

# DEPO: Dual-Efficiency Preference Optimization for LLM Agents

Sirui Chen<sup>1,2,3\*</sup>, Mengshi Zhao<sup>4\*</sup>, Lei Xu<sup>2,5</sup>, Yuying Zhao<sup>3</sup>,  
Beier Zhu<sup>3†</sup>, Hanwang Zhang<sup>3</sup>, Shengjie Zhao<sup>1†</sup>, Chaochao Lu<sup>2†</sup>

<sup>1</sup>Tongji University

<sup>2</sup>Shanghai Artificial Intelligence Laboratory

<sup>3</sup>Nanyang Technological University

<sup>4</sup>The University of Hong Kong

<sup>5</sup>École Polytechnique Fédérale de Lausanne (EPFL)

2111292@tongji.edu.cn, {beier.zhu, hanwangzhang}@ntu.edu.sg, luchaochao@pjlab.org.cn

## Abstract

Recent advances in large language models (LLMs) have greatly improved their reasoning and decision-making abilities when deployed as agents. Richer reasoning, however, often comes at the cost of longer chain of thought (CoT), hampering interaction efficiency in real-world scenarios. Nevertheless, there still lacks systematic definition of LLM agent efficiency, hindering targeted improvements. To this end, we introduce dual-efficiency, comprising (i) step-level efficiency, which minimizes tokens per step, and (ii) trajectory-level efficiency, which minimizes the number of steps to complete a task. Building on this definition, we propose DEPO, a dual-efficiency preference optimization method that jointly rewards succinct responses and fewer action steps. Experiments on WebShop and BabyAI show that DEPO cuts token usage by up to 60.9% and steps by up to 26.9%, while achieving up to a 29.3% improvement in performance. DEPO also generalizes to three out-of-domain math benchmarks and retains its efficiency gains when trained on only 25% of the data.

**Project page** — <https://opencausalab.github.io/DEPO>

## 1 Introduction

LLMs have shown impressive capabilities in complex tasks like long-context reasoning (Bai et al. 2024), coding (Zhuo et al. 2025), and mathematics (Luo et al. 2025b). To bring these abilities closer to real-world applications, recent studies have started investigating the LLM agents, aiming to build autonomous systems capable of task planning, tool use, and multi-turn interaction with dynamic environments (Shi et al. 2024; Song et al. 2024; Chen et al. 2024; Zeng et al. 2024; Yuan et al. 2025; Fu et al. 2025; Wang et al. 2025c; Feng et al. 2025a). As advanced models pursue greater reasoning accuracy, their CoT (Wei et al. 2022; Zhao et al. 2025) often becomes longer, which leads to increased response latency and diminished execution efficiency. Efficiency is crucial for transitioning LLM agents from prototypes to practical applications, as slow response times are a

\*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

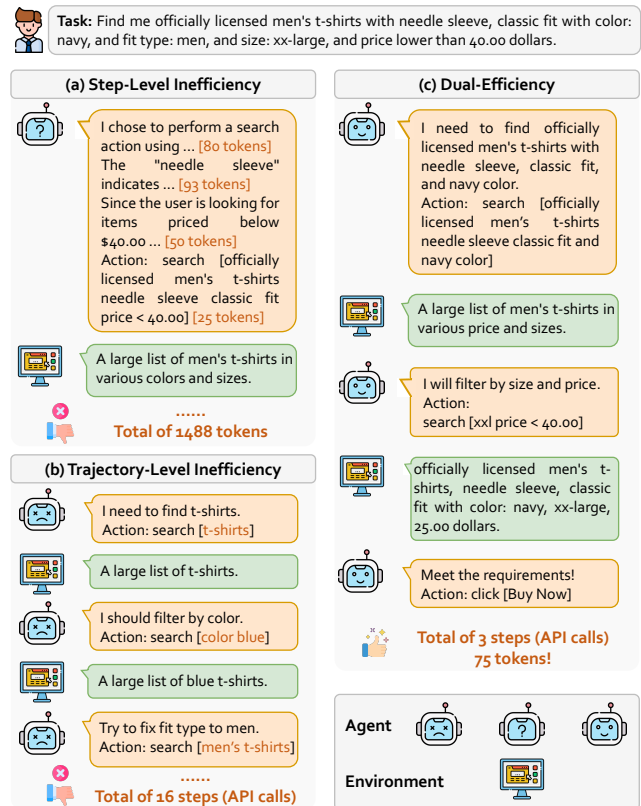


Figure 1: A comparison between (a) step-level inefficiency, arising from latency and cost in LLM token generation; (b) trajectory-level inefficiency, arising from latency and cost in environment interactions such as API calls, and our defined (c) dual-efficiency. For LLM agents, achieving genuine efficiency requires joint optimization across both dimensions.

major barrier to real-world adoption. However, speed cannot come at the expense of accuracy—a fast but unreliable agent is useless, or even harmful. We therefore seek response efficiency while preserving reasoning quality.

Current discussions on LLM efficiency predominantly fo-

cus on reducing the number of tokens generated (Wang et al. 2025a; Qu et al. 2025; Feng et al. 2025b). However, this perspective overlooks the interaction dynamics of LLM agents. In practice, agents often need to make API calls or access web services at each decision step, where latency and cost scale with the number of steps, not just tokens. Therefore, we first define LLM agent efficiency along two key dimensions, which we term *dual-efficiency*. (1) *Step-level efficiency*: Aims to minimize the number of tokens generated in each interaction step, promoting concise responses. (2) *Trajectory-level efficiency*: Aims to minimize the total number of interaction steps required to complete a task. Figure 1 presents examples comparing LLM agents with step-level and trajectory-level inefficiencies, alongside agents achieving dual-efficiency. A step-level inefficient agent generates too many tokens per step due to overthinking (Figure 1(a)). A trajectory-level inefficient agent is concise at each step but requires numerous steps to complete the task because of inaccurate reasoning (Figure 1(b)). In contrast, a dual-efficient agent keeps both tokens per step and total steps low, achieving fast and effective task completion (Figure 1(c)).

Existing reinforcement learning (RL) methods serve as a key tool for aligning LLMs with human preferences, with PPO as a representative approach (Schulman et al. 2017; Luo et al. 2025a; Shen et al. 2025; Aggarwal and Welleck 2025). Recent advances further improve training stability (e.g., DPO (Rafailov et al. 2023; Chen et al. 2025); GRPO (Shao et al. 2024)) and reduce computational costs (e.g., KTO (Ethayarajh et al. 2024)). However, these methods focus primarily on learning dynamics rather than agent efficiency—that is, how effectively the agent interacts with the environment. To fill this gap, we propose DEPO, a Dual-Efficiency Preference Optimization method that explicitly targets both step-level and trajectory-level efficiency in LLM agents. Concretely, we extend KTO with an efficiency bonus—a parameter-independent offset added to the desirable log-ratio—rewarding samples that use fewer steps and tokens. Then we apply the sigmoid preference loss with the Kullback–Leibler (KL) divergence (Kullback and Leibler 1951). The bonus encourages a larger decision margin for efficient trajectories, amplifying their preference score while leaving the loss formulation and the undesirable branch unchanged. Overall, DEPO relies solely on offline desirable and undesirable labels, without the need for paired annotations, reward model training, or on-policy sampling. The method is stable to train and both compute- and data-efficient.

We conduct comprehensive experiments using multiple LLM agents on five widely adopted benchmarks, Webshop, BabyAI, GSM8K, MATH, and SimulEq. The results show that DEPO significantly enhances agent efficiency, achieving up to a 60.9% reduction in token usage compared to models trained with behavioral cloning (BC), and reducing step count by up to 26.9% over vanilla KTO. Notably, DEPO enhances efficiency without sacrificing performance, and even achieves additional performance gains—the largest observed improvement being 29.3% relative to BC. DEPO also demonstrates strong generalizability, delivering efficiency improvements on three out-of-domain math benchmarks. Finally, DEPO exhibits excellent sample efficiency, maintain-

ing efficiency gains even when trained with only 25% of the training data.

To summarize, our main contributions are as follows:

- We formally define dual-efficiency for LLM agents in terms of both step-level and trajectory-level efficiency, providing a foundation for developing effective algorithms and evaluating their performance.
- We introduce DEPO, which incorporates dual-efficiency preferences into vanilla KTO, guiding models to generate responses that use fewer tokens and require fewer steps.
- We conduct comprehensive experiments to validate the effectiveness of DEPO, and demonstrate its strong generalizability and sample efficiency.

## 2 Related Work

**LLM agent.** As LLMs continue to improve in reasoning and planning capabilities, recent research begins to explore the potential of LLMs as agents and investigate ways to further enhance their abilities. Existing approaches can be broadly categorized into three groups: SFT, offline RL, and online RL. For SFT-based methods, AgentTuning (Zeng et al. 2024) uses a trajectory dataset and general-domain instructions to boost LLMs’ agent abilities without harming their general performance. Agent-FLAN (Chen et al. 2024) fine-tunes LLMs by decomposing and redesigning the agent training corpus. AgentRefine (Fu et al. 2025) uses a data-generation pipeline that simulates diverse environments to overcome overfitting in existing agent-tuning methods. For offline RL, ETO (Song et al. 2024) features an exploration–training loop: agents collect their failure trajectories, generate contrastive preference pairs for DPO training. DMPO (Shi et al. 2024) adapts DPO for multi-turn agent tasks by using a state-action occupancy measure constraint and adding length normalization to the Bradley–Terry model. Xiong et al. (2024) introduce the IPR framework, which employs Monte Carlo–estimated step-level rewards to construct contrastive action pairs for DPO training. As for online RL, Wang et al. (2025c) propose StarPO, a general reinforcement learning framework designed to optimize the full multi-turn interaction trajectories of LLM agents. Feng et al. (2025a) design GiGPO, an RL algorithm that uses a two-level grouping approach to estimate relative advantage. However, while these methods primarily focus on improving the performance of LLM agents, research on enhancing their efficiency remains largely unexplored.

**Efficient reasoning.** The phenomenon of overthinking in LLMs has sparked interest in investigating their efficient reasoning. In addition to controlling inference (Pan et al. 2024; Aytes, Baek, and Hwang 2025; Wang et al. 2025b; Li et al. 2025; Han et al. 2025) or constructing datasets for supervised fine-tuning (Liu et al. 2024b; Yu et al. 2024; Kang et al. 2025; Xia et al. 2025; Munkhbat et al. 2025), some studies address this problem using RL. The central idea of RL-based approaches is to incorporate a length penalty into the loss function, thereby training LLMs to produce more concise responses (Luo et al. 2025a; Team et al. 2025; Shen et al. 2025; Aggarwal and Welleck 2025; Arora and Zanette

2025). Most of these approaches focus on reducing the tokens generated in a single response, without addressing the inefficiency caused by redundant steps during dynamic interactions between LLM agents and environments.

### 3 Method

#### 3.1 Setup

The interaction between an LLM agent and an environment can be modeled as a partially observable Markov decision process (POMDP):  $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R})$  (Song et al. 2024; Yuan et al. 2025). Here,  $\mathcal{S}$  the state space;  $\mathcal{A}$  the action space;  $\mathcal{O}$  the observation space; the transition function  $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ ; and the reward function  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ . Given a task  $u$  and an LLM agent  $\pi_\theta$  parameterized by  $\theta$ , the historical trajectory  $\tau_t$  and the action  $a_t$  at time  $t$  are defined as:

$$\tau_t := (o_1, a_1, \dots, o_t), \quad a_t \sim \pi_\theta(\cdot | u, o_{1:t}, a_{1:t-1}). \quad (1)$$

The task terminates upon success or after exceeding the maximum number of rounds, at which point the environment returns a reward  $r(\tau) \in [0, 1]$ .

In this paper, our goal is to train an agent  $\pi_\theta$  that maximizes the expected reward  $r(\cdot)$ , which involves two objectives: (i) achieving high success rates across tasks, (ii) promoting **dual efficiency**, which comprises two desiderata — minimizing the number of tokens generated per step and the total number of steps required to complete a task. Efficiency is optimized only among successful trajectories, ensuring necessary long reasoning on difficult tasks remains intact.

To this end, we first generate a large set of trajectories using Monte Carlo Tree Search, and label each as desirable or undesirable based on task success and dual efficiency (Sec. 3.2). We then adopt a two-stage training procedure: supervised fine-tuning (SFT) on high-quality desirable trajectories (Sec. 3.3), followed by our dual-efficiency preference optimization (DEPO), which contrasts desirable and undesirable samples to improve generalization under distribution shifts (Sec. 3.4).

#### 3.2 Data Generation

To prepare training data, we follow a two-step process. First, we employ Monte Carlo Tree Search (MCTS) to generate a large set of trajectories. Then, we assign each trajectory as desirable or undesirable based on the corresponding reward.

**Step 1: MCTS.** Inspired by Wang et al. (2024); Fu et al. (2025); Zhang et al. (2025), we employ MCTS to generate ReAct-style data (Yao et al. 2023). Each step consists of a natural-language `Thought` that reasons over the observation, followed by a corresponding `Action`. MCTS consists of the following four components:

- 1) *Selection.* Starting from the root node, iteratively select child nodes along the trajectory using the upper confidence bounds (UCT) criterion (Kocsis and Szepesvári 2006).

$$\text{UCT}(s) = Q(s) + w \sqrt{\frac{\ln N(\text{Pa}(s))}{N(s)}}, \quad (2)$$

where  $Q(s)$  is the average reward of state  $s$ ,  $w$  controls the degree of exploration, with higher values encouraging the selection of less frequently visited nodes.  $N(\text{Pa}(s))$  refers to the visit count of the parent node of  $s$ , whereas  $N(s)$  represents the visit count of  $s$ .

- 2) *Expansion.* Upon reaching a non-fully expanded node, add a new child node for an available action.
- 3) *Rollout.* From the new node, perform a rollout where the model generates actions either until the task is completed or the maximum depth is reached.
- 4) *Backpropagation.* Propagate the final reward from the rollout back through the trajectory, updating the values of each node. Through this process, MCTS yields a rich collection of trajectories, which serve as the foundation for the subsequent data filtering and model training stages.

**Step 2: data labeling.** Based on the  $r(\tau)$ , we design the following protocol to distinguish whether a trajectory is desirable  $\tau \in \mathcal{D}$  or undesirable  $\tau \in \mathcal{U}$ :

$$\begin{cases} \tau \in \mathcal{D}, & \text{if } r(\tau) \geq \kappa_0 \\ \tau \in \mathcal{U}, & \text{if } \kappa_1 > r(\tau) \geq \kappa_2, \end{cases} \quad (3)$$

where  $1 > \kappa_0 > \kappa_1 > \kappa_2 > 0$  are thresholds used to filter trajectories according to  $r(\tau)$ . A trajectory is **desirable** if its reward is greater than or equal to the upper threshold  $\kappa_0$ , indicating strong task performance. Conversely, trajectories with rewards below the lower threshold  $\kappa_2$  are discarded due to poor quality and limited training value. To facilitate effective preference learning, we retain only **undesirable** samples whose rewards fall in the intermediate range  $[\kappa_2, \kappa_1]$ . By enforcing a margin  $(\kappa_0 - \kappa_1)$  between desirable trajectories ( $r(\tau) > \kappa_0$ ) and undesirable ones ( $\kappa_1 > r(\tau)$ ), we ensure a clearer quality separation, making it easier for the model to learn consistent preference signals.

To further improve data quality, we introduce a rephrasing model, which polishes the `Thought` component at each step using contextual information, while keeping the `Action` unchanged. Additionally, we ensure that desirable trajectories contain fewer tokens per step after rephrasing. We denote  $\mathcal{D}$  as the set of all desirable trajectories and  $\mathcal{U}$  the set of all undesirable trajectories. The concrete threshold and rephrasing settings are detailed in Section 4.1.

#### 3.3 Behavioral Cloning

Our pipeline starts with the standard Behavioral Cloning (BC) method for model fine-tuning, a technique that has been widely utilized in prior work (Zeng et al. 2024; Song et al. 2024). Specifically, we randomly select a subset  $\mathcal{D}_{\text{BC}}$  from our desirable set  $\mathcal{D}$  to serve as the expert demonstrations. Using this dataset, we train an LLM agent to learn a base policy  $\pi_{\text{BC}}$  by minimizing the loss function:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{\tau \sim \mathcal{D}_{\text{BC}}} [\log \pi_\theta(\tau | u)]. \quad (4)$$

The expectation is taken with respect to the empirical distribution of trajectories in the dataset  $\mathcal{D}_{\text{BC}}$ .

### 3.4 Dual-Efficiency Preference Optimization

Behavior cloning enables an agent to mimic expert behavior by directly imitating demonstrated thoughts and actions. However, this approach lacks negative signals, which results in brittleness under distribution shift. As outlined in the Section 1, our objective is to improve the agent’s *dual efficiency*: (1) minimizing the average number of tokens generated per step, and (2) reducing the total number of steps required to complete a task. To enhance both the efficiency and performance of the agent, we propose to post-train the agent with DEPO, an efficiency-aware extension of vanilla KTO that rewards concise thoughts and fewer steps. We start with a brief overview of vanilla KTO.

**Vanilla KTO.** Vanilla KTO builds on prospect theory (Tversky and Kahneman 1992), incorporating reference dependence, diminishing sensitivity, and loss aversion to steer policies toward human-preferred behavior. Formally, KTO has the following optimization objective:

$$\mathcal{L}_{\text{KTO}}(\theta) = \mathbb{E}_{\tau \sim \mathcal{D}, \mathcal{U}} [\lambda(\tau) - v(\tau)], \quad (5)$$

where the weighting function  $\lambda(\tau)$  is defined as

$$\lambda(\tau) = \begin{cases} \lambda_D, & \text{if } \tau \in \mathcal{D}, \\ \lambda_U, & \text{if } \tau \in \mathcal{U}, \end{cases} \quad (6)$$

where  $\lambda_D, \lambda_U$  are hyper-parameters reflecting the importance of desirable versus undesirable trajectories. The value function  $v(\tau)$  measures the desirability of a trajectory under the current policy, which consists of a implied reward  $r_\theta(\tau)$  and a regularization term  $z_0(\tau)$ :

$$v(\tau) = \begin{cases} \lambda(\tau) \cdot \sigma(\beta(r_\theta(\tau) - z_0(\tau))), & \text{if } \tau \in \mathcal{D}, \\ \lambda(\tau) \cdot \sigma(\beta(z_0(\tau) - r_\theta(\tau))), & \text{if } \tau \in \mathcal{U}, \end{cases} \quad (7)$$

where  $\sigma(\cdot)$  is the sigmoid function, and  $\beta$  is a temperature parameter controlling the sharpness of the sigmoid. The regularization term  $z_0(\tau)$  is the KL divergence between the current policy and the baseline policy over the trajectory:

$$z_0(\tau) = \text{KL}(\pi_\theta(\cdot | \tau_t) \| \pi_{\text{BC}}(\cdot | \tau_t)). \quad (8)$$

**DEPO.** Our DEPO extends vanilla KTO by incorporating an efficiency-aware bonus directly into the reward function  $r_\theta(\cdot)$ , encouraging trajectories that require fewer tokens per step and fewer steps overall. Specifically, our implied reward  $r_\theta(\tau)$  combines (i) how much the current policy  $\pi_\theta$  prefers the observed action, relative to a baseline policy  $\pi_{\text{BC}}$ , given the trajectory  $\tau_t$ , and (ii) an efficiency-aware bonus  $b(\tau)$  that explicitly favors trajectories with fewer steps and tokens:

$$r_\theta(\tau) = \log \frac{\pi_\theta(a_t | \tau_t)}{\pi_{\text{BC}}(a_t | \tau_t)} + b(\tau), \quad (9)$$

where  $b(\tau)$  is defined as:

$$b(\tau) = \begin{cases} \frac{\alpha_1}{\bar{T}_{\text{token}}(\tau)} + \frac{\alpha_2}{T_{\text{step}}(\tau)}, & \text{if } \tau \in \mathcal{D}, \\ 0, & \text{if } \tau \in \mathcal{U}, \end{cases} \quad (10)$$

where  $\bar{T}_{\text{token}}(\tau)$  denotes the average number of tokens per step in  $\tau$ ,  $T_{\text{step}}(\tau)$  is the total number of steps, and  $\alpha_1, \alpha_2 > 0$  are hyper-parameters. Larger values of  $\bar{T}_{\text{token}}(\tau)$

or  $T_{\text{step}}(\tau)$  reduce the efficiency bonus  $b(\tau)$ , thereby lowering the overall reward  $r_\theta(\tau)$  and discouraging inefficient trajectories during preference optimization.

Note that the bonus term  $b(\tau)$  is applied only to desirable trajectories  $\tau \in \mathcal{D}$ , and set to zero for undesirable ones  $\tau \in \mathcal{U}$ . This selective design simplifies implementation and avoids injecting unnecessary signal into low-quality samples. Such design is also supported by our empirical results (Section 4.5), which show that penalizing undesirable trajectories offers no additional performance benefit.

## 4 Experiment

### 4.1 Setup

**Baselines.** We consider a wide range of baselines, including Deepseek-V3 (Liu et al. 2024a), DeepSeek-R1-Distill-{Qwen-7B, Llama-8B} (Guo et al. 2025), Llama-3.1-8B-Instruct (Dubey et al. 2024), Llama-3.2-3B-Instruct (Meta 2024), Qwen2.5-{3B,7B,72B}-Instruct (Qwen et al. 2025). In terms of training paradigms, we also compare BC and vanilla KTO. Additionally, we include Token Budget (TB) (Han et al. 2025), which compresses output length by including a reasonable budget directly in the prompt.

**Datasets.** Following previous work (Xi et al. 2025; Fu et al. 2025; Yuan et al. 2025), we evaluate our DEPO on five diverse datasets: Webshop (Yao et al. 2022), BabyAI (Chevalier-Boisvert et al. 2019), GSM8K (Cobbe et al. 2021), MATH (Hendrycks et al. 2021), and SimulEq (Kushman et al. 2014). Webshop is an interactive environment that simulates online shopping, where the agent must follow natural language instructions and perform actions such as clicking or searching to locate and purchase the desired item. In contrast, BabyAI is a classic grid-world environment focused on embodied interaction, designed to assess an agent’s ability to comprehend and carry out basic commands such as “go to the red ball” or “pick up the key”. GSM8K, MATH, and SimulEq are three widely used math benchmarks.

**Prompts.** We use the ReAct-style 0-shot prompt (Yao et al. 2023), directing the model to respond in the “Thought: ...; Action: ...” format.

**Metrics.** We evaluate the agents from two perspectives: efficiency and performance. For efficiency, we consider token count and step count. Token count is measured with the `cl100k_base` encoding from tiktoken (OpenAI 2025b), while step count refers to the number of responses the model generates to complete an entire trajectory. For performance, we report both success rate and reward. For Webshop and BabyAI, we define task success based on environment feedback on the final step. The reward measures the agent’s average reward obtained across episodes (Xi et al. 2025).

**Implementation details.** For MCTS, we set the maximum search depth to 50 and employ Deepseek-V3 for the search process. For data filtering, in BabyAI, trajectories with  $0.9 \leq r(\tau)$  are assigned to the desirable set  $\mathcal{D}$ , while those with  $0.7 \leq r(\tau) < 0.9$  are included in the undesirable set  $\mathcal{U}$ . In Webshop, trajectories with  $r(\tau) = 1$  are considered desirable, and the criteria for undesirable trajectories follow the same range as in BabyAI. Next, we filter the data based on the number of steps: trajectories with steps  $< 7$  are included

Method	Webshop						BabyAI					
	Succ. ( $\uparrow$ )	Re. ( $\uparrow$ )	T@All ( $\downarrow$ )	S@All ( $\downarrow$ )	T@Succ. ( $\downarrow$ )	S@Succ. ( $\downarrow$ )	Succ. ( $\uparrow$ )	Re. ( $\uparrow$ )	T@All ( $\downarrow$ )	S@All ( $\downarrow$ )	T@Succ. ( $\downarrow$ )	S@Succ. ( $\downarrow$ )
DeepSeek-V3	0.11	0.27	6769	36.19	439	6.8	0.61	0.57	2235	24.53	489	7.69
R1-Distill-Qwen-7B	0.00	0.00	6639	3.09	0	0.00	0.09	0.08	5719	15.30	317	3.75
R1-Distill-Llama-8B	0.00	0.00	7360	5.32	0	0.00	0.02	0.02	6047	19.60	784	6.00
Llama-3.1-8B-Instruct	0.03	0.31	501	13.57	232	8.00	0.78	0.74	604	16.89	304	8.13
Llama-3.2-3B-Instruct	0.00	0.00	831	17.06	0	0.00	0.68	0.62	571	19.89	219	7.56
Qwen2.5-3B-Instruct	0.00	0.02	1483	18.80	0	0.00	0.40	0.37	2187	23.78	487	7.92
Qwen2.5-7B-Instruct	0.02	0.13	2094	17.82	372	5.50	0.57	0.52	1271	20.09	431	7.49
Qwen2.5-14B-Instruct	0.09	0.19	990	12.89	319	6.25	0.64	0.51	1026	19.26	449	8.33
Qwen2.5-72B-Instruct	0.06	0.08	1622	17.51	361	8.80	0.39	0.36	2295	26.72	667	8.57
Llama-3.1-8B-BC	0.47	<b>0.79</b>	840	6.38	812	6.50	0.77	0.71	836	13.09	365	5.61
+ TB	0.42	0.76	791	<b>6.32</b>	793	<b>6.34</b>	0.82	0.70	783	12.21	399	6.30
+ KTO	0.48	0.67	776	8.96	567	6.88	0.87	0.81	342	10.14	<b>145</b>	4.35
+ <b>DEPO</b>	<b>0.50</b>	0.72	<b>633</b>	7.80	<b>500</b>	6.83	<b>0.88</b>	<b>0.82</b>	<b>327</b>	<b>9.32</b>	<b>145</b>	<b>4.09</b>
Qwen2.5-7B-BC	0.44	0.75	1014	<b>6.34</b>	828	<b>6.25</b>	0.47	0.43	2062	22.51	641	5.98
+ TB	0.42	0.73	803	6.52	803	6.29	0.59	0.53	1780	19.60	589	6.96
+ KTO	0.54	0.76	886	8.18	<b>645</b>	6.88	0.58	0.57	1199	22.81	<b>350</b>	<b>5.67</b>
+ <b>DEPO</b>	<b>0.56</b>	<b>0.80</b>	<b>726</b>	7.73	669	7.35	<b>0.75</b>	<b>0.69</b>	<b>893</b>	<b>16.67</b>	464	7.03

Table 1: Comparison of DEPO with a wide range of baselines. “-BC” denotes models cold-started via behavioral cloning. Metrics use the following abbreviations: Succ. (success rate), Re. (mean trajectory reward), T@All (mean tokens per trajectory over all trajectories), S@All (mean steps per trajectory over all trajectories), and T@Succ., S@Succ. (computed only over successful trajectories). Higher is better for Succ. and Reward; lower is better for T and S. “R1-Distill” denotes DeepSeek-R1-Distill models. Models with Succ. below 0.2 are excluded from the final comparison, as assessing efficiency under such low Succ. is not meaningful. Best results in each group are in bold.

in set  $\mathcal{D}$ , while those with steps  $\geq 7$  are assigned to set  $\mathcal{U}$ . Finally, rephrasing is performed using GPT-4.1 mini (OpenAI 2025a) with a temperature setting of 0.7. For BC, we use 972 trajectories from BabyAI and 3,732 from Webshop. Fine-tuning is performed using LoRA (Hu et al. 2022) with a learning rate of  $1.0e-4$  for 3 epochs. For dual-efficiency preference optimization, the BabyAI contains 512 desirable and 471 undesirable trajectories, while Webshop provides 1,567 for each category. We set  $\beta$  to 0.2,  $\lambda_{\mathcal{D}}$  and  $\lambda_{\mathcal{U}}$  to 1, with a learning rate of  $2.0e-5$  and 3 epochs. For Llama3.1-8B-BC, both  $\alpha_1$  and  $\alpha_2$  are set to 3; for Qwen2.5-7B-BC, both are set to 2. All experiments are conducted with eight NVIDIA Tesla A800 GPUs, each with 80GB of memory.

## 4.2 Main Results

Table 1 presents a detailed comparison between DEPO and various baselines. We draw the following conclusions:

**DEPO achieves a significant improvement in the efficiency of LLM agents.** In terms of token count, DEPO achieves the lowest token usage across both the Webshop and BabyAI datasets. Compared to vanilla KTO, Llama3.1-8B-BC+DEPO reduces T@All, T@Succ. on Webshop and T@All on BabyAI by 18.4%, 11.8%, and 4.4% respectively. When compared to Llama3.1-8B-BC, the reductions in T@All and T@Succ. on both datasets are even more pronounced, with decreases of 24.6%, 38.4%, 60.9%,

and 60.3%. Qwen2.5-7B-BC+DEPO achieves reductions of 18.1% and 25.5% in T@All on Webshop and BabyAI compared to vanilla KTO. When compared to Qwen2.5-7B-BC, T@All and T@Succ. decrease by 28.4%, 19.2%, 56.7%, and 27.6% across the two datasets. Regarding step count, compared to vanilla KTO, Llama3.1-8B-BC+DEPO reduces S@All and S@Succ. by 12.9% and 0.7% on Webshop, and by 8.1% and 6.0% on BabyAI. As for Qwen2.5-7B-BC+DEPO, S@All decreases by 5.5% on Webshop and by 26.9% on BabyAI.

**DEPO not only ensures efficiency but also maintains—or even improves—the performance of LLM agents.** Compared with KTO and TB, DEPO offers a better performance/efficiency trade-off: it reduces T@All/T@Succ. and S@All/S@Succ. by similar or larger margins while delivering higher Succ. and Reward. Qwen2.5-7B-BC+DEPO achieves the highest Succ. and Reward on Webshop, while Llama3.1-8B-BC+DEPO attains the best results on BabyAI. Additionally, Llama3.1-8B-BC+DEPO improves Succ. and Reward on Webshop by 4.2% and 7.5% over vanilla KTO, whereas Qwen2.5-7B-BC+DEPO boosts these metrics on BabyAI by 29.3% and 21.1%.

**The primary advantage of DEPO on Webshop lies in its substantial reduction of tokens per step.** This trend holds for both Llama3.1-8B-BC+DEPO and Qwen2.5-7B-BC+DEPO: although the number of steps increases




Figure 2: Model generalizability across math benchmarks. The left y-axis shows accuracy and the right y-axis shows average tokens. Qwen refers to Qwen2.5-7B-Instruct, Qwen+ to Qwen2.5-7B-BC+DEPO, Llama to Llama3.1-8B-Instruct, Llama+ to Llama3.1-8B-BC+DEPO.

slightly, the average tokens per step (i.e., T@All/S@All and T@Succ./S@Succ.) decrease significantly, leading to a notable reduction in total tokens. For Llama3.1-8B-BC+DEPO, T@All/S@All drops by 6.4% and T@Succ./S@Succ. by 11.2%. For Qwen2.5-7B-BC+DEPO, the reductions are 13.3% and 2.8%. All values are relative to vanilla KTO.

### 4.3 Generalizability

To evaluate the generalizability and performance retention of DEPO, we compare models trained with DEPO to their original Instruct versions on GSM8K, MATH, and SimulEq. Figure 2 presents the results, showing Accuracy (left axis, higher is better) and Tokens (right axis, lower is better) for each model. We find that:

**DEPO demonstrates strong generalizability, achieving higher average accuracy while reducing generation tokens on most tasks.** DEPO consistently improves the accuracy of both models on SimulEq as well as on the average across all three datasets. In terms of efficiency, Llama3.1-8B-BC+DEPO achieves a clear reduction in token usage, while Qwen2.5-7B-BC+DEPO shows a slight increase on average; for the MATH dataset, both models generally generate fewer tokens.

### 4.4 Sample Efficiency

We evaluate the sample efficiency of our method by training with 25%, 50%, 75%, and 100% of the full dataset, and compare both agent’s performance and efficiency across these settings. We report the average results of Llama3.1-8B-BC and Qwen2.5-7B-BC on BabyAI, as shown in Figure 3. Our key findings are as follows:

**DEPO exhibits excellent sample efficiency.** With just 25% of the training data (only 245 samples for BabyAI and 783 for Webshop), it delivers substantial improvements over the BC baseline. In particular, T@All increases efficiency by more than 10%. This demonstrates that DEPO is able to rapidly acquire preferences and boost both performance and efficiency, even when training data is highly limited.

**With larger training sets, DEPO consistently yields further gains.** All four metrics improve as the amount of data increases, and T@All, in particular, achieves nearly a 60% improvement with the full dataset. These results suggest that




Figure 3: Sample efficiency of DEPO. Relative improvement  $\Delta\%$  over the BC baseline for four metrics as the training-set proportion increases from 25% to 100%.

DEPO is able to effectively utilize more data for continued improvement, and is less prone to early saturation.

### 4.5 Ablation Studies

**Effect of hyper-parameters  $\alpha_1$  and  $\alpha_2$ .** As defined by Eq. (10),  $\alpha_1$  and  $\alpha_2$  are parameters that adjust step-level and trajectory-level efficiency, respectively. In Table 2, we compare the impact of different values of these parameters on both performance and efficiency. We observe that:

**Jointly considering both  $\alpha_1$  and  $\alpha_2$  achieves the best overall trade-off between performance and efficiency.** For Llama3.1-8B-BC, setting  $\alpha_1 = 3$  and  $\alpha_2 = 3$  on Webshop yields the highest Succ. and Reward, along with the best T@All and S@All. On BabyAI, this configuration matches the performance of  $\alpha_1 = 3$ ,  $\alpha_2 = 0$  but demonstrates superior efficiency across multiple metrics. For Qwen2.5-7B-BC,  $\alpha_1 = 2$  and  $\alpha_2 = 2$  achieves the best Succ., Reward, and S@All on Webshop, with T@All being the second best. On BabyAI, it attains the highest scores across Succ., Reward, T@All, and S@All.

**Optimizing a single parameter can further improve individual efficiency metrics, but this often comes at the expense of performance.** When emphasizing step-level efficiency ( $\alpha_1 > 0$ ,  $\alpha_2 = 0$ ), Llama3.1-8B-BC achieves the

Parameter	Webshop						BabyAI					
	Succ. (↑)	Re. (↑)	T@All (↓)	S@All (↓)	T@Succ. (↓)	S@Succ. (↓)	Succ. (↑)	Re. (↑)	T@All (↓)	S@All (↓)	T@Succ. (↓)	S@Succ. (↓)
Llama3.1-8B-BC												
$\alpha_1 = 3, \alpha_2 = 3$	<b>0.50</b>	<b>0.72</b>	<b>633</b>	<b>7.80</b>	500	6.83	<b>0.88</b>	<b>0.82</b>	327	<b>9.32</b>	<b>145</b>	<b>4.09</b>
$\alpha_1 = 3, \alpha_2 = 0$	0.43	0.66	818	8.86	<b>459</b>	<b>6.51</b>	<b>0.88</b>	<b>0.82</b>	<b>326</b>	9.71	160	4.47
$\alpha_1 = 0, \alpha_2 = 3$	0.43	0.67	849	9.27	549	6.62	0.84	0.78	458	12.48	233	6.01
Qwen2.5-7B-BC												
$\alpha_1 = 2, \alpha_2 = 2$	<b>0.56</b>	<b>0.80</b>	726	<b>7.73</b>	669	7.35	<b>0.75</b>	<b>0.69</b>	<b>893</b>	<b>16.67</b>	464	7.03
$\alpha_1 = 2, \alpha_2 = 0$	0.49	0.72	757	8.00	670	7.30	0.66	0.61	1149	19.68	410	<b>5.81</b>
$\alpha_1 = 0, \alpha_2 = 2$	0.51	0.74	<b>703</b>	7.98	<b>657</b>	<b>7.00</b>	0.63	0.59	1059	21.26	<b>391</b>	6.58

Table 2: The impact of different  $\alpha_1, \alpha_2$  values on model performance and efficiency. The best result for each model is highlighted in bold, with comparisons made within each model group rather than across models. The results show that jointly considering both  $\alpha_1$  and  $\alpha_2$  achieves the best trade-off between performance and efficiency.

lowest T@Succ. on Webshop and the lowest T@All on BabyAI. However, its performance is only comparable to, or slightly lower than, the balanced setting. Conversely, when focusing on trajectory-level efficiency ( $\alpha_1 = 0, \alpha_2 > 0$ ), Qwen2.5-7B-BC attains the lowest S@Succ. on Webshop, but both Succ. and Reward drop noticeably.

**Undesirable penalty.** As defined in Section 3.4, we only apply the bonus to desirable samples. A natural extension is to ask whether we can also impose a penalty on undesirable samples. In this way, penalty can be formulated as follows:

$$p(\tau) = \begin{cases} \frac{\alpha_1}{T_{\text{token}}(\tau)} + \frac{\alpha_2}{T_{\text{step}}(\tau)}, & \text{if } \tau \in \mathcal{U}, \\ 0, & \text{if } \tau \in \mathcal{D}, \end{cases} \quad (11)$$

We can then similarly add the penalty offset back to the undesirable trajectory. We apply equally strong bonuses and penalties, with  $\alpha_1 = 3, \alpha_2 = 3$  for Llama3.1-8B-BC and  $\alpha_1 = 2, \alpha_2 = 2$  for Qwen2.5-7B-BC. Figure 4 shows a comparison between DEPO and DEPO + penalty. We find that: **applying an additional penalty to undesirable trajectories with the same intensity does not yield overall improvements.** In terms of performance, both models exhibit noticeable declines on Webshop and BabyAI when the penalty is applied. Regarding efficiency, T@All and S@All generally rise markedly, with the most pronounced increases observed for Llama3.1-8B-BC on BabyAI (T@All +46.5%, S@All +39.4%). Although Qwen2.5-7B-BC exhibits slight reductions in T@Succ. and S@Succ., these gains are insufficient to counterbalance the broader negative impact.

## 5 Conclusion

In this work, we introduce dual-efficiency for LLM agents, encompassing both step-level efficiency (fewer tokens per turn) and trajectory-level efficiency (fewer steps per task). Based on this definition, we design DEPO, a preference-based optimization method that incorporates an efficiency bonus into the desirable log-ratio, rewarding samples that achieve tasks with fewer steps and tokens. Extensive experiments demonstrate that DEPO effectively improves agent




Figure 4: Desirable bonus only (DEPO) vs. Desirable bonus (DEPO) + Undesirable penalty. Each point shows the relative change of DEPO + undesirable penalty against the DEPO, with  $\Delta\% = (\text{Penalty} - \text{DEPO})/\text{DEPO} \times 100$ .

efficiency while enhancing task performance. Moreover, DEPO exhibits strong generalizability across domains and maintains high sample efficiency.

## Acknowledgments

We thank all the anonymous reviewers and area chair for their valuable feedback throughout the review process. This work is supported in part by the National Natural Science Foundation of China under Grant 625B2131, in part by Shanghai Artificial Intelligence Laboratory and in part by the National Key R&D Program of China 2023YFC3806000.

## References

- Aggarwal, P.; and Welleck, S. 2025. L1: Controlling how long a reasoning model thinks with reinforcement learning. *arXiv preprint arXiv:2503.04697*.
- Arora, D.; and Zanette, A. 2025. Training language models to reason efficiently. *arXiv preprint arXiv:2502.04463*.
- Aytes, S. A.; Baek, J.; and Hwang, S. J. 2025. Sketch-of-thought: Efficient llm reasoning with adaptive cognitive-inspired sketching. *arXiv preprint arXiv:2503.05179*.
- Bai, Y.; Lv, X.; Zhang, J.; Lyu, H.; Tang, J.; Huang, Z.; Du, Z.; Liu, X.; Zeng, A.; Hou, L.; Dong, Y.; Tang, J.; and Li, J. 2024. LongBench: A Bilingual, Multitask Benchmark for Long Context Understanding. In *ACL*.
- Chen, X.; Xu, J.; Liang, T.; He, Z.; Pang, J.; Yu, D.; Song, L.; Liu, Q.; Zhou, M.; Zhang, Z.; Wang, R.; Tu, Z.; Mi, H.; and Yu, D. 2025. Do NOT Think That Much for 2+3=? On the Overthinking of Long Reasoning Models. In *ICML*.
- Chen, Z.; Liu, K.; Wang, Q.; Zhang, W.; Liu, J.; Lin, D.; Chen, K.; and Zhao, F. 2024. Agent-FLAN: Designing Data and Methods of Effective Agent Tuning for Large Language Models. In *Findings of ACL*.
- Chevalier-Boisvert, M.; Bahdanau, D.; Lahlou, S.; Willems, L.; Saharia, C.; Nguyen, T. H.; and Bengio, Y. 2019. BabyAI: First Steps Towards Grounded Language Learning With a Human In the Loop. In *ICLR*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv e-prints*.
- Ethayarajh, K.; Xu, W.; Muennighoff, N.; Jurafsky, D.; and Kiela, D. 2024. Model Alignment as Prospect Theoretic Optimization. In *ICML*.
- Feng, L.; Xue, Z.; Liu, T.; and An, B. 2025a. Group-in-group policy optimization for llm agent training. *arXiv preprint arXiv:2505.10978*.
- Feng, S.; Fang, G.; Ma, X.; and Wang, X. 2025b. Efficient reasoning models: A survey. *arXiv preprint arXiv:2504.10903*.
- Fu, D.; He, K.; Wang, Y.; Hong, W.; GongQue, Z.; Zeng, W.; Wang, W.; Wang, J.; Cai, X.; and Xu, W. 2025. AgentRefine: Enhancing Agent Generalization through Refinement Tuning. In *ICLR*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Han, T.; Wang, Z.; Fang, C.; Zhao, S.; Ma, S.; and Chen, Z. 2025. Token-Budget-Aware LLM Reasoning. In *ACL*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. In *NeurIPS Datasets and Benchmarks Track*.
- Hu, E. J.; yelong shen; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*.
- Kang, Y.; Sun, X.; Chen, L.; and Zou, W. 2025. C3ot: Generating shorter chain-of-thought without compromising effectiveness. In *AAAI*.
- Kocsis, L.; and Szepesvári, C. 2006. Bandit based monte-carlo planning. In *European conference on machine learning*, 282–293. Springer.
- Kullback, S.; and Leibler, R. A. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86.
- Kushman, N.; Artzi, Y.; Zettlemoyer, L.; and Barzilay, R. 2014. Learning to Automatically Solve Algebra Word Problems. In *ACL*.
- Li, Y.; Hu, X.; Qu, X.; Li, L.; and Cheng, Y. 2025. Test-Time Preference Optimization: On-the-Fly Alignment via Iterative Textual Feedback. In *ICML*.
- Liu, A.; Feng, B.; Xue, B.; Wang, B.; Wu, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Liu, T.; Guo, Q.; Hu, X.; Jiayang, C.; Zhang, Y.; Qiu, X.; and Zhang, Z. 2024b. Can language models learn to skip steps? *NeurIPS*.
- Luo, H.; Shen, L.; He, H.; Wang, Y.; Liu, S.; Li, W.; Tan, N.; Cao, X.; and Tao, D. 2025a. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning. *arXiv preprint arXiv:2501.12570*.
- Luo, H.; Sun, Q.; Xu, C.; Zhao, P.; Lou, J.-G.; Tao, C.; Geng, X.; Lin, Q.; Chen, S.; Tang, Y.; and Zhang, D. 2025b. WizardMath: Empowering Mathematical Reasoning for Large Language Models via Reinforced Evol-Instruct. In *ICLR*.
- Meta. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. <https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct>.
- Munkhbat, T.; Ho, N.; Kim, S. H.; Yang, Y.; Kim, Y.; and Yun, S.-Y. 2025. Self-Training Elicits Concise Reasoning in Large Language Models. In *Findings of ACL*.
- OpenAI. 2025a. Introducing GPT-4.1 in the API. <https://openai.com/index/gpt-4-1/>.
- OpenAI. 2025b. Tiktoken. <https://github.com/openai/tiktoken>.
- Pan, J.; Zhang, Y.; Zhang, C.; Liu, Z.; Wang, H.; and Li, H. 2024. DynaThink: Fast or Slow? A Dynamic Decision-Making Framework for Large Language Models. In *EMNLP*.

- Qu, X.; Li, Y.; Su, Z.; Sun, W.; Yan, J.; Liu, D.; Cui, G.; Liu, D.; Liang, S.; He, J.; et al. 2025. A survey of efficient reasoning for large reasoning models: Language, multimodality, and beyond. *arXiv preprint arXiv:2503.21614*.
- Qwen; ; Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Tang, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2025. Qwen2.5 Technical Report.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shen, Y.; Zhang, J.; Huang, J.; Shi, S.; Zhang, W.; Yan, J.; Wang, N.; Wang, K.; Liu, Z.; and Lian, S. 2025. Dast: Difficulty-adaptive slow-thinking for large reasoning models. *arXiv preprint arXiv:2503.04472*.
- Shi, W.; Yuan, M.; Wu, J.; Wang, Q.; and Feng, F. 2024. Direct Multi-Turn Preference Optimization for Language Agents. In *EMNLP*.
- Song, Y.; Yin, D.; Yue, X.; Huang, J.; Li, S.; and Lin, B. Y. 2024. Trial and Error: Exploration-Based Trajectory Optimization of LLM Agents. In *ACL*.
- Team, K.; Du, A.; Gao, B.; Xing, B.; Jiang, C.; Chen, C.; Li, C.; Xiao, C.; Du, C.; Liao, C.; et al. 2025. Kimi k1.5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*.
- Tversky, A.; and Kahneman, D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4): 297–323.
- Wang, P.; Li, L.; Shao, Z.; Xu, R.; Dai, D.; Li, Y.; Chen, D.; Wu, Y.; and Sui, Z. 2024. Math-Shepherd: Verify and Reinforce LLMs Step-by-step without Human Annotations. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *ACL*.
- Wang, R.; Wang, H.; Xue, B.; Pang, J.; Liu, S.; Chen, Y.; Qiu, J.; Wong, D. F.; Ji, H.; and Wong, K.-F. 2025a. Harnessing the reasoning economy: A survey of efficient reasoning for large language models. *arXiv preprint arXiv:2503.24377*.
- Wang, X.; Feng, S.; Li, Y.; Yuan, P.; Zhang, Y.; Tan, C.; Pan, B.; Hu, Y.; and Li, K. 2025b. Make Every Penny Count: Difficulty-Adaptive Self-Consistency for Cost-Efficient Reasoning. In *Findings of NAACL*.
- Wang, Z.; Wang, K.; Wang, Q.; Zhang, P.; Li, L.; Yang, Z.; Jin, X.; Yu, K.; Nguyen, M. N.; Liu, L.; et al. 2025c. Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*.
- Xi, Z.; Ding, Y.; Chen, W.; Hong, B.; Guo, H.; Wang, J.; Guo, X.; Yang, D.; Liao, C.; He, W.; Gao, S.; Chen, L.; Zheng, R.; Zou, Y.; Gui, T.; Zhang, Q.; Qiu, X.; Huang, X.; Wu, Z.; and Jiang, Y.-G. 2025. AgentGym: Evaluating and Training Large Language Model-based Agents across Diverse Environments. In *ACL*.
- Xia, H.; Leong, C. T.; Wang, W.; Li, Y.; and Li, W. 2025. Tokenskip: Controllable chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*.
- Xiong, W.; Song, Y.; Zhao, X.; Wu, W.; Wang, X.; Wang, K.; Li, C.; Peng, W.; and Li, S. 2024. Watch Every Step! LLM Agent Learning via Iterative Step-level Process Refinement. In *EMNLP*.
- Yao, S.; Chen, H.; Yang, J.; and Narasimhan, K. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. In *NeurIPS*.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K. R.; and Cao, Y. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. In *ICLR*.
- Yu, P.; Xu, J.; Weston, J.; and Kulikov, I. 2024. Distilling system 2 into system 1. *arXiv preprint arXiv:2407.06023*.
- Yuan, S.; Chen, Z.; Xi, Z.; Ye, J.; Du, Z.; and Chen, J. 2025. Agent-R: Training Language Model Agents to Reflect via Iterative Self-Training. *arXiv preprint arXiv:2501.11425*.
- Zeng, A.; Liu, M.; Lu, R.; Wang, B.; Liu, X.; Dong, Y.; and Tang, J. 2024. AgentTuning: Enabling Generalized Agent Abilities for LLMs. In *Findings of ACL*.
- Zhang, Y.; Wang, T.; Chen, S.; Wang, K.; Zeng, X.; Lin, H.; Han, X.; Sun, L.; and Lu, C. 2025. ARise: Towards Knowledge-Augmented Reasoning via Risk-Adaptive Search. In *ACL*.
- Zhao, K.; Zhu, B.; Sun, Q.; and Zhang, H. 2025. Unsupervised visual chain-of-thought reasoning via preference optimization. *arXiv preprint arXiv:2504.18397*.
- Zhuo, T. Y.; Chien, V. M.; Chim, J.; Hu, H.; Yu, W.; Widyasari, R.; Yusuf, I. N. B.; Zhan, H.; He, J.; Paul, I.; Brunner, S.; GONG, C.; Hoang, J.; Zebaze, A. R.; Hong, X.; Li, W.-D.; Kaddour, J.; Xu, M.; Zhang, Z.; Yadav, P.; Jain, N.; Gu, A.; Cheng, Z.; Liu, J.; Liu, Q.; Wang, Z.; Lo, D.; Hui, B.; Muennighoff, N.; Fried, D.; Du, X.; de Vries, H.; and Werra, L. V. 2025. BigCodeBench: Benchmarking Code Generation with Diverse Function Calls and Complex Instructions. In *ICLR*.