

Svyatoslav Pchelintsev, Maxim Patratskiy, Anatoly Onishchenko Alexandr Korchemnyi, Aleksandr Medvedev, Uliana Vinogradova, Ilya Galuzinskiy, Aleksey Postnikov, Alexey K. Kovalev, and Aleksandr I. Panov

Motivation

LLM-based planners often fail in dynamic environments because they rely solely on textual representations and cannot perceive real-world changes or execution errors. To overcome this limitation, we propose a vision-driven replanning approach that enables robots to interpret visual feedback and recover from failures without requiring object detection or predefined conditions.

Method

The LERa module performs a sequence of Visual Language Model queries — **Look**, **Explain**, and **Replan** — to identify execution errors and adjust the current task plan accordingly. LERa is **not** a standalone task planner; instead, it acts as a **corrective layer** that refines and updates existing plans when failures occur. This modular design allows LERa to be seamlessly integrated into any existing agent and planning system, enhancing their robustness through vision-based feedback.

Experiments

Agent	ALFRED-ChaOS (Seen)			ALFRED-ChaOS (Unseen)		
	SR↑	GCR↑	SRep↑	SR↑	GCR↑	SRep↑
Oracle	33.04	50.04	-	31.65	51.71	-
O-Ra	34.38	51.19	7.69	34.17	54.08	14.16
O-ERa	40.18	56.40	37.08	42.81	61.81	33.33
O-LRa	34.38	51.00	6.73	33.45	53.57	11.01
Baseline	33.04	50.15	3.15	32.01	51.89	0.98
O-LERa	49.55	64.55	73.39	53.60	70.23	74.57


Robot Experiments



1 Move **apple** and **orange** into container



2 Pick apple



3 Fail



4 Replan: Pick apple



5 Fail



6 Replan: **apple** finally recognised as **pear**



7 Pick lemon



8 Place lemon into container

LERa Overview

