

RoboGPT-R1: Enhancing Robot Task Planning with Reinforcement Learning

Jinrui Liu*

Institute of Automation, CASIA
Beijing, China
School of Artificial Intelligence, UCAS
Beijing, China
lijinrui2024@ia.ac.cn

Bingyan Nie*

Institute of Automation, CASIA
Beijing, China
School of Artificial Intelligence, UCAS
Beijing, China
niebingyan2025@ia.ac.cn

Boyu Li

Institute of Automation, CASIA
Beijing, China
School of Artificial Intelligence, UCAS
Beijing, China
liboyu2021@ia.ac.cn

Yaran Chen

Institute of Automation, CASIA
Beijing, China
School of Artificial Intelligence, UCAS
Beijing, China
chenyaran2013@ia.ac.cn

Yuze Wang

Huawei Cloud Technology Co., Ltd
Shenzhen, China
wangyuze1@hisilicon.com

Shunsen He

Huawei Cloud Technology Co., Ltd
Shenzhen, China
heshunsen@huawei.com

Haoran Li[†]

Institute of Automation, CASIA
Beijing, China
School of Artificial Intelligence, UCAS
Beijing, China
Beijing Academy of
Artificial Intelligence
Beijing, China
lihaoran2015@ia.ac.cn

ABSTRACT

Improving the reasoning capabilities of embodied agents is crucial for robots to complete complex human instructions in long-view manipulation tasks successfully. Despite the success of large language models and vision language models based on Supervised Fine-Tuning (SFT) in planning tasks, they continue facing challenges in performing long-horizon manipulation tasks in complex real-world environments, owing to their restricted common sense and reasoning capabilities. Considering that aligning general-purpose vision language models to robotic planning tasks via supervised fine-tuning suffers from poor generalization and insufficient physical understanding, we propose RoboGPT-R1, a two-stage fine-tuning framework for embodied planning. In this framework, supervised training acquires foundational knowledge through expert sequences, followed by RL to address the model’s shortcomings in visual-spatial understanding and reasoning. To achieve physical

understanding and action sequence consistency in multi-step reasoning tasks, we design a rule-based reward function that simultaneously considers long-horizon performance and action constraint in the environment. The reasoning model, trained on Qwen2.5-VL-3B, significantly outperforms the larger-scale model, GPT-4o-mini, by 21.33% and surpasses other work trained on Qwen2.5-VL-7B by 20.33% on the EmbodiedBench benchmark.

KEYWORDS

Robot Task Planning; Reasoning Planning; Reinforcement Learning; Vision-Language-Model

ACM Reference Format:

Jinrui Liu, Bingyan Nie, Boyu Li, Yaran Chen, Yuze Wang, Shunsen He, and Haoran Li. 2026. RoboGPT-R1: Enhancing Robot Task Planning with Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 11 pages. <https://doi.org/10.65109/NOXT1107>

1 INTRODUCTION

Recently, vision language models (VLMs) have been increasingly employed as high-level planners for embodied tasks [6, 17, 24], given their emerging capability to ground natural language instructions into long-horizon robotic action sequences. Nevertheless, in real-world environments, VLMs still fail to meet the demands of robustness and generalization [6, 7, 68]. Two major challenges remain. First, the prevailing supervised fine-tuning (SFT)-only paradigm primarily imitates expert demonstrations, yet lacks mechanisms

*Equal contribution.

[†]Corresponding author.

Code Page: <https://github.com/Jinrui-Liu137/RoboGPT-R1>



This work is licensed under a Creative Commons Attribution International 4.0 License.

for adaptation or self-correction in dynamic environments [10, 28]. Second, the design of long-horizon reward functions remains inadequate—existing rewards are often sparse or poorly aligned with the execution of grounded action, ultimately hindering planning performance [54].

In real-world long-horizon embodied tasks, VLMs still exhibit limited planning capability [54], as they are not well aligned with the physical realities of robotic embodiments or with accurate state transition dynamics [6, 61]. Although existing approaches based on the SFT-only paradigm can enhance the performance of VLMs [8, 66], they remain ineffective when confronted with scenarios or instructions that fall outside the distribution of the SFT dataset [59]. The lack of physical common sense and feasibility constraints often results in ambiguous object recognition and biased state estimation [14, 27]. Moreover, the absence of feedback and error-correction signals [58] in the SFT paradigm encourages models to memorize answers rather than learn generalizable reasoning strategies [10], thereby failing to mitigate the accumulation of local errors over extended task horizons.

Reinforcement learning (RL) has proven effective for VLMs in domains such as video reasoning [26], object detection [30, 39], and mathematical reasoning [50, 51, 60], where tasks provide clear and verifiable answers. However, when transferred to open-ended embodied planning tasks, RL-based methods face challenges in designing dense and interpretable reward functions [24, 64, 68], as the outcomes are often ambiguous and context-dependent. For instance, in embodied planning tasks, when the reference plan is "pick up an apple and put it on the table", using a straightforward RL reward such as string matching or accuracy calculation allows the model to gain higher rewards simply by generating more actions, as some subsequences are likely to overlap with the reference plan. This mechanism misleads the model to produce overly long yet logically incorrect reasoning chains, masking its true deficiencies in action ordering and planning coherence. Therefore, a dense and sequence-aware reward is required to directly capture whether a multi-step plan is executed fully or partially correctly [6, 61, 64], particularly in long-horizon and complex action sequences, rather than rewarding superficial token overlap.

To address the above problems, we propose RoboGPT-R1, a two-stage training framework designed to enhance robotic planning with small-scale models. In contrast to the SFT stage, which learns predefined answers, the Group Relative Policy Optimization (GRPO) algorithm explores optimal solutions, addressing the shortcomings of SFT in generalization, task understanding, spatial perception, and planning consistency [10]. Moreover, in second-stage RL training, unlike conventional RL approaches in reasoning tasks that typically rely on sparse or single-point accuracy rewards, our method introduces a rule-based variable reward function specifically designed for long-horizon embodied reasoning and planning. This reward function consists of two complementary components: a format reward and an accuracy reward. As illustrated in Fig. 1, the format reward integrates multiple dimensions, including structural completeness of reasoning, action type correctness, and action validity. The accuracy reward is based on the longest common subsequence (LCS) between predicted and reference action sequences, effectively preserving action order and enhancing long-horizon performance.

The results on EmbodiedBench [59] show that RoboGPT-R1 significantly outperforms GPT-4o-mini and is competitive with closed-source models such as GPT-4o and Gemini-2.0-flash. Furthermore, its performance surpasses open-source models like Llama-3.2-90B, achieving a 23.33% higher overall score. Compared to the previous state-of-the-art [54], it yields a 20.33% improvement. On long-horizon tasks, it leads with an accuracy of 50%, demonstrating the superior reasoning capabilities of small models.

In summary, our contributions are as follows:

- We propose RoboGPT-R1, a two-stage training paradigm for embodied multi-step reasoning tasks. With RL training, RoboGPT-R1 develops the reasoning capability in complex tasks and environments, thereby enhancing its physical commonsense and error correction abilities.
- We design a reward function based on perception-reasoning-planning-action loop, with LCS reward effectively enhancing the model’s understanding and self-correction. This enables efficient and high-quality reward computation at a very low cost and demonstrates good reasoning capabilities on long-horizon tasks.
- We conduct extensive experiments on 6 tasks in 2 scenarios, including spatial perception, long-horizon reasoning, common-sense questions, and visual understanding. In seen scenarios, our method outperforms open-source general models and existing embodied planning models, achieving competitive performance compared to closed-source general models. Moreover, in unseen scenarios, it demonstrates superior reasoning ability and surpasses the state-of-the-art embodied planning model.

2 RELATED WORK

2.1 Embodied Planning

Embodied agents require not only active exploration, manipulation, and scene perception, but also embodied task planning capabilities [14, 15, 55, 69]. Embodied planning aims to decompose high-level natural language instructions into executable subtask sequences [16, 41], enabling the embodied agent to generate actionable steps within an interactive environment to complete complex behaviors. With the advent of large language models [29, 56, 58], natural language offers greater expressive flexibility than structured languages, making it possible to utilize LLMs to decompose complex plans into sub-plans in a fully automated manner [5, 52, 57, 62, 68]. For example, TaPA introduces an embodied task planner that grounds free-form instructions into executable plans, trained on a new multimodal benchmark (80 scenes, 15K instruction–plan pairs) [55]. It fine-tunes LLaMA with object lists from multi-view open-vocabulary detection. Additionally, SayCan [3] combines an LLM with reinforcement learning, leveraging the high-level reasoning capabilities of LLM to complement the value assessment of pre-trained skills, thereby laying a foundation for language in robotics and enabling the feasible scoring of actions. This can generate executable long-term plans suitable for real robots. While LLMs can generate preliminary plans based on commonsense reasoning, they lack constraints on the physical environment and the feasibility of actions [22, 43, 49, 62]. The emergence of VLMs [40, 55, 61] has led to their use as high-level planners, with the current mainstream

approach being to fine-tune VLMs based on demonstration data. Zhang et al. [68] extend the O1-style deep reasoning to embodied interactive tasks by coupling visual search with step-by-step planning, reflection, and verification, trained on synthesized Observation–Thought–Action trajectories to improve performance on AI2–THOR-style [25] tasks. Moreover, Reflective Planning [18] proposes a test-time computation framework that augments a pre-trained VLM with a reflection loop. It imagines future world states via a diffusion-based dynamics model, critiques potential suboptimality, and revises plans to improve multi-stage, long-horizon manipulation.

2.2 Reinforcement Learning for LLMs and VLMs

In recent years, with the emergence of reasoning models like OpenAI’s o1 [36], research on large language models (LLMs) has gradually shifted towards enhancing their reasoning capabilities through reinforcement learning (RL) [19, 32, 46, 48, 53, 67]. Numerous studies have explored ways to enhance the performance of LLMs in reasoning tasks, including solving mathematical problems [38, 47, 60] and coding [63]. A notable breakthrough in this field is DeepSeek-R1 [20], which experienced an “aha moment” during GRPO-based training, enabling the model to independently reflect and reevaluate its initial policy without any explicit guidance. Subsequently, several works [26, 30, 31, 39, 45, 65] have used reinforcement learning to enhance the reasoning capabilities of models in multimodal settings. For example, VLM-R1 [39] and Visual-RFT [30] extend R1-style reinforcement learning to vision-language models, sampling multiple answer outputs for each input and optimizing with GRPO using verifiable, task-specific rewards, resulting in stronger visual reasoning and perception enhancements over the SFT baseline. These advances demonstrate the potential of RL to propel large models from “imitation learning” to “emergent intelligence” [10]. Inspired by the R1 paradigm, this paper employs GRPO-based reinforcement learning to perform two-stage training on the model, systematically improving the planning ability and long-term consistency of embodied agents in multimodal task planning.

3 METHODOLOGY

3.1 Overview

In this section, we provide a brief introduction to the proposed RoboGPT-R1 framework. In contrast to previous approaches that solely rely on SFT, this study explores the incorporation of RL and reasoning techniques to better align the model with embodied planning tasks. Section 3.3 will introduce the two-stage learning scheme. As demonstrated in Fig. 1, the training of agents is comprised of two phases: an initial SFT phase, the purpose of which is to instill fundamental knowledge and elementary reasoning capabilities into the agent; and a subsequent reinforcement learning phase, the function of which is to utilize the GRPO policy optimization algorithm to enable the agent to continually explore, think, and learn independently. Next, to address issues such as unstable step-to-step coherence in multi-step reasoning, limited error recovery, and poor performance on long-term tasks, we designed a rule-based, verifiable reward to incentivize the VLM’s planning capabilities in Section 3.4.

3.2 Data Preparation

Embodied task planning in real-world indoor scenarios requires a substantial amount of multimodal data, encompassing both perceptual and physical knowledge. Given the impressive inference performance of large, closed-source models, data distillation can be employed to generate high-quality datasets. Following REBP [54], we employ the SFT dataset distilled from Gemini-2.0-flash, in the SFT phase. In the RL phase, the RFT dataset from REBP is augmented with task-relevant examples and unsuccessful exploratory tasks. In our preliminary studies, we observed that the quantity of examples contained within a dataset can result in contradictory outcomes regarding the training and testing of models. For instance, the direct application of n -shot (n denotes the number of examples) to the SFT phase results in the model overfitting to the provided demonstrations. However, during testing, there is a significant drop in the n -shot score. In contrast, an untrained model significantly improves its score when tested using n -shot, demonstrating its firm reliance on examples during the testing process. Consequently, it can be concluded that the injection of knowledge into a trained model may result in the model’s rigid adherence to predefined answer templates. This phenomenon impedes the model’s capacity to adapt to the complexity of the problem and the variability of the environment. Experiments have demonstrated that zero-shot training and testing models exhibit superior performance, while concurrently reducing the number of input tokens by approximately one-third (from 9,000+ to less than 6,000), thereby enhancing the efficiency of the training process. Consequently, we employ zero-shot processing uniformly for both training and testing data.

3.3 Two-stage Training Scheme

Stage 1: Initial Planning via SFT. To equip the base VLM with the fundamental capacity to generate multi-step reasoning, it is first necessary to undertake a supervised fine-tuning phase. This step is crucial because the reasoning patterns learned in subsequent reinforcement learning are significantly affected by the capabilities of the base model. Furthermore, the use of reinforcement learning as the sole training method is found to be significantly affected by data distribution, resulting in instability during the initial training stages. Therefore, we use a small amount of the data for an SFT-based warm-up stage. Following this, reinforcement learning training is conducted using the entire dataset. The findings of this study demonstrate that this approach enhances stability in the initial stages of training and rapidly integrates the base model with relevant knowledge, thereby enabling it to acquire a certain level of embodied reasoning knowledge and planning capabilities.

Stage 2: Enhancing reasoning with GRPO. The DeepSeek R1-zero algorithm employs the GRPO framework. Unlike reinforcement learning algorithms such as PPO, which require an additional critic model to estimate policy performance, GRPO directly compares groups of candidate responses, eliminating the need for a separate critic. Given a question q , GRPO samples N candidate responses $\{o_1, o_2, \dots, o_N\}$ from the policy π_θ and evaluates the quality of each response o_i using a reward function $R(q, o_i)$. To determine the relative quality of these responses, the algorithm normalizes the rewards by computing their mean and standard deviation and

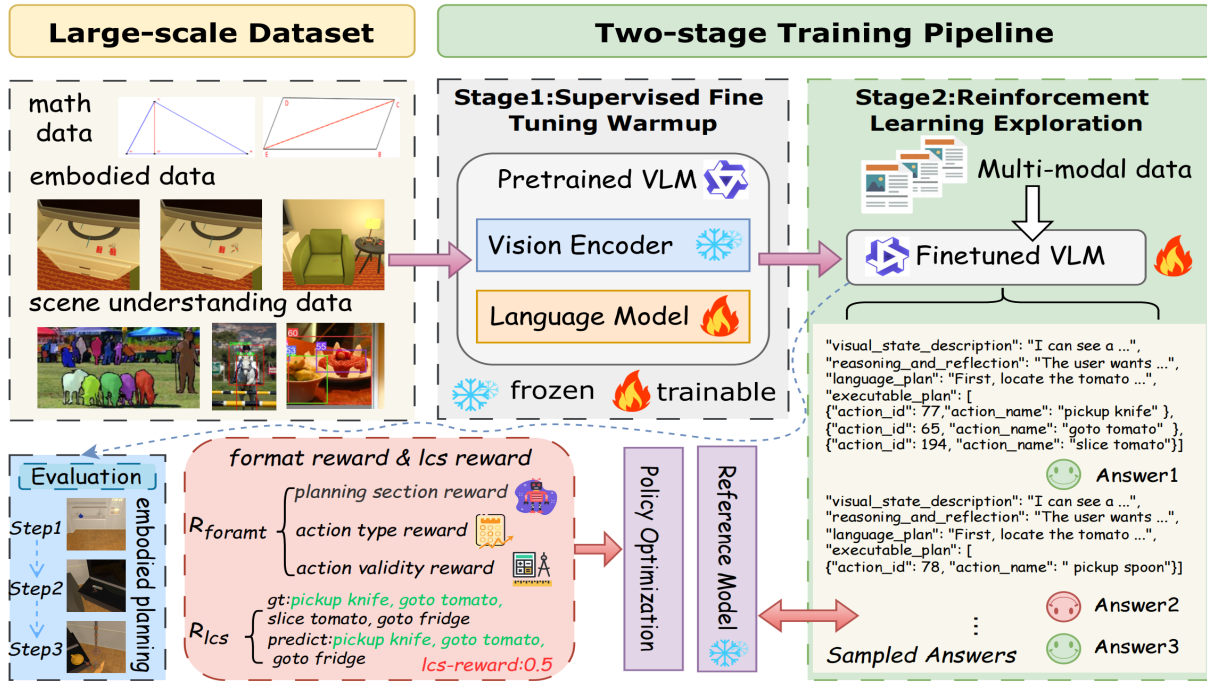


Figure 1: An overview of RoboGPT-R1. RoboGPT-R1 adopts a two-stage learning paradigm. In the initial phase, supervised fine-tuning introduces the model to data from mathematics, embodied tasks, and visual understanding, establishing a foundation for embodied reasoning. In the second phase, we apply GRPO-based reinforcement fine-tuning guided by a tailored reward function. The model is subsequently evaluated across six categories of tasks, including long-horizon planning and spatial reasoning.

subsequently derives the advantage as:

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_N\})}{\text{std}(\{r_1, r_2, \dots, r_N\})} \quad (1)$$

where A_i denotes the advantage of candidate o_i measured against the other samples in the group. To encourage the model to generate responses with higher advantages within the group, the group penalizes large deviations from the reference policy π_{ref} . The objective is as follows:

$$\mathcal{J}_{grpo}(\theta) = \mathbb{E}_{q, \{o_i\}_{i=1}^N \sim \pi_{\theta_{old}}(q)} \left[\frac{1}{N} \sum_{i=1}^N \{\min(\ell_1 \cdot A_i, \ell_2 \cdot A_i) - \beta D_{KL}(\pi_{\theta} \parallel \pi_{ref})\} \right] \quad (2)$$

$$\ell_1 = \frac{\pi_{\theta}(o_i | q)}{\pi_{\theta_{old}}(o_i | q)}; \ell_2 = \text{clip}\left(\frac{\pi_{\theta}(o_i | q)}{\pi_{\theta_{old}}(o_i | q)}, 1 - \epsilon, 1 + \epsilon\right) \quad (3)$$

where ϵ is the clipping hyperparameter and β controls the KL penalty against a reference policy π_{ref} . Inspired by DeepSeek-R1, our approach incorporates two complementary rewards: the accuracy reward and the format reward. Next, we introduce the reward function R designed for robotic planning tasks.

3.4 Reward Design

In Section 3.3, we have briefly introduced the general GRPO algorithm. Given that embodied planning requires an agent to complete

complex tasks in a real or simulated physical environment based on natural language instructions, we design a set of reward functions specifically for this purpose. The following sections describe the format reward and the accuracy reward, respectively.

3.4.1 Format Reward. To ensure output regularity and facilitate the extraction of the reasoning process and the final result, most R1-style approaches enclose the reasoning within `<think></think>` tags and the final plan within `<answer></answer>` tags. If the generated output deviates from this structure, the format reward is assigned a value of zero. Inspired by REBP [54], in embodied multi-step planning, the agent must generate responses that are not only semantically meaningful but also structurally executable. Unlike conventional text generation tasks, where free-form output may be acceptable, embodied planning requires a higher level of structural rigour to ensure that each response remains interpretable and executable by downstream systems. First, the model should follow a clear cognitive loop—just like how humans plan their actions. Before acting, we reflect on the task, observe the environment, make a plan, and then execute it. Second, the generation of invalid or fabricated actions should be penalized. In summary, we set the following reward format, which consists of three parts:

$$R_{format} = 0.3 \cdot R_{section} + 0.3 \cdot R_{type} + 0.4 \cdot R_{validity} \quad (4)$$

The section reward $R_{section}$ evaluates whether all required fields (`visual_state_description`, `reasoning_and_reflection`, `language_plan`, `executable_plan`) are present and correctly typed, and is computed

as:

$$R_{\text{section}} = \frac{1}{|S|} \sum_{s \in S} \mathbf{1}[\text{type}(o_s) = T_s] \quad (5)$$

where $\mathbf{1}$ is an indicator function, $S = \{s_1, s_2, s_3, s_4\}$ includes the four fields, T_s denotes the expected type of each field and o_s is the value of field s in the output object. The type reward R_{type} checks whether each action step is well-formed, defined as:

$$R_{\text{type}} = \frac{1}{m} \sum_{i=1}^m \mathbf{1}[\hat{y}_i^{\text{id}} \in \text{Int} \wedge \hat{y}_i^{\text{name}} \in \text{Str}] \quad (6)$$

In this formulation, m denotes the number of action steps, \hat{y}_i^{id} and \hat{y}_i^{name} represent the action id and name of step i respectively. Here Int denotes the set of integers, while Str represents the set of non-empty strings. Finally, the validity reward R_{validity} measures whether each action id–name pair matches the predefined action dictionary, given by:

$$R_{\text{validity}} = \frac{1}{|C|} \sum_{i \in C} \mathbf{1}[\text{norm}(\hat{y}_i^{\text{name}}) = \text{norm}(\mathcal{D}_{\text{action}}[\hat{y}_i^{\text{id}}])] \quad (7)$$

where C is the index set of steps whose action ids are defined in the action dictionary $\mathcal{D}_{\text{action}}$ and $\text{norm}(\cdot)$ is a normalization function that lowercases and trims whitespace. The action dictionary $\mathcal{D}_{\text{action}}$ defines a mapping from action ids to their corresponding names.

Unlike REBP [54], our setup employs dynamic action IDs that vary across tasks and environments, preventing the model from relying on memorization of a fixed action set. Instead, we want it to learn the meaning of actions and use them correctly. In summary, such a reward design can guide the model to generate structured outputs that conform to the task execution closed loop, and avoid hallucinations or inconsistent outputs through self-understanding and thinking.

3.4.2 LCS Reward. In embodied multi-step planning, the correctness of individual actions is not sufficient—the order in which actions are executed is often critical to task success. For example, placing an object before picking it up may involve the right actions but in a logically invalid sequence. Traditional token-level matching or step-wise accuracy metrics fail to penalize such disorder, treating unordered but correct actions as equally valid. Besides, in long-horizon tasks, plans can extend over dozens of steps, where strict reward strategies like exact matching become too rigid to reflect realistic performance. A model might make early mistakes yet recover in later steps to complete the task successfully. However, prefix-based accuracy rewards, such as those used in REBP [54], overlook this “error recovery” behavior, leading to sparse and less informative reward signals. To address these problems, we design the accuracy reward based on the Longest Common Subsequence (LCS) between the predicted and reference action sequences. By computing the LCS over action names, we enforce both content accuracy and sequence coherence. This approach is robust to local deviations while maintaining global alignment and remains effective even as task length increases. We define the model generation sequence as $\hat{Y} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_m)$ and the reference sequence as

$Y = (y_1, y_2, \dots, y_n)$. The detailed accuracy reward is as follows:

$$\text{LCS}(i, j) = \begin{cases} 0, & \text{if } i = 0 \text{ or } j = 0 \\ \text{LCS}(i-1, j-1) + 1, & \text{if } \hat{y}_i = y_j \\ \max(\text{LCS}(i-1, j), \text{LCS}(i, j-1)), & \text{otherwise} \end{cases} \quad (8)$$

$$R_{\text{lcs}} = \frac{k}{n}, \quad (9)$$

where n is the length of the reference sequence $|Y|$, k is the length of the longest common subsequence $\text{LCS}(m, n)$. To evaluate the effectiveness of our proposed accuracy reward, we conduct an ablation study. Compared to prefix accuracy (as used in REBP) and standard step-wise matching, the LCS-based accuracy reward shows stronger performance in multi-step planning tasks.

3.4.3 Overall Reward. To jointly encourage structural correctness and sequential accuracy, we define the overall reward as a weighted combination of the format reward and the LCS-based accuracy reward:

$$R = 0.2 \cdot R_{\text{format}} + 0.8 \cdot R_{\text{lcs}} \quad (10)$$

This overall reward formulation has two key purposes in embodied planning: enforcing structural correctness and promoting action-level accuracy. The format reward R_{format} enforces a fixed output structure and ensures each action step is well-formed and consistent with the current scene and instructions. Meanwhile, the accuracy reward R_{lcs} evaluates whether the predicted action sequence aligns with the reference plan, not only in content but also in order. By combining these two aspects, the overall reward encourages the model to reason more effectively and generate executable plans with reasonable length and structure.

4 EXPERIMENTS

4.1 Experimental Settings

Evaluation. We evaluate RoboGPT-R1 in EmbodiedBench [59], a unified benchmarking suite for multimodal embodied planning agents. EmbodiedBench offers standardized protocols and interfaces that cover multi-sensory input, language instructions, and long-horizon decision making, enabling consistent and reproducible comparisons across task suites. Specifically, we focus on its two constituent suites: EB-ALFRED and EB-Habitat. The former is rooted in the ALFRED ecosystem and targets instruction-following household tasks (e.g., pick-and-place, cleaning, and mobile manipulation) that emphasize object state tracking and stepwise dependencies; the latter builds on the Habitat ecosystem and emphasizes navigation and interaction in 3D environments, with observation distributions, scene layouts, and action semantics that differ markedly from EB-ALFRED [37, 42]. Because our training data are primarily associated with EB-ALFRED/ALFRED, we treat performance on EB-ALFRED as in-domain and use it to gauge method effectiveness. Results on EB-Habitat are regarded as out-of-domain and used to assess generalization.

Unless otherwise noted, all evaluation settings follow the benchmark defaults; the only test variable we modify is the number of examples in context (n_{shots}) in the control input.

Table 1: Success rates of diverse models on EB-ALFRED and EB-Habitat. Entries without any symbol are sourced from the EMBench leaderboard. Symbol † indicates results directly cited from the REBP paper, based on their evaluations. Symbol ‡ marks results obtained through our own reproduction by querying the official API. All scores are reported as percentages (%). For each metric, the best-performing result is highlighted with a gray background.

Method	Params	EB-ALFRED (seen)						EB-Habitat (unseen)							
		Avg.	Base	Common	Complex	Visual	Spatial	Long	Avg.	Base	Common	Complex	Visual	Spatial	Long
• Type: Closed-Source General Model															
Gemini-2.0-flash [13]	-	52.30	62	48	54	46	46	58	42.30	82	38	38	36	34	26
Qwen-VL-Max [11]	-	41.30	44	48	44	42	38	32	45.30	74	40	50	42	30	36
GPT-4.1 ‡ [35]	-	64.67	70	64	70	62	62	60	50.67	90	38	50	36	46	44
GPT-4o ‡ [34]	-	51.67	54	46	58	52	52	48	57.00	84	42	62	38	62	54
GPT-4o-mini ‡ [33]	-	34.00	66	56	68	14	0	0	35.00	70	24	36	30	30	20
• Type: Open-Source General Model															
Llama-3.2-90B-Vision-Ins [4]	90B	32.00	38	34	44	28	32	16	40.30	94	24	50	32	28	14
InternVL2.5-8B [9]	8B	2.00	4	6	2	0	0	0	11.30	36	4	0	10	16	2
Gemma-3-12b-its [44]	12B	25.70	32	26	38	26	20	12	23.00	58	10	24	18	24	4
Qwen2.5-VL-72B-Ins ‡ [12]	72B	43.67	62	36	48	40	44	32	50.33	92	38	48	34	46	44
Qwen2.5-VL-7B-Ins ‡ [12]	7B	2.67	6	2	6	0	2	0	15.00	42	6	22	12	4	4
Qwen2.5-VL-3B-Ins ‡ [12]	3B	1.33	2	2	0	0	4	0	14.67	34	0	18	18	14	4
• Type: Embodied Planning Model															
RoboBrain † [23]	7B	0.33	2	0	0	0	0	0	15.30	38	6	18	8	18	4
Tapa † [55]	7B	0.00	0	0	0	0	0	0	0.00	0	0	0	0	0	0
REBP ‡ [54]	7B	35.00	52	46	46	28	32	6	18.33	50	6	18	14	14	8
RoboGPT-R1 (ours) ‡	3B	55.33	62	56	64	50	50	50	22.00	64	8	18	20	12	10

Baselines. We evaluate both the Qwen and GPT model series, including Qwen2.5-VL-72B-Ins, Qwen2.5-VL-7B-Ins, Qwen2.5-VL-3B-Ins, as well as GPT-4.1, GPT-4o and GPT-4o-mini. Closed-source models are assessed via their official APIs, while open-source models are tested through local deployment. Performance results for additional baselines are obtained from the REBP and EmbodiedBench leaderboards. We group baselines into three categories, and present the corresponding results below:

- (1) **General closed-source models:** including five representative proprietary multimodal models from the GPT, Gemini, and Qwen families: Gemini-2.0-Flash [13], Qwen-VL-Max [11], GPT-4.1 [35], GPT-4o [34] and GPT-4o-mini [33].
- (2) **General open-source models:** Comprising open-source models from multiple series—including Qwen2.5-VL-3B/7B/72B-Ins [12], LLaMA-3.2-90B-Vision-Ins [4], InternVL2.5-8B [9], and Gemma-3-12B-it [44].
- (3) **Embodied domain-specific models:** focusing on models tailored for embodied reasoning and planning, such as REBP [54], RoBoBrain [23], TaPa [55] and ours.

Dataset. We process the REBP public dataset [54] to generate a base dataset and an augmented dataset. The base dataset is directly distilled from the EB-ALFRED tasks in EmbodiedBench [59], serving as a benchmark-aligned dataset used in the SFT phase to endow the model with initial multimodal embodied-planning skills. The main body of the augmented dataset originates from the open-source ALFRED trajectory dataset [42]. While its content is similar to EB-ALFRED, it exhibits significant differences in details such as action space, visual appearance, task types, and task length, making it a benchmark-adjacent (near-domain) dataset. The augmented dataset includes all embodied planning data in the base dataset to prevent catastrophic forgetting. It is employed in the RFT phase to improve reasoning robustness under near-domain distributions, thereby further enhancing planning performance on EmbodiedBench.




Figure 2: Reward Curves in RFT. Our LCS-based accuracy reward provides an appropriate and dense learning signal for RFT, rising steadily from 0.30 to 0.80. The format reward, already aligned by SFT, starts around 0.95 at the onset of RFT and saturates within ~20 steps, stabilizing near 1.0. The overall reward is computed as a weighted sum of the accuracy and format rewards with weights 0.8 and 0.2, respectively, and it increases in tandem with the steady improvement of the accuracy reward.

Training. We adopt Qwen2.5-VL-3B-Instruct as the multimodal base model and employ a two-stage training scheme. In the first stage, we perform full-parameter SFT on the *base* dataset to endow the model with initial planning skills aligned with *EB-ALFRED*. In the second stage, we conduct reinforcement fine-tuning (RFT) with GRPO on the *augmented* dataset, aiming to improve reasoning and generalization from data that are not strictly benchmark-matched. SFT is implemented with LLaMA-Factory [21] and trained on 8× Ascend 910B3 64GB NPUs for about 1.5 hours; RFT is implemented with VERL [1] and trained on 4× NVIDIA H20 96GB GPUs for about 25 hours.

Table 2: Training-Strategy Ablation Result (success rate, %). After SFT, the model acquires core embodied planning competence, raising the average success rate from 1.33% to 42.00% while still lagging on long-horizon tasks. With subsequent RFT, performance further increases to 55.33%, with especially pronounced gains on long-horizon tasks (26% → 50%).

Base	SFT	RFT	Avg.	Base	Common	Complex	Visual	Spatial	Long
✓	✗	✗	1.33 _(+0.00)	2 _(+0.00)	2 _(+0.00)	0 _(+0.00)	0 _(+0.00)	4 _(+0.00)	0 _(+0.00)
✓	✓	✗	42.00 _(+41.33)	48 ₍₊₄₆₎	44 ₍₊₄₂₎	58 ₍₊₅₈₎	38 ₍₊₃₈₎	38 ₍₊₃₄₎	26 ₍₊₂₆₎
✓	✓	✓	55.33 _(+54.00)	62 ₍₊₆₀₎	56 ₍₊₅₄₎	64 ₍₊₆₄₎	50 ₍₊₅₀₎	50 ₍₊₄₆₎	50 ₍₊₅₀₎

Table 3: Data-Source Ablation Result (success rate, %). Base denotes the in-domain dataset distilled from EMBench; Aug is the near-domain ALFRED-derived set. Training only with SFT on Base reaches 40.00% on average, whereas replacing Base with Aug for SFT collapses performance to 6.00%, indicating poor transfer under pure supervision. Continuing RFT on Base after SFT on Base yields only a modest gain (40.00% → 44.33%), while using Aug during RFT achieves the best results (55.33%), showing that only the RL (RFT) stage can effectively absorb near-domain data and transfer it to the target task.

Model	Avg.	Base	Common	Complex	Visual	Spatial	Long
Base Model (Qwen2.5-VL-3B)	1.33	2	2	0	0	4	0
Only SFT w Base (ours)	42.00	48	44	58	38	38	26
Only SFT w Aug	6.00	14	6	12	2	2	0
SFT+ RFT w Base	44.33	56	56	54	32	36	32
SFT+ RFT w Aug (ours)	55.33	62	56	64	50	50	50

4.2 Main Results

Training Results in EB-ALFRED. Table 1 reports the evaluation results in EB-ALFRED. RoboGPT-R1 attains an average success rate of 55.33% across six sub-task suites. This performance **significantly outperforms** multiple strong baselines: it surpasses the closed-source **GPT-4o (51.67%)** and **GPT-4o-mini (34.00%)**, and trails only **GPT-4.1 (64.67%)**. Among open-source general models, it notably outperforms the larger-parameter **Qwen2.5-VL-72B-Instruct (43.67%)**. Relative to the small-scale model with similar parameters **Qwen2.5-VL-3B-Instruct**, our approach yields an approximately 54% relative improvement in average success. Compared with the embodied specialist **REBP (35.00%)**, RoboGPT-R1 leads by nearly 20 percentage points overall and shows a pronounced advantage on long-horizon tasks: REBP achieves only 6%, whereas RoboGPT-R1 reaches 50%. Notably, despite using only 3B parameters, RoboGPT-R1 delivers this level of performance under a small-model, low-inference-cost setting, highlighting the parameter efficiency and effectiveness of our method.

Generalization Results in EB-Habitat. Table 1 summarizes the EB-Habitat results. RoboGPT-R1 attains an average success rate of 22% across six sub-task suites, representing a 7% improvement over the base model **Qwen2.5-VL-3B-Instruct (14.67%)** and outperforming both the 7B-parameter models **Qwen2.5-VL-7B-Instruct (15.00%)** and **REBP (18.33%)**. Although a performance gap remains compared to **Qwen2.5-VL-72B-Instruct** and several closed-source general models, these *out-of-domain* results indicate that our approach substantially improves generalization and transferability of the embodied model.

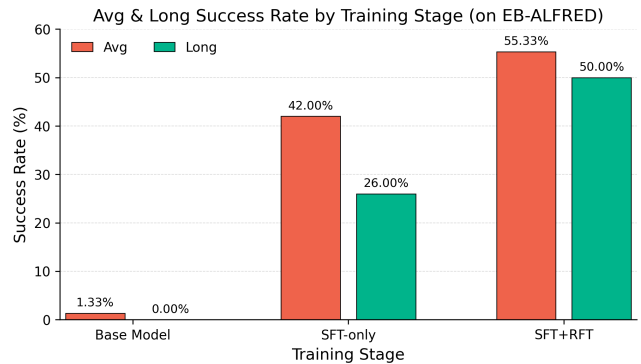


Figure 3: Success rates with different stages. Bars show the macro average and the long-horizon score for the base model, SFT-only, and SFT+RFT model. SFT establishes initial embodied planning (Avg: 1.33%→42.00%) but leaves long-horizon performance limited (26%). Adding RFT lifts the averages further (to 55.33%) and markedly improves long-horizon planning (to 50%), validating the effectiveness of our two-stage framework—SFT for foundational competence and RFT for the additional gains needed to solve long-horizon tasks.

4.3 Ablation Study

We conduct two sets of ablations on EB-ALFRED to disentangle the effects of the two training stages and the two data sources used in our framework. For brevity, we refer to the *Base dataset* and the *Aug dataset* as **Base** and **Aug**, respectively.

Table 4: Accuracy Reward Comparison Result (success rates, %). We compare three GRPO accuracy rewards used during RFT—Step Accuracy, REBP Acc., and LCS (ours)—with “RFT Base” as the reference row for deltas. Bracketed scores denote the change relative to the reference. LCS (ours) delivers the strongest overall performance (Avg. 55.33%) and the largest increase on long-horizon tasks (+24%), outperforming the alternatives across most sub-suites.

Accuracy Reward	Avg.	Base	Common	Complex	Visual	Spatial	Long
RFT Base	42.00 (+0.00)	48 (+0.00)	44 (+0.00)	58 (+0.00)	38 (+0.00)	38 (+0.00)	26 (+0.00)
Step Accuracy	43.67 (+1.67)	52 (+4.00)	56 (+12.00)	54 (-4.00)	40 (+2.00)	40 (+2.00)	20 (-6.00)
REBP Acc.	48.33 (+6.33)	62 (+14.00)	52 (+8.00)	58 (+0.00)	46 (+8.00)	38 (+0.00)	34 (+8.00)
LCS (ours)	55.33 (+13.33)	62 (+14.00)	56 (+12.00)	64 (+8.00)	50 (+12.00)	50 (+12.00)	50 (+24.00)

Training-Strategy Ablation. We compare the performance changes on *EB-ALFRED* among the base model (no fine-tuning), the SFT-only model, and the final SFT+RFT model, as shown in Table 2 and Fig. 3. After SFT, the average success rate rises from **1.33%** (base model) to **42.00%**, indicating that SFT learns the initial embodied planning competence. However, performance on long-horizon tasks remains limited (**26%**). With subsequent RFT, performance improves across all sub-task suites, yielding an average of **55.33%**; notably, the long-horizon score increases from **26%** to **50%**, highlighting the particular effectiveness of RFT for complex, extended plans.

Data-Source Ablation. As shown in Table 3, these ablations highlight the importance of near-domain data during RFT. (1) *SFT on Aug only*: Replacing Base with Aug for SFT reduces the average success from **42.00%** to **6.00%**, suggesting that, under supervised learning, near-domain data alone does not transfer effectively to the target task. (2) *RFT on Base*: Starting from SFT on Base and continuing RFT on Base (instead of Aug) yields only a modest increase from **42.00%** to **44.33%**. Compared with the SFT+RFT setting that uses Aug during RFT (**55.33%**), this indicates that incorporating near-domain data in the RFT stage is essential.

Ablation Analysis. We summarize three observations: (i) SFT establishes the initial planning ability but leaves a gap on long-horizon tasks; (ii) RFT closes this gap with the strongest gains on long-horizon sub-task suites; (iii) combining RFT with Aug is essential, as Aug brings little benefit in SFT but provides substantial gains when used during RFT.

4.4 Accuracy Reward Comparison

During the RFT stage, we instantiate the accuracy reward in the GRPO algorithm [2] as the normalized *Longest Common Subsequence (LCS)* length between the generated and reference trajectories. Thanks to the appropriate and dense learning signal provided by the LCS-based accuracy reward, the reward curves rise steadily throughout RFT, as shown in Fig. 2. This section conducts a head-to-head comparison of three *accuracy rewards* to assess the suitability and advantages of our LCS reward for embodied planning.

Compared rewards. We set the accuracy reward in GRPO to one of:

- **LCS reward (ours):** the ratio of the longest common subsequence between the generated and reference trajectories (Normalized LCS);
- **Step Accuracy:** the ratio of strictly matched actions/instructions step-by-step;

- **REBP reward:** the multi-step, progress-style signal used in REBP [54].

We evaluate on *EB-ALFRED* with the same settings as in the main experiments. To ensure fairness, we fix the backbone model, training data, total update budget, optimization hyperparameters, and the format reward term, and vary only the accuracy reward. The primary metrics are the *average success rate* and the scores of each sub-task suite (with emphasis on *long-horizon* tasks).

Comparison Results. As summarized in Table 4, using **Step Accuracy** yields only a slight improvement in the average success rate (from **42.00%** to **43.67%**), while the long-horizon score decreases (from **26%** to **20%**). The **REBP reward** improves the average to **48.33%** and raises the long-horizon score from **26%** to **34%**. In contrast, our **LCS reward** achieves the largest gains: the average increases from **42.00%** to **55.33%**, and the long-horizon score from **26%** to **50%**. These findings demonstrate the suitability of the LCS-based accuracy reward for embodied planning and its advantages over alternatives under the same training budget.

5 CONCLUSIONS

In this work, we propose RoboGPT-R1, a two-stage training framework for embodied planning. Stage 1 (SFT) equips the model with initial instruction-following and planning priors. Stage 2 (RFT) performs reinforcement fine-tuning with GRPO and an LCS-based, accuracy reward (paired with format constraints), providing dense and verifiable feedback. This design overcomes the limitations of SFT-only behavior cloning, which fails to adequately elicit the reasoning capabilities of VLMs and often undermines in-domain performance when leveraging near-domain data. These shortcomings result in poor performance on long-horizon tasks and brittle out-of-domain generalization. Evaluated on EmbodiedBench, the 3B-parameter model trained with our framework surpasses general-purpose VLM baselines such as Qwen2.5-VL-72B and GPT-4o, and substantially outperforms other 7B-scale embodied planners, with especially pronounced gains on long-horizon subtasks.

REFERENCES

- [1] Alibaba DAMO Academy. 2024. EasyR1: A unified framework for reward modeling and RLHF. <https://github.com/alibaba/EasyR1>.
- [2] Alibaba DAMO Academy. 2024. GRPO: Generalized Reward Preference Optimization for LLM alignment. <https://github.com/alibaba/GRPO>.
- [3] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil J

- Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. 2022. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. arXiv:2204.01691 [cs.RO] <https://arxiv.org/abs/2204.01691>
- [4] Meta AI. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/>.
- [5] Hao Bai, Yifei Zhou, Mert Cemri, Jiayi Pan, Alane Suhr, Sergey Levine, and Aviral Kumar. 2024. DigiRL: Training In-The-Wild Device-Control Agents with Autonomous Reinforcement Learning. arXiv:2406.11896 [cs.LG] <https://arxiv.org/abs/2406.11896>
- [6] Zitong Bo, Yue Hu, Jiming Ma, Mingliang Zhou, Junhui Yin, Yachen Kang, Yuqi Liu, Tong Wu, Diyun Xiang, and Hao Chen. 2025. Reinforced Embodied Planning with Verifiable Reward for Real-World Robotic Manipulation. arXiv:2509.25852 [cs.RO] <https://arxiv.org/abs/2509.25852>
- [7] Shaoyu Chen, Bo Jiang, Hao Gao, Bencheng Liao, Qing Xu, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggong Wang. 2024. VADv2: End-to-End Vectorized Autonomous Driving via Probabilistic Planning. arXiv:2402.13243 [cs.CV] <https://arxiv.org/abs/2402.13243>
- [8] Yaran Chen, Wenbo Cui, Yuanwen Chen, Mining Tan, Xinyao Zhang, Jinrui Liu, Haoran Li, Dongbin Zhao, and He Wang. 2025. RoboGPT: An LLM-Based Long-Term Decision-Making Embodied Agent for Instruction Following Tasks. *IEEE Transactions on Cognitive and Developmental Systems* 17, 5 (2025), 1163–1174. <https://doi.org/10.1109/TCDSS.2025.3543364>
- [9] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, Lixin Gu, Xuehui Wang, Qingyun Li, Yiming Ren, Zixuan Chen, Jiapeng Luo, Jiahao Wang, Tan Jiang, Bo Wang, Conghui He, Botian Shi, Xingcheng Zhang, Han Lv, Yi Wang, Wenqi Shao, Pei Chu, Zhongying Tu, Tong He, Zhiyong Wu, Huipeng Deng, Jiaye Ge, Kai Chen, Kaipeng Zhang, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhao Wang. 2025. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling. arXiv:2412.05271 [cs.CV] <https://arxiv.org/abs/2412.05271>
- [10] Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V. Le, Sergey Levine, and Yi Ma. 2025. SFT Memorizes, RL Generalizes: A Comparative Study of Foundation Model Post-training. arXiv:2501.17161 [cs.AI] <https://arxiv.org/abs/2501.17161>
- [11] Alibaba Cloud. 2024. Qwen-VL-Max: Large Vision-Language Model. <https://github.com/QwenLM/Qwen-VL>
- [12] Alibaba Cloud. 2025. Qwen2.5-VL Technical Report. arXiv:2502.13923 [cs.CL] <https://arxiv.org/abs/2502.13923>
- [13] Google DeepMind. 2024. Introducing Gemini 2.0: Our new AI model for the agentic era. <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/>.
- [14] Jiafei Duan, Samson Yu, Hui Li Tan, Hongyuan Zhu, and Cheston Tan. 2022. A Survey of Embodied AI: From Simulators to Research Tasks. arXiv:2103.04918 [cs.AI] <https://arxiv.org/abs/2103.04918>
- [15] Zane Durante, Qiuyuan Huang, Naoki Wake, Ran Gong, Jae Sung Park, Bidipta Sarkar, Rohan Taori, Yusuke Noda, Demetri Terzopoulos, Yejin Choi, Katsushi Ikeuchi, Hoi Vo, Li Fei-Fei, and Jianfeng Gao. 2024. Agent AI: Surveying the Horizons of Multimodal Interaction. arXiv:2401.03568 [cs.AI] <https://arxiv.org/abs/2401.03568>
- [16] Zhaoye Fei, Li Ji, Siyin Wang, Junhao Shi, Jingjing Gong, and Xipeng Qiu. 2025. Unleashing Embodied Task Planning Ability in LLMs via Reinforcement Learning. arXiv:2506.23127 [cs.CL] <https://arxiv.org/abs/2506.23127>
- [17] Sicheng Feng, Kaiwen Tuo, Song Wang, Lingdong Kong, Jianke Zhu, and Huan Wang. 2025. RewardMap: Tackling Sparse Rewards in Fine-grained Visual Reasoning via Multi-Stage Reinforcement Learning. arXiv:2510.02240 [cs.CV] <https://arxiv.org/abs/2510.02240>
- [18] Yunhai Feng, Jiaming Han, Zhuoran Yang, Xiangyu Yue, Sergey Levine, and Jianlan Luo. 2025. Reflective Planning: Vision-Language Models for Multi-Stage Long-Horizon Robotic Manipulation. arXiv:2502.16707 [cs.RO] <https://arxiv.org/abs/2502.16707>
- [19] Yuqian Fu, Tinghong Chen, Jiajun Chai, Xihuai Wang, Songjun Tu, Guojun Yin, Wei Lin, Qichao Zhang, Yuanheng Zhu, and Dongbin Zhao. 2025. SRFT: A Single-Stage Method with Supervised and Reinforcement Fine-Tuning for Reasoning. arXiv:2506.19767 [cs.CL] <https://arxiv.org/abs/2506.19767>
- [20] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirog Ma, Xiao Bi, Xiaokang Zhang, Xingka Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiusi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shutang Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaoshua Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shao, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyuan Lu, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhanli Sha, Zhe Fu, Zhean Xu, Zhenda Xi, Zhengyan Zhang, Zhenwen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature* 645, 8081 (Sept. 2025), 633–638. <https://doi.org/10.1038/s41586-025-09422-z>
- [21] Hiyouga. 2023. LLaMA Factory: Open-source instruction tuning framework for LLMs. <https://github.com/hiyouga/LLaMA-Factory>.
- [22] Wenlong Huang, Fei Xia, Dhruv Shah, Danny Driess, Andy Zeng, Yao Lu, Pete Florence, Igor Mordatch, Sergey Levine, Karol Hausman, and Brian Ichter. 2023. Grounded decoding: guiding text generation with grounded models for embodied agents. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 2606, 26 pages.
- [23] Yuheng Ji, Huajie Tan, Jiayu Shi, Xiaoshuai Hao, Yuan Zhang, Hengyuan Zhang, Pengwei Wang, Mengdi Zhao, Yao Mu, Pengju An, Xinda Xue, Qinghang Su, Huaihai Lyu, Xiaolong Zheng, Jiaming Liu, Zhongyuan Wang, and Shanghang Zhang. 2025. RoboBrain: A Unified Brain Model for Robotic Manipulation from Abstract to Concrete. arXiv:2502.21257 [cs.RO] <https://arxiv.org/abs/2502.21257>
- [24] Bo Jiang, Shaoyu Chen, Qian Zhang, Wenyu Liu, and Xinggong Wang. 2025. AlphaDrive: Unleashing the Power of VLMs in Autonomous Driving via Reinforcement Learning and Reasoning. arXiv:2503.07608 [cs.CV] <https://arxiv.org/abs/2503.07608>
- [25] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, Aniruddha Kembhavi, Abhinav Gupta, and Ali Farhadi. 2022. AI2-THOR: An Interactive 3D Environment for Visual AI. arXiv:1712.05474 [cs.CV] <https://arxiv.org/abs/1712.05474>
- [26] Xinhao Li, Ziang Yan, Desen Meng, Lu Dong, Xiangyu Zeng, Yan He, Yali Wang, Yu Qiao, Yi Wang, and Limin Wang. 2025. VideoChat-R1: Enhancing Spatio-Temporal Perception via Reinforcement Fine-Tuning. arXiv:2504.06958 [cs.CV] <https://arxiv.org/abs/2504.06958>
- [27] Wenlong Liang, Rui Zhou, Yang Ma, Bing Zhang, Songlin Li, Yijia Liao, and Ping Kuang. 2025. Large Model Empowered Embodied AI: A Survey on Decision-Making and Embodied Learning. arXiv:2508.10399 [cs.RO] <https://arxiv.org/abs/2508.10399>
- [28] Xin Liu, Haoran Li, and Dongbin Zhao. 2025. Videos are Sample-Efficient Supervisions: Behavior Cloning from Videos via Latent Representations. arXiv:2512.21586 [cs.LG] <https://arxiv.org/abs/2512.21586>
- [29] Yang Liu, Weixing Chen, Yongjie Bai, Xiaodan Liang, Guanbin Li, Wen Gao, and Liang Lin. 2025. Aligning Cyber Space with Physical World: A Comprehensive Survey on Embodied AI. arXiv:2407.06886 [cs.CV] <https://arxiv.org/abs/2407.06886>
- [30] Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. 2025. Visual-RFT: Visual Reinforcement Fine-Tuning. arXiv:2503.01785 [cs.CV] <https://arxiv.org/abs/2503.01785>
- [31] Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Tiancheng Han, Botian Shi, Wenhao Wang, Junjun He, Kaipeng Zhang, Ping Luo, Yu Qiao, Qiaosheng Zhang, and Wenqi Shao. 2025. MM-Eureka: Exploring the Frontiers of Multimodal Reasoning with Rule-based Reinforcement Learning. arXiv:2503.07365 [cs.CV] <https://arxiv.org/abs/2503.07365>
- [32] Youssef Mroueh. 2025. Reinforcement Learning with Verifiable Rewards: GRPO's Effective Loss, Dynamics, and Success Amplification. arXiv:2503.06639 [cs.LG] <https://arxiv.org/abs/2503.06639>
- [33] OpenAI. 2024. GPT-4o mini: Advancing cost-efficient intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>.
- [34] OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>.
- [35] OpenAI. 2025. Introducing GPT-4.1 in the API. <https://openai.com/index/gpt-4-1/>.

- [36] OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Itfimej, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich, Andrew Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghorbani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botao Hao, Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, Camillo Lugaresi, Cary Bassin, Cary Hudson, Chak Ming Li, Charles de Bourcy, Chelsea Voss, Chen Shen, Chong Zhang, Chris Koch, Chris Orsinger, Christopher Hesse, Claudia Fischer, Clive Chan, Dan Roberts, Daniel Kappler, Daniel Levy, Daniel Selsam, David Dohan, David Farhi, David Mely, David Robinson, Dimitris Tsipras, Doug Li, Dragos Oprica, Eben Freeman, Eddie Zhang, Edmund Wong, Elizabeth Proehl, Enoch Cheung, Eric Mitchell, Eric Wallace, Erik Ritter, Evan Mays, Fan Wang, Felipe Petroski Such, Filippo Raso, Florencia Leoni, Foivos Tsimpouras, Francis Song, Fred von Lohmann, Freddie Sulit, Geoff Salmon, Giambattista Parascandolo, Gildas Chabot, Grace Zhao, Greg Brockman, Guillaume Leclerc, Hadi Salman, Jakob Pachocki, James Lennon, Jason Wei, Jean Harb, Jerry Twore, Jiacheng Feng, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joaquin Quiñero Candela, Joe Palermo, Joel Parish, Johannes Heidecke, John Hallman, John Rizzo, Jonathan Gordon, Jonathan Uesato, Jonathan Ward, Joost Huizinga, Julie Wang, Kai Chen, Kai Xiao, Karan Singhal, Karina Nguyen, Karl Cobbe, Katy Shi, Kayla Wood, Kendra Rimbach, Keren Gu-Lemberg, Kevin Liu, Kevin Lu, Kevin Stone, Kevin Yu, Lama Ahmad, Lauren Yang, Leo Liu, Leon Maksin, Leyton Ho, Liam Fedus, Lilian Weng, Linden Li, Lindsay McCallum, Lindsey Held, Lorenz Kuhn, Lukas Kondraciuk, Lukasz Kaiser, Luke Metz, Madelaine Boyd, Maja Trebacz, Manas Joglekar, Mark Chen, Marko Tintor, Mason Meyer, Matt Jones, Matt Kaufner, Max Schwarzer, Meghan Shah, Mehmet Yatbaz, Melody Y. Guan, Mengyuan Xu, Mengyuan Yan, Mia Glaese, Mianna Chen, Michael Lampe, Michael Malek, Michele Wang, Michelle Fradin, Mike McClay, Mikhail Pavlov, Miles Wang, Mingxuan Wang, Mira Murati, Mo Bavarian, Mostafa Rohaninejad, Nat McAleese, Neil Chowdhury, Neil Chowdhury, Nick Ryder, Nikolas Tezak, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Patrick Chao, Paul Ashbourne, Pavel Izmailov, Peter Zhokhov, Rachel Dias, Rahul Arora, Randall Lin, Rapha Gontijo Lopes, Raz Gaon, Reah Miyara, Reimar Leike, Renny Hwang, Rhythm Garg, Robin Brown, Roshan James, Rui Shu, Ryan Cheu, Ryan Greene, Saachi Jain, Sam Altman, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Santiago Hernandez, Sasha Baker, Scott McKinney, Scottie Yan, Shengjia Zhao, Shengli Hu, Shibani Santurkar, Shraman Ray Chaudhuri, Shuyuan Zhang, Siyuan Fu, Spencer Papay, Steph Lin, Suchir Balaji, Suvansh Sanjeev, Szymon Sidor, Tal Broda, Aidan Clark, Tao Wang, Taylor Gordon, Ted Sanders, Tejal Patwardhan, Thibault Sottiaux, Thomas Degry, Thomas Dimson, Tianhao Zheng, Timur Garipov, Tom Stasi, Trapit Bansal, Trevor Creech, Troy Peterson, Tyna Eloundou, Valerie Qi, Vineet Kosaraju, Vinnie Monaco, Vitvityr Pong, Vlad Fomenko, Weiye Zheng, Wenda Zhou, Wes McCabe, Wojciech Zaremba, Yann Dubois, Yinghai Lu, Yiming Chen, Young Cha, Yu Bai, Yuchen He, Yuchen Zhang, Yunyun Wang, Zheng Shao, and Zhuohan Li. 2024. OpenAI o1 System Card. arXiv:2412.16720 [cs.AI] <https://arxiv.org/abs/2412.16720>
- [37] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. 2019. Habitat: A Platform for Embodied AI Research. arXiv:1904.01201 [cs.CV] <https://arxiv.org/abs/1904.01201>
- [38] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. DeepSeek-Math: Pushing the Limits of Mathematical Reasoning in Open Language Models. arXiv:2402.03300 [cs.CL] <https://arxiv.org/abs/2402.03300>
- [39] Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, Ruochen Xu, and Tiancheng Zhao. 2025. VLM-R1: A Stable and Generalizable R1-style Large Vision-Language Model. arXiv:2504.07615 [cs.CV] <https://arxiv.org/abs/2504.07615>
- [40] Lucy Xiaoyang Shi, Brian Ichter, Michael Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, Adrian Li-Bell, Danny Driess, Lachy Groom, Sergey Levine, and Chelsea Finn. 2025. Hi Robot: Open-Ended Instruction Following with Hierarchical Vision-Language-Action Models. arXiv:2502.19417 [cs.RO] <https://arxiv.org/abs/2502.19417>
- [41] Suyeon Shin, Sujin jeon, Junghyun Kim, Gi-Cheon Kang, and Byoung-Tak Zhang. 2025. Socratic Planner: Self-QA-Based Zero-Shot Planning for Embodied Instruction Following. arXiv:2404.15190 [cs.AI] <https://arxiv.org/abs/2404.15190>
- [42] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winsong Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. arXiv:1912.01734 [cs.CV] <https://arxiv.org/abs/1912.01734>
- [43] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M. Sadler, Wei-Lun Chao, and Yu Su. 2023. LLM-Planner: Few-Shot Grounded Planning for Embodied Agents with Large Language Models. arXiv:2212.04088 [cs.AI] <https://arxiv.org/abs/2212.04088>
- [44] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Keaneley, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhanshu Sharma, Abheesht Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Pettrini, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Paparas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szepesktor, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Pöder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry, Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Viktor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sansevier, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. 2025. Gemma 3 Technical Report. arXiv:2503.19786 [cs.CL] <https://arxiv.org/abs/2503.19786>
- [45] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chunming Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu, Shengling Li, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong, Weiran He, Weixiao Huang, Weixin Xu, Wenhao Wu, Wenyang He, Xianghui Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, Zonghan Yang, and Zongyu Lin. 2025. Kimi k1.5: Scaling Reinforcement Learning with LLMs. arXiv:2501.12599 [cs.AI] <https://arxiv.org/abs/2501.12599>
- [46] Songjun Tu, Jiahao Lin, Xiangyu Tian, Qichao Zhang, Linjing Li, Yuqian Fu, Nan Xu, Wei He, Xiangyuan Lan, Dongmei Jiang, and Dongbin Zhao. 2025. Enhancing LLM Reasoning with Iterative DPO: A Comprehensive Empirical Investigation. arXiv:2503.12854 [cs.CL] <https://arxiv.org/abs/2503.12854>
- [47] Songjun Tu, Jiahao Lin, Qichao Zhang, Xiangyu Tian, Linjing Li, Xiangyuan Lan, and Dongbin Zhao. 2025. Learning When to Think: Shaping Adaptive Reasoning in R1-Style Models via Multi-Stage RL. arXiv:2505.10832 [cs.CL] <https://arxiv.org/abs/2505.10832>
- [48] Songjun Tu, Jingbo Sun, Qichao Zhang, Xiangyuan Lan, and Dongbin Zhao. 2024. Online Preference-based Reinforcement Learning with Self-augmented Feedback from Large Language Model. arXiv:2412.16878 [cs.LG] <https://arxiv.org/abs/2412.16878>
- [49] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An Open-Ended Embodied Agent with Large Language Models. arXiv:2305.16291 [cs.AI] <https://arxiv.org/abs/2305.16291>

- abs/2305.16291
- [50] Hongcheng Wang, Yinuo Huang, Sukai Wang, Guanghui Ren, and Hao Dong. 2025. GRPO-MA: Multi-Answer Generation in GRPO for Stable and Efficient Chain-of-Thought Training. arXiv:2509.24494 [cs.CL] <https://arxiv.org/abs/2509.24494>
- [51] Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. 2024. Measuring Multimodal Mathematical Reasoning with MATH-Vision Dataset. arXiv:2402.14804
- [52] Ruoyao Wang, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. 2022. ScienceWorld: Is your Agent Smarter than a 5th Grader? arXiv:2203.07540 [cs.CL] <https://arxiv.org/abs/2203.07540>
- [53] Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, Lingming Zhang, Daniel Fried, Gabriel Synnaeve, Rishabh Singh, and Sida I. Wang. 2025. SWE-RL: Advancing LLM Reasoning via Reinforcement Learning on Open Software Evolution. arXiv:2502.18449 [cs.SE] <https://arxiv.org/abs/2502.18449>
- [54] Di Wu, Jiaxin Fan, Junzhe Zang, Guanbo Wang, Wei Yin, Wenhao Li, and Bo Jin. 2025. Reinforced Reasoning for Embodied Planning. arXiv:2505.22050 [cs.AI] <https://arxiv.org/abs/2505.22050>
- [55] Zhenyu Wu, Ziwei Wang, Xiuwei Xu, Jiwen Lu, and Haibin Yan. 2023. Embodied Task Planning with Large Language Models. arXiv:2307.01848 [cs.CV] <https://arxiv.org/abs/2307.01848>
- [56] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu, Xuanjing Huang, and Tao Gui. 2023. The Rise and Potential of Large Language Model Based Agents: A Survey. arXiv:2309.07864 [cs.AI] <https://arxiv.org/abs/2309.07864>
- [57] Kaixuan Xu, Jiajun Chai, Sicheng Li, Yuqian Fu, Yuanheng Zhu, and Dongbin Zhao. 2025. DipLLM: Fine-Tuning LLM for Strategic Decision-making in Diplomacy. arXiv:2506.09655 [cs.AI] <https://arxiv.org/abs/2506.09655>
- [58] Zhiyuan Xu, Kun Wu, Junjie Wen, Jinming Li, Ning Liu, Zhengping Che, and Jian Tang. 2024. A Survey on Robotics with Foundation Models: toward Embodied AI. arXiv:2402.02385 [cs.RO] <https://arxiv.org/abs/2402.02385>
- [59] Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, Heng Ji, Huan Zhang, and Tong Zhang. 2025. EmbodiedBench: Comprehensive Benchmarking Multi-modal Large Language Models for Vision-Driven Embodied Agents. arXiv:2502.09560 [cs.AI] <https://arxiv.org/abs/2502.09560>
- [60] Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. 2025. Demystifying Long Chain-of-Thought Reasoning in LLMs. arXiv:2502.03373 [cs.CL] <https://arxiv.org/abs/2502.03373>
- [61] Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. 2025. Robotic Control via Embodied Chain-of-Thought Reasoning. arXiv:2407.08693 [cs.RO] <https://arxiv.org/abs/2407.08693>
- [62] Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2023. AgentTuning: Enabling Generalized Agent Abilities for LLMs. arXiv:2310.12823
- [63] Huaye Zeng, Dongfu Jiang, Haozhe Wang, Ping Nie, Xiaotong Chen, and Wenhui Chen. 2025. ACECODER: Acing Coder RL via Automated Test-Case Synthesis. arXiv:2502.01718 [cs.SE] <https://arxiv.org/abs/2502.01718>
- [64] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. 2024. Vision-Language Models for Vision Tasks: A Survey. arXiv:2304.00685 [cs.CV] <https://arxiv.org/abs/2304.00685>
- [65] Jingyi Zhang, Jiaying Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. 2025. R1-VL: Learning to Reason with Multimodal Large Language Models via Step-wise Group Relative Policy Optimization. arXiv:2503.12937 [cs.AI] <https://arxiv.org/abs/2503.12937>
- [66] Wenqi Zhang, Mengna Wang, Gangao Liu, Xu Huixin, Yiwei Jiang, Yongliang Shen, Guiyang Hou, Zhe Zheng, Hang Zhang, Xin Li, Weiming Lu, Peng Li, and Yueting Zhuang. 2025. Embodied-Reasoner: Synergizing Visual Search, Reasoning, and Action for Embodied Interactive Tasks. arXiv:2503.21696 [cs.CL] <https://arxiv.org/abs/2503.21696>
- [67] Yaocheng Zhang, Haohuan Huang, Zijun Song, Yuanheng Zhu, Qichao Zhang, Zijie Zhao, and Dongbin Zhao. 2025. CriticSearch: Fine-Grained Credit Assignment for Search Agents via a Retrospective Critic. arXiv:2511.12159 [cs.CL] <https://arxiv.org/abs/2511.12159>
- [68] Zijiang Zhang, Ziyang Chen, Mingxiao Li, Zhaopeng Tu, and Xiaolong Li. 2025. RLVMR: Reinforcement Learning with Verifiable Meta-Reasoning Rewards for Robust Long-Horizon Agents. arXiv:2507.22844 [cs.LG] <https://arxiv.org/abs/2507.22844>
- [69] Zijie Zhao, Honglei Guo, Shengqian Chen, Kaixuan Xu, Bo Jiang, Yuanheng Zhu, and Dongbin Zhao. 2025. Empowering Multi-Robot Cooperation via Sequential World Models. arXiv:2509.13095 [cs.RO] <https://arxiv.org/abs/2509.13095>