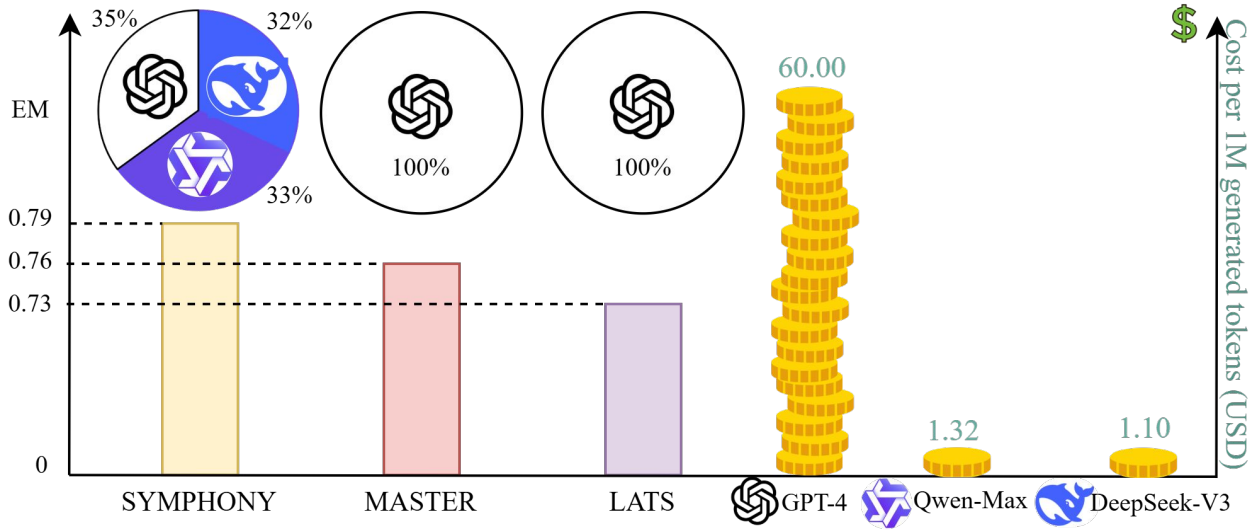

Conclusion

- We introduced **SYMPHONY**: The first heterogeneous multi-agent planning framework for MCTS.
- Core mechanisms (UCB, Pool-wise Memory, EMCS) establish structural diversity and collective learning.
- SOTA results on HotpotQA, WebShop, and MBPP with high efficiency.

Code:

https://github.com/ZHUWEI-hub/SYMPHONY

Thank You