LLaMAR: Long-Horizon Planning for Multi-Agent Robots in Partially Observable Environments

Siddharth Nayak* [1] Adelmo Orozco* [1] Marina Ten Have [1] Vittal Thirumalai [1] Jackson Zhang [1] Darren Chen [1] Aditya Kapoor [2]
Eric Robinson [3] Karthik Gopalakrishnan [4] James Harrison [5] Brian Ichter [5] Anuj Mahajan [6] Hamsa Balakrishnan [1]

[1]Massachusetts Institute of Technology [2]TCS RnI [3]USAF-MIT AI Accelerator [4]Stanford University [5]Google DeepMind [6]Apple

Planning with LLMs

Current Approaches to Planning with LLMs

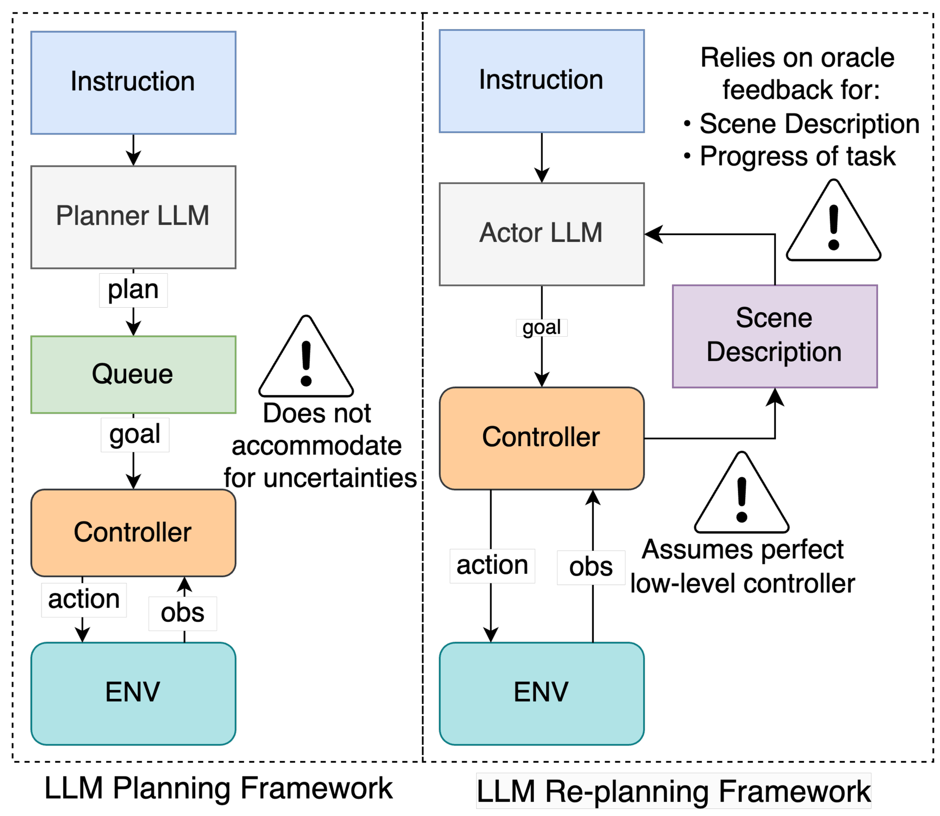

	Two Step	Smart LLM	S-ATLAS	CoELA	LLaMAR
Dynamic Planning	✗	✗	✓	✓	✓
Local Information	✗	✗	✗	✓	✓
Failure Correction	✗	✗	✗	✗	✓
Self Verification	✗	✗	✗	✗	✓

Table 1. Comparison of planning methods and their capabilities.

Need a method that can accommodate uncertainties, and does not rely on perfect low-level control and oracle feedback.

Results

LLaMAR vs Baseline Performance

Algorithm	LM	SR↑	TR↑	C↑	B↑
Act	GPT-4V	0.33	0.67	0.91	0.59
ReAct	GPT-4V	0.34	0.72	0.92	0.67
CoT	GPT-4V	0.14	0.59	0.87	0.62
SmartLLM	GPT-4V	0.11	0.23	0.91	0.45
CoELA	GPT-4V	0.25	0.46	0.76	0.73
LLaMAR	GPT-4	0.51	0.85	0.95	0.83
LLaMAR	LLaVA	0.54	0.84	0.91	0.75
LLaMAR	IDEFICS-2	0.57	0.86	0.94	0.78
LLaMAR	CogVLM	0.61	0.89	0.95	0.80
LLaMAR (w/o expl)	GPT-4V	0.62	0.87	0.95	0.82
LLaMAR (w/ expl)	GPT-4V	**0.66**	**0.91**	**0.97**	**0.82**

Table 2. Comparison of LLaMAR against baselines averaged across all tasks.

Modules Used	SR↑	TR↑	C↑	B↑
Actor	0.33	0.67	0.91	0.59
Planner + Actor + Verifier	0.45	0.78	0.92	0.69
Planner + Actor + Corrector‡	**0.67**	**0.91**	**0.97**	**0.84**
Planner + Actor + Corrector + Verifier +	0.66	**0.91**	**0.97**	0.82

Table 3. Performance in the 2-agent scenarios in MAP-THOR obtained by ablating different modules in LLaMAR

Partitioning the planning process into distinct modules enhances performance in MAP-THOR

LLaMAR is able to perform as well as a model with an oracle-based verifier

LLaMAR Architecture

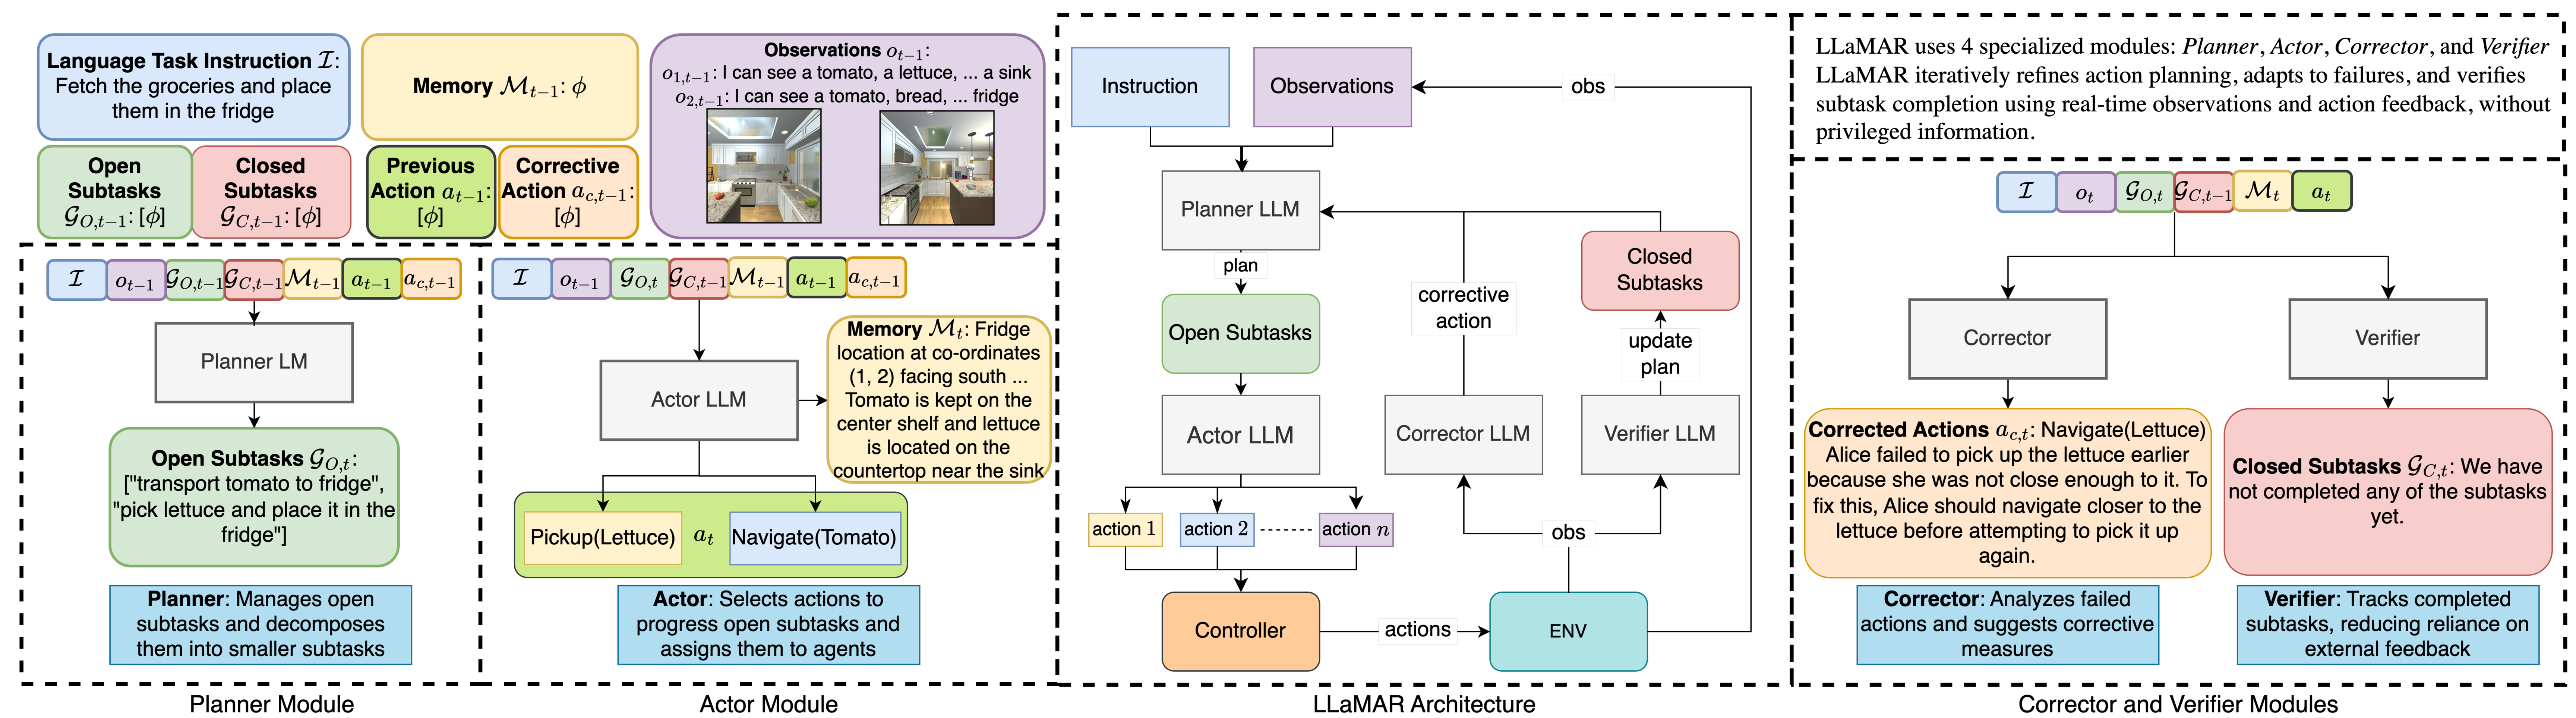

Task Environments

1. MAP-THOR: A test-suite and benchmark on language-based multi-agent robotic planning based on AI2THOR.
2. Search & Rescue (SAR): Agents are tasked with extinguishing fires before they spread and rescuing missing humans.
3. LLaMAR: Creates performant long-horizon planning in multi-agent tasks by generating subtasks and assigning them to different agents.

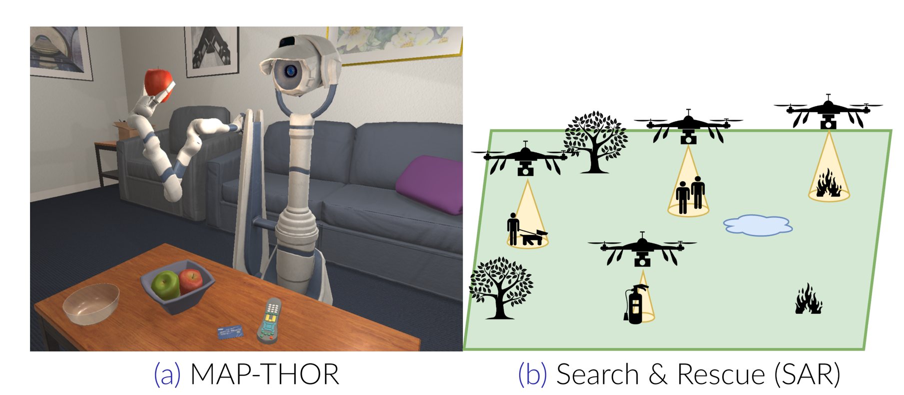

(a) MAP-THOR (b) Search & Rescue (SAR)

We perform experiments in two different multi-agent robotics environments.

Conclusions

- LLaMAR creates long-horizon plans in multi-agent tasks by creating and assigning subtasks to different agents.
- LLaMAR enables performant multi-agent planning in uncertain environments.
- We introduce MAP-THOR a benchmark dataset consisting of language-conditioned multi-agent robotic tasks with various difficulty levels.

Results

# of agents	MAP-THOR SR↑	MAP-THOR TR↑	MAP-THOR C↑	SAR SR↑	SAR TR↑	SAR C↑
1	0.37	0.67	0.87	0.28	0.75	0.86
2	0.62	0.87	0.95	0.44	0.86	0.94
3	**0.70**	**0.91**	0.98	0.68	0.92	0.96
4	0.68	0.90	**0.99**	0.72	0.94	0.98
5	0.62	0.90	**0.99**	**0.74**	**0.96**	**1.00**

Table 4. LLaMAR with various number of agents in the scenario in both MAP-THOR and SAR environments

Increasing number of agents improves performance in SAR

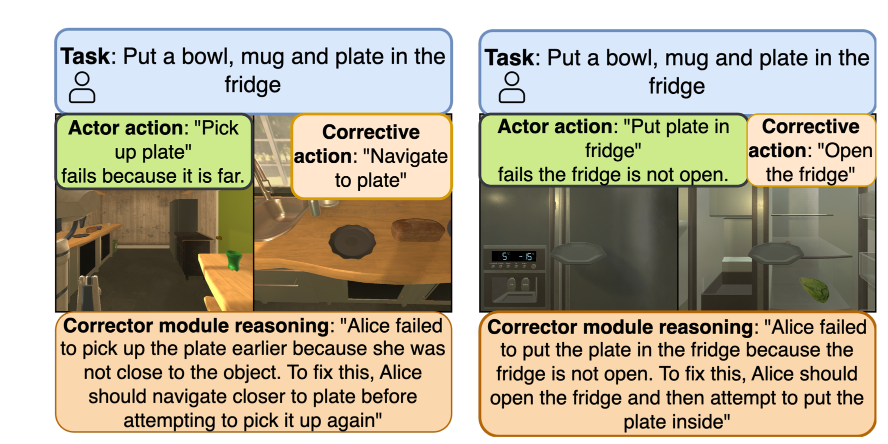

Figure 1. Failure Correction with visual feedback

Corrector helps with reasoning on failures at low-level execution

Limitations & Assumptions

- Low-level controllers can perform language-conditioned actions. e.g., `Pickup(Bowl)`
- When an agent observes an object, its location is inherently stored in the navigation module's memory.
- Performance is limited by the underlying VLM and hence has limited spatial reasoning
- LLaMAR incurs a higher computational cost at each step due to the involvement of multiple language models during each iteration.
- The experiments are conducted in a simulated environment rather than in the real world.

This research was sponsored by the U.S. AFRL/U.S. Air Force Artificial Intelligence Accelerator (Cooperative Agreement Number FA8750-19-2-1000).
The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies of the Department of the Air Force or the U.S. Government.