

# Rethinking the Role of Entropy in Optimizing Tool-Use Behaviors for Large Language Model Agents

Zeping Li<sup>1</sup>, Hongru Wang<sup>2,†</sup>, Yiwon Zhao<sup>1</sup>, Guanhua Chen<sup>3</sup>, Yixia Li<sup>3</sup>, Keyang Chen<sup>1</sup>,  
Yixin Cao<sup>1</sup>, Guangnan Ye<sup>1</sup>, Hongfeng Chai<sup>1,†</sup>, Zhenfei Yin<sup>4,5,†</sup>

<sup>1</sup>Fudan University <sup>2</sup>University of Edinburgh <sup>3</sup>Southern University of Science and Technology  
<sup>4</sup>Oxford University <sup>5</sup>Haven <sup>†</sup>Corresponding authors

## Abstract

Tool-using agents based on Large Language Models (LLMs) excel in tasks such as mathematical reasoning and multi-hop question answering. However, in long trajectories, agents often trigger excessive and low-quality tool calls, increasing latency and degrading inference performance, making managing tool-use behavior challenging. In this work, we conduct entropy-based pilot experiments and observe a strong positive correlation between entropy reduction and high-quality tool calls. Building on this finding, we propose using entropy reduction as a supervisory signal and design two reward strategies to address the differing needs of optimizing tool-use behavior. Sparse outcome rewards provide coarse, trajectory-level guidance to improve efficiency, while dense process rewards offer fine-grained supervision to enhance performance. Experiments across diverse domains show that both reward designs improve tool-use behavior: the former reduces tool calls by 72.07% compared to the average of baselines, while the latter improves performance by 22.27%. These results position entropy reduction as a key mechanism for enhancing tool-use behavior, enabling agents to be more adaptive in real-world applications.

## 1 Introduction

Large language models are increasingly deployed as tool-using agents, where solving a task requires interleaving natural language reasoning with external actions such as code execution, web search, or API calls (Qian et al., 2025a; Li et al., 2025b; Jin et al., 2025a). This agentic paradigm has demonstrated strong performance on complex reasoning benchmarks, including mathematical problem solving (Cobbe et al., 2021), multi-hop question answering (Yang et al., 2018; Ho et al., 2020; Trivedi et al., 2022), and deep searching tasks (Mialon et al., 2023; Wu et al., 2025; Phan et al., 2025).

However, in long trajectories, agents may invoke tools excessively or inappropriately, increas-

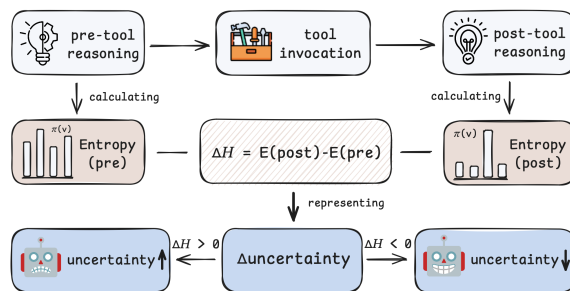


Figure 1: Changes in entropy reflect shifts in uncertainty within the agent. High-quality tool calls help the model reduce uncertainty, as indicated by a decrease in entropy.

ing computation cost and derailing the reasoning process (Qian et al., 2025b). Therefore, effective tool use poses a new challenge: an agent must not only produce a correct final answer, but also decide when and how to invoke tools throughout a long-horizon reasoning process (Zhang et al., 2025b), highlighting the importance of tool-invocation behavior itself (Wang et al., 2025a).

Most recent existing works adopt final-answer correctness as the outcome reward (Gao et al., 2025; Zhang et al., 2025a): while this supervision can teach models to use tools, it provides only endpoint feedback and thus fails to make the agent attend to tool-invocation behavior itself. In contrast, recent process-reward reinforcement learning (RL) methods (Feng et al., 2025; Zhang et al., 2025b), which focus on interactive environments like navigation tasks or object manipulation, deliver fine-grained supervision at each tool-use step, better guiding agents toward more accurate tool usage. Nonetheless, these process-reward designs typically rely on externally specified notions of “good” intermediate behaviors (e.g., hand-crafted rules, structured annotations), which are difficult to design for more complex long-term reasoning tasks. Consequently, a principled and lightweight signal for evaluating tool-call quality during long-horizon reasoning remains underexplored.

In this work, we revisit model uncertainty as a source of supervision for tool-augmented reasoning (Li et al., 2025a; Zhao et al., 2025; Cheng et al., 2025). Through entropy-based pilot experiments on tool-using models, we observe a consistent pattern across multiple domains: high-quality tool calls are generally followed by a reduction in entropy, whereas low-quality tool calls tend to increase it. This can be understood through cognitive load theory (Sweller, 2011): the quality of a tool call depends on whether it reduces unnecessary processing or introduces irrelevant information, thus affecting the model’s uncertainty. From this, we derive a general insight: the change in entropy before and after a tool call, quantified as delta segment entropy, serves as an intrinsic signal that reflects tool effectiveness, independent of task-specific annotations or external evaluators.

Building on this insight, we use delta segment entropy as a lightweight supervisory signal. Accordingly, we propose **Tool-enhanced Entropy-guided Policy Optimization (TEPO)**, introducing two reward designs to address different aspects of optimizing tool-use behavior. Specifically,  $\text{TEPO}_{\text{sparse}}$  focuses on improving tool-use efficiency by applying a sparse outcome reward that modulates the final-task reward based on the proportion of entropy-decreasing tool calls; whereas  $\text{TEPO}_{\text{dense}}$  emphasizes performance gains, using a dense process reward that grants an additional bonus whenever a tool call reduces entropy, thereby providing fine-grained supervision. Extensive experiments across domains demonstrate the effectiveness of both designs.  $\text{TEPO}_{\text{sparse}}$  improves efficiency, reducing tool calls by 72.07% with comparable results, whereas  $\text{TEPO}_{\text{dense}}$  focuses on boosting overall performance with an 22.27% improvement.

In summary, main contributions are as follows:

- We empirically observe that high-quality tool calls often induce entropy reduction, yielding a model-agnostic signal of tool-call quality across diverse reasoning domains.
- The proposed reward designs,  $\text{TEPO}_{\text{sparse}}$  and  $\text{TEPO}_{\text{dense}}$ , enhance tool behavior with distinct focuses on efficiency or performance, addressing the diverse needs across tasks.
- Comprehensive experiments demonstrate that entropy reduction is a crucial factor in optimizing tool-use behavior, empowering agents to adapt more effectively to real-world tasks.

## 2 Preliminary

### 2.1 Formalization of Agents’ Tool Use

We formalize tool-augmented generation as an iterative interaction between an agent  $A$  and a tool executor  $E$ . Given a query  $Q$ , the system state at step  $k$  is defined as

$$s_k = (r_1, a_1, o_1, \dots, r_k, a_k, o_k), \quad (1)$$

where  $r_i$  denotes the natural-language reasoning generated at step  $i$ ,  $a_i$  denotes the tool call issued at step  $i$ , and  $o_i$  denotes the observation returned by the executor  $E$  after executing  $a_i$ .

At each step, the agent conditions on  $Q$  and the interaction history  $s_{k-1}$  to produce the next reasoning segment and tool call, and the executor returns the corresponding observation. The interaction terminates when the agent emits a designated *finish* action, yielding the final answer.

### 2.2 Formalization of Delta Segment Entropy

At each decoding step  $t$ , the agent model (parameterized by  $\theta$ ) produces a hidden representation  $h_t$ , which yields a next-token distribution  $\pi_\theta(\cdot | h_t)$  over the vocabulary  $V$ . We define the token-level entropy as:

$$H(h_t) = - \sum_{v \in V} \pi_\theta(v | h_t) \log \pi_\theta(v | h_t), \quad (2)$$

which measures the uncertainty of the model’s next-token prediction at step  $t$ .

Given a reasoning segment  $r$  consisting of a contiguous sequence of tokens, we define the segment-level entropy as the token-entropy averaged over all tokens within the segment:

$$H(r) = \frac{1}{|I|} \sum_{t \in I} H(h_t), \quad (3)$$

where  $I$  denotes the index set of tokens belonging to the reasoning segment  $r$ . This formulation captures the overall uncertainty exhibited by the model throughout the reasoning process.

Let  $r_{k-1}$  denote the reasoning segment generated immediately before the  $k$ -th tool call, and let  $r_k$  denote the subsequent reasoning segment generated after receiving the tool observation  $o_k$ . We define the delta segment entropy  $\Delta H_k$  between these two segments as:

$$\Delta H_k = H(r_k) - H(r_{k-1}). \quad (4)$$

Entropy	Score	Math	Search	DeepS.	Avg
<i>Qwen3-8B-SFT</i>					
$\Delta H_k$	0	17.81	37.90	2.72	19.47
$\Delta H_k$	1	-37.27	-14.48	-23.07	<b>-24.94</b>
$\Delta H_k^{\text{ratio}}$	0	48.62	18.77	16.14	27.51
$\Delta H_k^{\text{ratio}}$	1	50.89	26.77	20.47	32.04
<i>Llama3.1-8B-SFT</i>					
$\Delta H_k$	0	-0.96	32.95	2.08	11.35
$\Delta H_k$	1	-41.03	1.38	-30.29	<b>-23.64</b>
$\Delta H_k^{\text{ratio}}$	0	47.84	19.92	21.25	29.34
$\Delta H_k^{\text{ratio}}$	1	42.98	26.54	25.72	31.41

Table 1: Entropy-based pilot experiment results: Average delta segment entropy ( $\Delta H_k$ ) and delta segment entropy ratio ( $\Delta H_k^{\text{ratio}}$ ) of SFT models, grouped by tool score (0/1) across the Math, Search, DeepSearch(DeepS.) domains. For readability,  $\Delta H_k$  values are scaled by a factor of  $10^3$ .

When  $\Delta H_k$  is negative, indicating an entropy reduction, we define the delta segment entropy ratio  $\Delta H_k^{\text{ratio}}$  to further quantify the extent of this reduction and enable comparison across different queries and reasoning contexts as:

$$\Delta H_k^{\text{ratio}} = \frac{H(r_k) - H(r_{k-1})}{H(r_{k-1}) + \epsilon}, \quad (5)$$

where  $\epsilon = 10^{-8}$  is a small constant introduced for numerical stability and to avoid division by zero.

### 2.3 Entropy-Based Pilot Experiments

To obtain an initial understanding of how entropy dynamics relate to tool usage, we conduct a set of entropy-based pilot experiments. Following the ARPO/AEPO setup (Dong et al., 2025c,a), we use the same supervised fine-tuning dataset to train two models, resulting in Qwen3-8B-SFT and Llama3.1-8B-SFT. These SFT-trained models already exhibit basic tool-use capabilities, providing a suitable starting point for our entropy analysis.

We evaluate the SFT models on three domains, following the same domain partition as in ARPO (Dong et al., 2025c) and AEPO (Dong et al., 2025a). To quantify the relationship between tool-call quality and entropy dynamics, we score each tool call with an LLM-as-judge and compute per-call entropy changes after the tool interaction. We then compare entropy statistics between high-quality and low-quality calls (the evaluation details can be found in Appendix A)

The results shown in Table 1 reveal a clear association across all three domains and both SFT models: high-quality tool calls (score 1) consistently

yield negative  $\Delta H_k$ , indicating reduced uncertainty in subsequent reasoning, whereas low-quality calls (score 0) often increase entropy. Moreover, in the Search and DeepSearch domains, score 1 calls exhibit higher  $\Delta H_k^{\text{ratio}}$  than score 0 calls. Overall, these findings suggest that entropy reduction serves as a useful, model-agnostic signal correlated with tool effectiveness.

## 3 Method

In this section, we first reformulate GRPO from a token-level perspective to clarify reward attribution to generated tokens (Section 3.1). This unified objective underlies our two reward designs: Section 3.2 applies a sparse outcome reward, while Section 3.3 introduces a dense process reward.

### 3.1 Token-Level GRPO Perspective

For each input question  $x$ , we sample a group of  $N$  rollouts  $\{y_i\}_{i=1}^N$  under the GRPO framework. Let  $y_{i,t}$  denote the token generated at position  $t$  in rollout  $i$ . The token-level GRPO objective is formulated as

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{D}} \left[ \frac{1}{N} \sum_{i=1}^N \sum_t \rho_{i,t}(\theta) A_{i,t} \right] - \beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}), \quad (6)$$

where  $A_{i,t}$  denotes the advantage assigned to the  $t$ -th token in rollout  $i$ , and the token-level importance ratio is defined as

$$\rho_{i,t}(\theta) = \frac{\pi_{\theta}(y_{i,t} \mid y_{i,<t}, x)}{\pi_{\theta_{\text{old}}}(y_{i,t} \mid y_{i,<t}, x)}. \quad (7)$$

### 3.2 Sparse Outcome Reward Design

Building on the findings from the entropy-based pilot experiments, we adopt a straightforward approach by incorporating the proportion of entropy-decreasing tool calls within a rollout into the reward design. This design encourages the model to either increase the number of entropy-decreasing tool calls or reduce the total number of tool calls to achieve a higher reward. Let  $n_i$  denote the total number of tool calls in rollout  $i$ , and let  $m_i$  denote the number of tool calls that induce entropy decrease. We adopt the final-answer F1-score as the verifiable correctness signal, consistent with ARPO (Dong et al., 2025c) and AEPO (Dong et al., 2025a). When a rollout makes no tool calls (i.e.,  $n_i = 0$ ), the outcome reward reduces to the F1-score only. Otherwise, we define the sparse outcome reward for rollout  $i$  as:

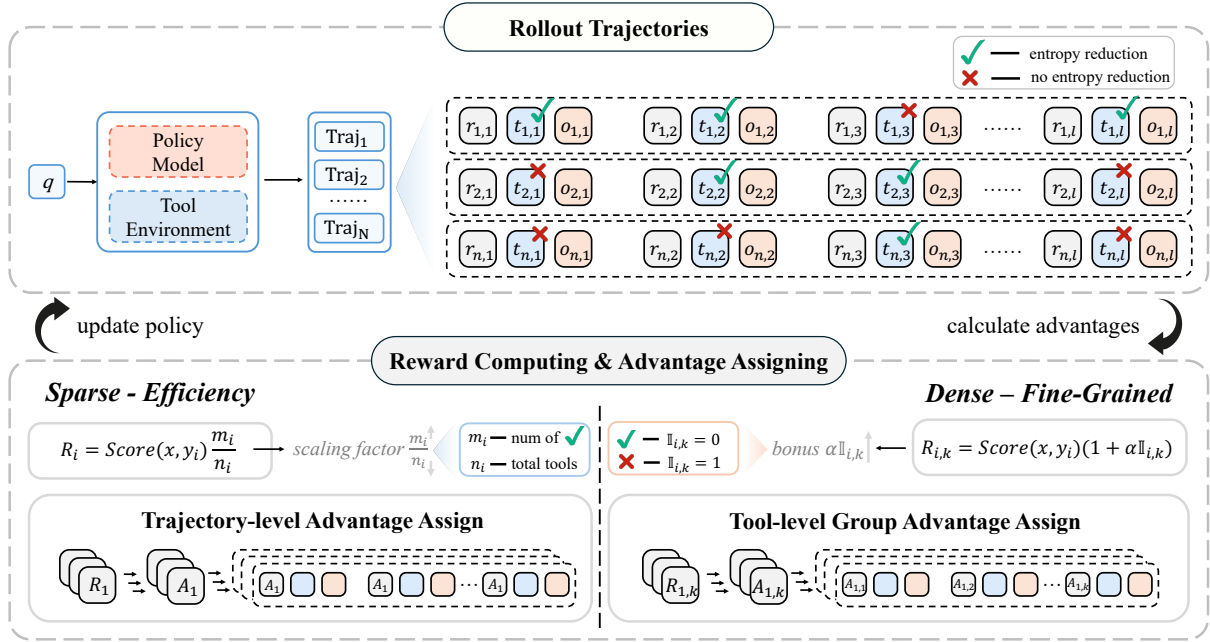


Figure 2: The overall framework of  $\text{TEPO}_{\text{sparse}}$  and  $\text{TEPO}_{\text{dense}}$ . In the sparse reward design, the reward and advantage are calculated and then uniformly assigned to each token within the trajectory (same  $A_{i,t}$  for all tokens). In contrast, the dense reward design assigns fine-grained tool rewards and advantages, resulting in different  $A_{i,t}$  values for different tokens within the same trajectory.

$$r_i^{\text{sparse}} = \text{F1}(x, y_i) \cdot \frac{m_i}{n_i}. \quad (8)$$

Following GRPO, the trajectory-level advantage can be calculated as:

$$A_i^{\text{traj}} = \frac{r_i^{\text{sparse}} - \mu^{\text{sparse}}}{\sigma^{\text{sparse}} + \epsilon}. \quad (9)$$

As a sparse outcome reward design, we assign the resulting trajectory-level advantage uniformly to all tokens in the rollout. This yields a token-wise assignment  $A_{i,t} = A_i^{\text{traj}}$  for every token position  $t$  in rollout  $i$ , providing a stable but coarse-grained credit signal over the entire trajectory.

### 3.3 Dense Process Reward Design

While outcome rewards provide an effective learning signal in reinforcement learning, incorporating process rewards can offer denser supervision and better guide the optimization process. However, obtaining reliable process-level feedback is often challenging for tool-augmented generation. Motivated by the entropy-based pilot experiments, we leverage entropy-decrease signals as a lightweight yet informative proxy for process reward. For the  $k$ -th tool call in rollout  $i$ , the tool-level reward can be calculated as:

$$r_{i,k}^{\text{tool}} = \text{F1}(x, y_i) \cdot (1 + \alpha \mathbb{I}_{i,k}), \quad (10)$$

where  $\mathbb{I}_{i,k} \in \{0, 1\}$  is an entropy-decrease indicator, which equals 1 if the  $k$ -th tool call in rollout  $i$  lowers entropy, and 0 otherwise. This formulation anchors each tool call to task correctness while explicitly rewarding those that reduce uncertainty.

Since the tools used across the  $N$  rollouts of the same question are often similar in type and function, we group these tool rewards and compute tool-level advantages as:

$$A_{i,k}^{\text{tool}} = \frac{r_{i,k}^{\text{tool}} - \mu_{\mathcal{R}(x)}}{\sigma_{\mathcal{R}(x)} + \epsilon}. \quad (11)$$

Here,  $\mathcal{R}(x)$  denotes the collection of all tool rewards for question  $x$ :

$$\mathcal{R}(x) = \left\{ r_{i,k}^{\text{tool}} \mid i \in [1, N], k \in [1, n_i] \right\}. \quad (12)$$

We assign token-wise advantages by propagating each tool-level advantage to the reasoning segment before the tool call:

$$A_{i,t} = A_{i,k}^{\text{tool}}, \quad \forall t \in I_{i,k}^{\text{pre}}, \quad (13)$$

where  $I_{i,k}^{\text{pre}}$  denotes the token indices of the reasoning segment before the  $k$ -th tool in rollout  $i$ .

As a result, different token spans within the same trajectory receive distinct advantages, providing targeted signals that guide the agent to internalize high-quality tool usage and optimize its behavior.

Table 2: Overall performance on reasoning domain tasks. The best results are in **bold**, and the second-best results are underlined.  $\overline{\text{acc}}$  represents the average accuracy, and  $\overline{\text{tools}}$  represents the average number of tool calls, both calculated across all five datasets.

Model	aime24	aime25	HotPotQA	2wiki	musique	$\overline{\text{acc}}$	$\overline{\text{tools}}$
<b>Qwen2.5-7B-Instruct</b>	8.67 ± 3.40	13.33 ± 4.71	47.70 ± 3.31	38.50 ± 2.81	19.20 ± 1.47	25.48	1.98
+ SFT	<u>24.67 ± 3.40</u>	15.33 ± 4.00	59.60 ± 1.59	56.20 ± 0.81	30.20 ± 1.81	37.20	3.98
+ GRPO	23.33 ± 2.98	20.00 ± 2.11	<u>60.30 ± 0.81</u>	<b>59.30 ± 0.68</b>	<u>31.70 ± 1.33</u>	<u>38.93</u>	1.97
+ ARPO	23.33 ± 2.72	<u>24.44 ± 1.57</u>	56.83 ± 1.55	55.17 ± 0.85	29.83 ± 1.65	37.92	<u>1.67</u>
+ AEPO	23.33 ± 4.71	22.67 ± 6.80	58.30 ± 1.12	54.20 ± 1.63	31.00 ± 1.67	37.90	2.96
+ TEPO <sub>sparse</sub>	20.00 ± 2.98	21.33 ± 1.63	60.10 ± 1.02	<u>57.50 ± 1.22</u>	29.90 ± 1.36	37.77	<b>1.64</b>
+ TEPO <sub>dense</sub>	<b>30.67 ± 4.42</b>	<b>30.00 ± 3.65</b>	<b>62.10 ± 1.77</b>	57.00 ± 1.00	<b>32.20 ± 1.78</b>	<b>42.39</b>	3.99
<b>Llama3.1-8B-Instruct</b>	4.67 ± 2.67	2.00 ± 1.63	46.70 ± 0.68	25.60 ± 2.87	17.30 ± 1.44	19.25	3.45
+ SFT	8.67 ± 2.67	7.33 ± 3.27	59.40 ± 1.32	54.30 ± 0.98	29.00 ± 1.05	31.74	5.42
+ GRPO	13.33 ± 3.65	11.33 ± 1.63	62.70 ± 0.68	<b>62.50 ± 0.45</b>	<u>32.10 ± 0.80</u>	36.39	2.02
+ ARPO	<b>20.00 ± 2.98</b>	<u>13.33 ± 3.65</u>	<u>63.20 ± 0.87</u>	61.10 ± 1.66	29.20 ± 1.91	<u>37.37</u>	<u>1.79</u>
+ AEPO	-	-	-	-	-	-	-
+ TEPO <sub>sparse</sub>	<u>18.67 ± 4.00</u>	12.00 ± 1.63	60.60 ± 0.58	60.70 ± 1.63	30.20 ± 0.68	36.43	<b>1.66</b>
+ TEPO <sub>dense</sub>	<u>18.67 ± 3.40</u>	<b>16.67 ± 0.00</b>	<b>64.60 ± 1.16</b>	<u>61.70 ± 0.93</u>	<b>33.40 ± 2.84</b>	<b>39.01</b>	3.60

## 4 Experiment

### 4.1 Datasets

To evaluate the effectiveness of our algorithm for training LLM-based tool-using agents, we follow the domain partition adopted in ARPO (Dong et al., 2025c) and AEPO (Dong et al., 2025a), testing on three domains: mathematical reasoning, knowledge-intensive reasoning, and deep information searching. For mathematical reasoning, we use AIME2024 and AIME2025. For knowledge-intensive reasoning, we utilize HotpotQA (Yang et al., 2018), 2WikiMultihopQA (Ho et al., 2020), and Musique (Trivedi et al., 2022). For deep information searching, we employ GAIA (Mialon et al., 2023), WebWalker (Wu et al., 2025), and HLE (Phan et al., 2025). Detailed descriptions of the datasets can be found in the Appendix B.1.

### 4.2 Experimental Settings

We adopt a two-stage training paradigm, consisting of the SFT stage followed by the RL stage, identical to ARPO/AEPO. This approach not only stabilizes the early-stage optimization process (Dong et al., 2025b) but also ensures a fair comparison. All models undergo the same SFT training phase, and for the RL stage, reasoning tasks are primarily tested on Qwen2.5 and Llama3.1 models, while deep search tasks are tested on the Qwen3 model. Details of the SFT and RL training datasets, together with the training parameters and additional information, can be found in Appendix B.2.

Model	HPQA	2Wiki	Musi.	$\overline{\text{EM}}$	$\overline{\text{tools}}$
SearchR1-PPO	30.00	31.50	10.00	23.83	4.51
OTC-GRPO†	36.60	31.10	13.00	26.90	<b>1.05</b>
OTC-PPO†	38.30	36.30	15.20	29.93	1.76
StepSearch†	38.60	36.60	<b>22.60</b>	32.60	-
TEPO <sub>sparse</sub>	42.75	<b>54.12</b>	19.62	<b>38.83</b>	2.24
TEPO <sub>dense</sub>	<b>43.60</b>	51.90	20.60	38.70	5.77

Table 3: Results on five evaluation runs. † indicates that the results are directly cited from the paper. All baselines are based on Qwen2.5-7B-Instruct.  $\overline{\text{EM}}$  denotes the average exact match, included to ensure comparability with prior work.

During training and testing, we primarily use two external tools. For computation, we integrate a Python compiler in a sandbox environment that allows safe execution of generated code for complex computation. For knowledge retrieval, we adopt a search setup inspired by the Search-R1 (Jin et al., 2025a) framework, where the model generates search queries during reasoning and retrieves relevant information from a wiki-18 corpus.

### 4.3 Baselines

We compare against two groups of baselines. (i) **Algorithm-level baselines:** methods trained under the same pipeline as ours (same SFT stage, training data, and hyperparameters unless stated otherwise), so that differences can be attributed to the optimization objective and reward design. This group includes GRPO, a standard group-based pol-




Figure 3: Results on GAIA dataset evaluated with Qwen3-8B using Bing Search API as search tool, including avg@5, pass@5, and average tool calls.

icity optimization method, and two entropy-driven RL objectives, ARPO (Dong et al., 2025c) and AEPO (Dong et al., 2025a). (ii) **Recent TIR baselines:** recent RL methods that also use wiki-18 as the knowledge source but may adopt different training recipes. This group includes SearchR1 (Jin et al., 2025a), StepSearch (Wang et al., 2025c), and OTC-PO (Wang et al., 2025b). To account for variance, we run each evaluation five times and report the mean and standard deviation.

#### 4.4 Main Results

**Results on Reasoning Tasks.** The results are presented in Table 2. TEPO<sub>sparse</sub> yields a substantial improvement in tool-use efficiency across different models, reducing tool calls by 72.07% compared to the average of baselines while still showing comparable performance. This is expected because TEPO<sub>sparse</sub> employs an outcome-level reward in which the number of tool calls appears as a denominator term, thereby providing a global training signal that encourages the agent to achieve correct outcomes with fewer calls.

In contrast, TEPO<sub>dense</sub> outperforms all baselines in reasoning performance with an average increase of 22.27%, highlighting the advantage of its process-level reward: by assigning fine-grained credit to individual tool calls, it better shapes step-wise tool-use decisions.

Additionally, the results in Table 3 show that TEPO<sub>sparse</sub> and TEPO<sub>dense</sub> exceed the performance of several recent baselines. The OTC methods, which focus on tool-use efficiency, show higher efficiency than TEPO<sub>sparse</sub>, while the latter achieves a performance boost. Overall, these results demonstrate that our proposed entropy-guided algorithms are both effective and robust.

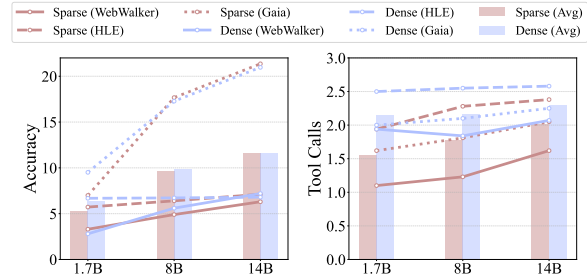


Figure 4: Results on Deep Search Tasks evaluated with different sizes of Qwen3 models using wiki-18 as search corpus, including avg@5 and average tool calls.

**Results on Deep Search Tasks.** The main experimental results are shown in Table 4. Additionally, we conducted extra experiments using the Bing Search API as a retrieval tool, which better supports deep search tasks. The results are shown in Figure 3, and the conclusions remain consistent: TEPO<sub>sparse</sub> demonstrates stronger tool-use efficiency, while TEPO<sub>dense</sub> exhibits superior reasoning performance.

## 5 Analysis

### 5.1 Scaling Analysis.

To further assess the scalability of our methods, we conducted an analysis across different model sizes, as shown in Figure 4. Our results show that both TEPO<sub>sparse</sub> and TEPO<sub>dense</sub> exhibit strong scalability. As the model size increases from 1.7B to 14B, performance consistently improves for both methods, following a clear scaling law. However, the scaling law is less pronounced on the HLE dataset, suggesting that the knowledge involved may exceed the capabilities of both the model and the wiki-18 knowledge base. This indicates the need for more powerful knowledge retrieval tools to further enhance performance. Notably, tool-use efficiency remains relatively stable across different scales, demonstrating the effectiveness in optimizing tool calls without compromising efficiency.

### 5.2 Training Efficiency Analysis

**Tool-Call Efficiency Analysis.** Figure 5 (a) shows the tool invocation curves during the training process. All three algorithms exhibit an initial increase in tool calls followed by a gradual decrease. Notably, TEPO<sub>sparse</sub> shows the fastest reduction in tool calls, which is consistent with the earlier finding that the TEPO<sub>sparse</sub> method demonstrates higher tool-call efficiency.

Table 4: Overall performance on deep search tasks. The best results are in **bold**, and the second-best results are underlined.  $\overline{\text{acc}}$  represents the average accuracy, and  $\overline{\text{tools}}$  represents the average number of tool calls, both calculated across all three datasets.

Model	webw.		hle		gaia		$\overline{\text{acc}}$	$\overline{\text{tools}}$
	acc	tools	acc	tools	acc	tools		
<b>Qwen3-8B</b>	1.00 ± 0.84	<b>1.67 ± 0.17</b>	5.76 ± 0.77	<u>0.34 ± 0.03</u>	12.04 ± 1.58	<u>1.41 ± 0.17</u>	6.27	<b>1.14</b>
+ SFT	5.30 ± 1.03	7.70 ± 0.34	6.36 ± 0.74	1.68 ± 0.06	15.53 ± 2.38	7.56 ± 0.39	9.06	5.65
+ GRPO	<u>5.80 ± 1.17</u>	2.87 ± 0.20	6.36 ± 0.73	0.57 ± 0.03	16.12 ± 2.50	1.85 ± 0.10	9.43	1.77
+ ARPO	5.40 ± 1.28	6.75 ± 0.26	6.16 ± 0.34	1.04 ± 0.07	16.89 ± 2.00	6.11 ± 0.28	9.48	4.63
+ AEPO	<b>6.00 ± 1.30</b>	5.06 ± 0.05	<b>6.76 ± 0.92</b>	0.92 ± 0.06	16.31 ± 0.95	3.64 ± 0.22	<u>9.69</u>	3.21
+ TEPO <sub>sparse</sub>	4.90 ± 0.73	<u>2.28 ± 0.20</u>	6.40 ± 0.42	<b>0.29 ± 0.02</b>	<b>17.67 ± 2.16</b>	<b>1.12 ± 0.22</b>	9.66	<u>1.23</u>
+ TEPO <sub>dense</sub>	5.60 ± 1.11	2.90 ± 0.16	<u>6.72 ± 0.48</u>	0.63 ± 0.02	<u>17.28 ± 1.13</u>	2.00 ± 0.09	<b>9.87</b>	1.84

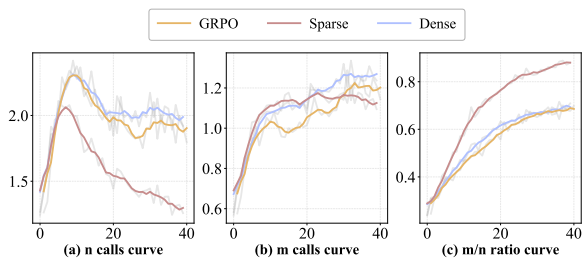


Figure 5: Visualization of training dynamics, showing (a)  $n$  curve, (b)  $m$  curve and (c)  $m/n$  ratio curve. Here,  $n$  is the total number of tool calls, and  $m$  is the number of tool calls inducing entropy reduction.

**Entropy-Reducing Tools Analysis.** We further investigated the relationship between the number of tool calls that induce entropy reduction across different algorithms during training. We define  $n$  as the total number of tool calls and  $m$  as the number of tool calls inducing entropy reduction, with the ratio  $m/n$  reflecting the proportion of entropy-reducing tool calls. The results are shown in Figure 5 (b), which presents the  $m$  calls curve, and Figure 5 (c), which shows the  $m/n$  ratio curve.

As can be seen from the figures, TEPO<sub>sparse</sub> exhibits a rapid increase in the  $m/n$  curve. This is largely driven by its sparse reward design, which uses  $m/n$  as a scaling factor: optimizing this ratio can be achieved either by increasing  $m$  or by reducing  $n$ . As a result, TEPO<sub>sparse</sub> improves tool-use efficiency while also raising the proportion of tool calls that induce entropy decrease. In contrast, TEPO<sub>dense</sub> shows a larger growth in  $m$ , which is directly attributable to its fine-grained reward design. By providing dense tool-level process rewards, TEPO<sub>dense</sub> explicitly encourages each tool call to be entropy-reducing, leading to more consistently effective tool use and ultimately better reasoning performance.

Score	Num.	Webwalker	HLE	Gaia	Avg
<i>Qwen3-8B-SFT</i>					
0	13002	16.33	-18.22	-5.90	-2.60
1	2122	-20.81	-27.72	-26.78	-25.10
<i>Qwen3-8B-Sparse</i>					
0	2814	-0.71	-65.23	-55.57	-40.50
1	551	-53.09	-71.85	-75.59	-66.84
<i>Qwen3-8B-Dense</i>					
0	4435	16.71	-39.75	-7.69	-10.24
1	1008	-17.34	-54.40	-34.74	-35.49

Table 5: Delta entropy experiment results: We report the average delta segment entropy ( $\Delta H_k$ ) of SFT, Sparse, and Dense models on deep search tasks across three datasets, aggregated by tool scores. For readability,  $\Delta H_k$  values are scaled by a factor of  $10^3$ .

### 5.3 Delta Entropy Analysis.

To further explore the impact of entropy reduction signals on tool usage, we examined the entropy changes induced by tool calls under different training methods. Delta entropy analysis was conducted on the three datasets from the deep search tasks, focusing on the SFT model as well as the TEPO<sub>sparse</sub> and TEPO<sub>dense</sub> models.

Specifically, we conducted five evaluation runs on these datasets, and following the experimental method outlined in section 2.3, we computed the average delta segment entropy ( $\Delta H_k$ ) for both high-quality and low-quality tool calls. The results are shown in Table 5.

After training with TEPO<sub>sparse</sub> and TEPO<sub>dense</sub>, both low-quality and high-quality tools exhibit a stronger tendency for entropy reduction compared to the SFT model, with TEPO<sub>sparse</sub> showing a more pronounced entropy reduction. However, the proportion of high-quality tools in TEPO<sub>sparse</sub> (16.4%) remains lower than that in TEPO<sub>dense</sub> (18.5%).

The above results provide complementary evidence for the effectiveness of both proposed algorithms from another perspective.  $\text{TEPO}_{\text{sparse}}$  is built upon an outcome-level reward and thus optimizes tool-use behavior at a global level, which manifests as fewer total tool calls and more entropy-reducing tool calls. On the other hand,  $\text{TEPO}_{\text{dense}}$  adopts a fine-grained reward design that encourages entropy-reducing tool calls at each step, also leading to a stronger entropy-reduction tendency than the SFT model. These findings suggest that entropy reduction can serve as an effective supervision signal for both  $\text{TEPO}_{\text{sparse}}$  and  $\text{TEPO}_{\text{dense}}$ .

## 6 Related Work

**Agentic RL for Tool-Integrated LLMs.** Tool-integrated LLM agents require policies that interleave reasoning with environment actions (e.g., search, API calls), shifting the learning problem from single-shot generation to long-horizon interaction (Yao et al., 2022; Schick et al., 2023; Qin et al., 2023; Chen et al., 2023; Gou et al., 2023). With reinforcement learning becoming a popular paradigm for training, recent work has utilized reinforcement learning to train models to use tools such as Python interpreters or search tools, enabling them to perform complex, multi-step, long-horizon reasoning tasks (Qian et al., 2025a; Gao et al., 2025; Li et al., 2025b; He et al., 2025; Zhang et al., 2025a; Wang et al., 2026).

However, in long trajectories, agents may invoke tools excessively or inappropriately, increasing computation cost and derailing the reasoning process (Qian et al., 2025b). Therefore, recent work has started to focus on the behavior of models when using tools (Wang et al., 2025a; Jin et al., 2025b). For example, Wang et al. (2025b) explores how to train models to learn the optimal number of tool calls, improving tool invocation efficiency. Similarly, Li et al. (2025c) introduces a reward signal based on tool-use completeness, which enhances the model’s ability to call tools effectively.

Although recent studies have explored process rewards, such as (Zhang et al., 2025b; Feng et al., 2025), how to provide step-level supervision signals for tool invocation behavior without a process reward model or handcrafted rules remains an open and valuable research question. The entropy reduction pattern serves as an effective lightweight supervision signal, enabling the model to learn what constitutes good tool-use behavior during training.

**Entropy-Based Signals for Agentic RL.** Entropy and information-theoretic metrics provide intrinsic uncertainty signals (Li et al., 2025a; Sharma and Chopra, 2025; Stoisser et al., 2025) that guide agent behavior. Yong et al. (2025) investigates step-wise information gain and adaptive termination for efficient reasoning. In agentic RL, Dong et al. (2025c) and Dong et al. (2025a) primarily leverage entropy for exploration and training stability, allocating more branching and sampling around high-uncertainty steps. Similarly, Cheng et al. (2025) demonstrates that higher entropy in reasoning correlates with exploratory behaviors, suggesting that entropy can drive deeper, more comprehensive reasoning chains.

While these studies use entropy to encourage exploration during high-uncertainty steps, this paper takes a different approach by directly rewarding tool calls that reduce entropy. Instead of promoting exploration, we reinforce behaviors that decrease uncertainty through tool use, with a reduction in entropy signaling positive information gain and improved reasoning performance.

## 7 Conclusion

In this work, we first delve into the relationship between tool usage and entropy. Through pilot experiments, we found that high-quality tool calls are often accompanied by entropy reduction. Building on this finding, we propose using entropy reduction as a supervisory signal and introduce two distinct reward strategies, each tailored to optimize tool-use behavior in different contexts.  $\text{TEPO}_{\text{sparse}}$  incorporates the proportion of entropy-reducing tool calls into final-answer correctness, reducing overall tool usage while improving efficiency and increasing the proportion of entropy-reducing tools for better reasoning performance. In contrast,  $\text{TEPO}_{\text{dense}}$  uses entropy reduction as a process-level supervision signal, guiding the model to recognize and optimize good tool usage behavior during training. Extensive experiments on several datasets validate the effectiveness of both methods.  $\text{TEPO}_{\text{sparse}}$  prioritizes efficiency, while  $\text{TEPO}_{\text{dense}}$  focuses on performance. This trade-off enables flexible tool usage strategies for different tasks. Our findings demonstrate that entropy reduction can serve as a powerful signal in reinforcement learning, paving the way for future research to refine entropy-based reward mechanisms and train more adaptive agents for complex environments.

## Limitations

Although the entropy reduction-based reward design proposed in this paper demonstrates promising results across various reasoning tasks, there are still some limitations. Firstly, while we validated the scalability from 1.7B to 14B models in our experiments, due to computational resource limitations, experiments on larger model sizes were not conducted. Future work could involve testing on larger models to further assess the performance of the proposed method across different model scales. Secondly, the primary experiments in this paper used the wiki-18 search tool, chosen for its stability and reproducibility in controlled settings. However, more challenging datasets and recent knowledge domains would benefit from the use of real-time search APIs, such as Bing or Google. While we conducted an experiment using the Bing Search API on the GAIA dataset, due to API cost and stability considerations, we did not test additional search APIs. In future work, we plan to expand the experimentation to include other real-time search APIs, ensuring broader validation and effectiveness in scenarios requiring up-to-date knowledge.

## Acknowledgement

This work was supported by grants from the National Natural Science Foundation of China(72595845).

## References

- Zhipeng Chen, Kun Zhou, Beichen Zhang, Zheng Gong, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Chatcot: Tool-augmented chain-of-thought reasoning on chat-based large language models. *arXiv preprint arXiv:2305.14323*.
- Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and Furu Wei. 2025. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Guanting Dong, Licheng Bao, Zhongyuan Wang, Kangzhi Zhao, Xiaoxi Li, Jiajie Jin, Jinghan Yang, Hangyu Mao, Fuzheng Zhang, Kun Gai, and 1 others. 2025a. Agentic entropy-balanced policy optimization. *arXiv preprint arXiv:2510.14545*.
- Guanting Dong, Yifei Chen, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Yutao Zhu, Hangyu Mao, Guorui Zhou, Zhicheng Dou, and Ji-Rong Wen. 2025b. Tool-star: Empowering llm-brained multi-tool reasoner via reinforcement learning. *arXiv preprint arXiv:2505.16410*.
- Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia Chen, Jiazhen Du, Huiyang Wang, Fuzheng Zhang, and 1 others. 2025c. Agentic reinforced policy optimization. *arXiv preprint arXiv:2507.19849*.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. Group-in-group policy optimization for llm agent training. *arXiv preprint arXiv:2505.10978*.
- Jiaxuan Gao, Wei Fu, Minyang Xie, Shusheng Xu, Chuyi He, Zhiyu Mei, Banghua Zhu, and Yi Wu. 2025. Beyond ten turns: Unlocking long-horizon agentic search with large-scale asynchronous rl. *arXiv preprint arXiv:2508.07976*.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen. 2023. Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint arXiv:2305.11738*.
- Guanzhong He, Zhen Yang, Jinxin Liu, Bin Xu, Lei Hou, and Juanzi Li. 2025. Webseer: Training deeper search agents through reinforcement learning with self-reflection. *arXiv preprint arXiv:2510.18798*.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps. *arXiv preprint arXiv:2011.01060*.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025a. [Search-rl: Training llms to reason and leverage search engines with reinforcement learning](#). *Preprint*, arXiv:2503.09516.
- Jiahe Jin, Abhijay Paladugu, and Chenyan Xiong. 2025b. Beneficial reasoning behaviors in agentic search and effective post-training to obtain them. *arXiv preprint arXiv:2510.06534*.
- Pengyi Li, Matvey Skripkin, Alexander Zubrey, Andrey Kuznetsov, and Ivan Oseledets. 2025a. Confidence is all you need: Few-shot rl fine-tuning of language models. *arXiv preprint arXiv:2506.06395*.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025b. Torl: Scaling tool-integrated rl. *arXiv preprint arXiv:2503.23383*.
- Zhiwei Li, Yong Hu, and Wenqing Wang. 2025c. Encouraging good processes without the need for good answers: Reinforcement learning for llm agent planning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1654–1666.

- Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2023. Gaia: a benchmark for general ai assistants. In *The Twelfth International Conference on Learning Representations*.
- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, and 1 others. 2025. Humanity’s last exam. *arXiv preprint arXiv:2501.14249*.
- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiushi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025a. Toolrl: Reward is all tool learning needs. *arXiv preprint arXiv:2504.13958*.
- Cheng Qian, Emre Can Acikgoz, Hongru Wang, Xiushi Chen, Avirup Sil, Dilek Hakkani-Tur, Gokhan Tur, and Heng Ji. 2025b. Smart: Self-aware agent for tool overuse mitigation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 4604–4621.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, and 1 others. 2023. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–68551.
- Aman Sharma and Paras Chopra. 2025. Think just enough: Sequence-level entropy as a confidence signal for llm reasoning. *arXiv preprint arXiv:2510.08146*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, pages 1279–1297.
- Josefa Lia Stoisser, Marc Boubnovski Martell, Lawrence Phillips, Gianluca Mazzoni, Lea Mørch Harder, Philip Torr, Jesper Ferkinghoff-Borg, Kaspar Martens, and Julien Fauqueur. 2025. Towards agents that know when they don’t know: Uncertainty as a control signal for structured reasoning. *arXiv preprint arXiv:2509.02401*.
- John Sweller. 2011. Cognitive load theory and e-learning. In *International Conference on Artificial Intelligence in Education*, pages 5–6. Springer.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. Musique: Multi-hop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554.
- Hongru Wang, Cheng Qian, Manling Li, Jiahao Qiu, Boyang Xue, Mengdi Wang, Heng Ji, and Kam-Fai Wong. 2025a. Toward a theory of agents as tool-use decision-makers. *arXiv preprint arXiv:2506.00886*.
- Hongru Wang, Cheng Qian, Wanjun Zhong, Xiushi Chen, Jiahao Qiu, Shijue Huang, Bowen Jin, Mengdi Wang, Kam-Fai Wong, and Heng Ji. 2025b. Otc: Optimal tool calls via reinforcement learning. *arXiv e-prints*, pages arXiv–2504.
- Tianyi Wang, Yixia Li, Long Li, Yibiao Chen, Shaohan Huang, Yun Chen, Peng Li, Yang Liu, and Guanhua Chen. 2026. Sppo: Sequence-level ppo for long-horizon reasoning tasks. *arXiv preprint arXiv:2604.08865*.
- Ziliang Wang, Xuhui Zheng, Kang An, Cijun Ouyang, Jialu Cai, Yuhang Wang, and Yichao Wu. 2025c. Stepsearch: Igniting llms search ability via stepwise proximal policy optimization. *arXiv preprint arXiv:2505.15107*.
- Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and 1 others. 2025. Web-walker: Benchmarking llms in web traversal. *arXiv preprint arXiv:2501.07572*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*.
- Xixian Yong, Xiao Zhou, Yingying Zhang, Jinlin Li, Yefeng Zheng, and Xian Wu. 2025. Think or not? exploring thinking efficiency in large reasoning models via an information-theoretic lens. *arXiv preprint arXiv:2505.18237*.
- Yabo Zhang, Yihan Zeng, Qingyun Li, Zhen Hu, Kavin Han, and Wangmeng Zuo. 2025a. Tool-rl: Sample-efficient reinforcement learning for agentic tool use. *arXiv preprint arXiv:2509.12867*.
- Zijing Zhang, Ziyang Chen, Mingxiao Li, Zhaopeng Tu, and Xiaolong Li. 2025b. Rlvmr: Reinforcement learning with verifiable meta-reasoning rewards for robust long-horizon agents. *arXiv preprint arXiv:2507.22844*.
- Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. 2025. Learning to reason without external rewards. *arXiv preprint arXiv:2505.19590*.

## A Details for Entropy-Based Pilot Experiments

### A.1 Domain Partition and Datasets

We follow the same domain partition as in ARPO (Dong et al., 2025c) and AEPO (Dong et al., 2025a). The Mathematical Reasoning domain includes AIME2024 and AIME2025, consisting of competition-level problems that require multi-step numerical and symbolic reasoning. The Knowledge-Intensive Reasoning domain includes open-domain multi-hop QA benchmarks that benefit from external search, including HotpotQA (Yang et al., 2018), 2WikiMultihopQA (Ho et al., 2020), and Musique (Trivedi et al., 2022). The Deep Information Searching domain includes web-search benchmarks that require iterative retrieval and long-horizon decision making, including GAIA (Mialon et al., 2023), WebWalker (Wu et al., 2025), and HLE (Phan et al., 2025).

### A.2 LLM-as-Judge for Tool-Call Quality

During evaluation, we adopt GPT-4o-mini as an LLM-as-judge to assess the quality of each tool call. For every tool call at step  $k$ , the judge is provided with: (i) the original question, (ii) the context preceding the call, (iii) the tool query, (iv) the tool result returned by the executor  $E$ , and (v) the subsequent response segment after observing the tool result. The judge assigns a binary tool score  $y_k \in \{0, 1\}$ , where  $y_k = 1$  indicates a high-quality tool call that provides relevant and useful information for solving the task, and  $y_k = 0$  indicates a low-quality call, such as malformed queries, failed executions, or irrelevant results. The full prompt used by the LLM-as-judge is provided in Appendix A.4.

### A.3 Evaluation Details and Entropy Statistics

For each domain and each SFT model, we run five independent evaluation trials. For each trial, we collect all tool calls made by the model and compute entropy-based statistics for each tool call.

**Per-call entropy change.** Following the entropy indicator defined in Sec. 2.2, we compute the delta segment entropy  $\Delta H_k$  for each tool call at step  $k$ , measuring the change of uncertainty in the subsequent response segment after the tool interaction. A negative  $\Delta H_k$  indicates entropy reduction (i.e., decreased uncertainty), while a positive value indicates increased uncertainty.

**Per-call entropy change ratio.** We also compute the delta segment entropy ratio  $\Delta H_k^{\text{ratio}}$  for each tool call that induces an entropy decrease, which normalizes the entropy change by the corresponding segment scale as defined in Sec. 2.2. This ratio enables comparisons of entropy dynamics across domains and models with different segment lengths and distributions.

### A.4 Judge Prompt

We include the prompt template used by GPT-4o-mini for assigning tool-call quality scores.

#### Utility Scoring Prompt

**You are an expert evaluator assessing the quality of tool calls in an AI agent’s problem-solving process.**

Original Question: {question}

Context previous to Tool Call: {previous\_context}

Search Query: {search\_query}

Tool Result: {tool\_result}

Agent Response subsequent to Tool Result: {subsequent\_response}

For the tool call, assign a score of 0 or 1. The criteria for evaluation are as follows:

- If the tool call itself fails, or if the search query derived from the previous context is unreasonable and cannot help answer the question, or if the tool result is irrelevant, unhelpful, or incorrect, assign a score of 0.
- If the tool call result provides help in reasoning about the problem or advances the reasoning process, assign a score of 1.

Respond with only the score (0 or 1) followed by a brief explanation classification: If the score is 0, give the classification in [unreasonable search query / search api failed / irrelevant or incorrect search result]. If the score is 1, no need to provide explanation.

Format: SCORE: <0 or 1>:  
<explanation classification>

Figure 6: LLM-as-Judge prompt for scoring tool quality.

Datasets	Num.	Domain
AIME2024	30	Math Reasoning
AIME2025	30	Math Reasoning
HotpotQA	200	Knowledge Intensive Reasoning
2wiki	200	Knowledge Intensive Reasoning
Musique	200	Knowledge Intensive Reasoning
GAIA	103	Deep Information Searching
WebWalker	200	Deep Information Searching
HLE	500	Deep Information Searching

Table 6: Details of datasets used in the main experiment.

## B Details for Experiment Settings

### B.1 Dataset Details

The datasets used in our experiments span three distinct domains, each focusing on different aspects of reasoning abilities. Mathematical Reasoning evaluates the model’s ability to solve competition-level problems that require multi-step numerical and symbolic reasoning. Knowledge-Intensive Reasoning tests the model’s capacity for multi-hop question answering, where external search is leveraged to retrieve relevant information. Finally, Deep Information Searching assesses the model’s performance in web-search tasks that require iterative retrieval and long-horizon decision-making, where the model needs to make decisions based on continuously evolving information.

Dataset sizes are shown in Table 6. Notably, in order to facilitate the evaluation, only the first 500 questions of the HLE dataset were selected. All tests involving the use of the HLE dataset in this paper were conducted with fair comparisons, where the results are based on the average of five evaluation runs on the first 500 questions.

### B.2 Implementation Details

**SFT Training Stage:** We initialize the tool-use behaviors via supervised fine-tuning using the LLaMAFactory framework. The SFT training corpus primarily sourced from the Tool-Star (Dong et al., 2025b) open-source dataset, along with some STILL dataset, totaling 54K data samples.

**RL Training Stage:** Building on the SFT training, we refine the model using reinforcement learning, selecting distinct RL training sets for different test domains, as in ARPO/AEPO. For reasoning tasks, we use 10K open-source RL samples from Tool-Star, primarily evaluating the performance on Qwen2.5 and Llama3.1 models. For deep search




Figure 7: Results on the GAIA dataset using the Bing search API. The reasons for each tool call labeled with a score of 0 are summarized. All values in the figure are presented as percentages.

tasks, we use a 1K RL training set, mainly assessing the model’s performance on Qwen3.

In our implementation, we adopt VERL (Sheng et al., 2025) framework to train the models. To stabilize reinforcement learning training, we exclude the tool-call results from loss computation to avoid bias, and we set the KL divergence coefficient to zero. We configure the rollout number to 8 to balance sample efficiency and training stability. During reinforcement learning, we train with a batch size of 128, a PPO minibatch size of 16, and a context window of 20K tokens. Training epochs are set to 2 and 5 separately for Reasoning tasks and Deep Search tasks. All experiments are conducted on four NVIDIA H200 GPUs. In the evaluation process, we adopt the VLLM framework with the inference parameters set as follows: top-p, temperature, and top-k are set to 0.95, 0.2, and 20, respectively. After extracting the answers, we use GPT-4o-mini to evaluate their correctness.

## C Additional Experimental Analyses

### C.1 Failed Cases Analysis

As shown in Figure 7, most low-quality tool calls were due to the inability to find relevant results. This suggests that while the model is trained to generate tool queries that aim to retrieve critical information for reasoning, these queries may not always be fully effective when interacting with external search APIs. This highlights the importance of future research on enhancing models’ ability to craft and refine search queries. Improving the query construction process can enable more precise interactions with external tools, helping models retrieve more relevant and accurate information.

Temp.	Score	WebWalker	HLE	GAIA	Avg
0.2	0	16.33	-18.22	-5.90	-2.60
0.2	1	-20.81	-27.72	-26.78	-25.10
0.7	0	26.10	-5.18	4.31	8.41
0.7	1	-30.86	-29.74	-34.62	-31.74
1.0	0	44.85	12.51	14.83	24.06
1.0	1	-44.80	5.96	-14.92	-17.92

Table 7: Entropy estimation stability under different decoding temperatures. We report the average entropy change for tool calls with different quality scores on WebWalker, HLE, and GAIA. The last column shows the average across the three benchmarks.

## C.2 Temperature Sensitivity Analysis

We further examine whether the observed entropy-quality relationship is sensitive to the decoding temperature used for entropy estimation. Since decoding temperature directly affects the sharpness of the output distribution, it can also influence the resulting entropy values and thus potentially act as a confounding factor in our analysis. To verify that our conclusion is not tied to a particular temperature choice, we conduct an additional analysis on Qwen3-8B under three decoding temperatures: a low temperature of 0.2, a medium temperature of 0.7, and a high temperature of 1.0. For each temperature, we compare tool calls with different quality scores and report the average entropy change on three search-intensive benchmarks: WebWalker, HLE, and GAIA.

The results are shown in Table 7. Across all three temperature settings, higher-quality tool calls remain consistently associated with entropy reduction, while lower-quality tool calls are generally associated with weaker reduction or even entropy increase. This trend holds under low, medium, and high temperatures, suggesting that the entropy-quality relationship is robust rather than tied to a particular temperature choice. These results provide additional evidence that the entropy signal captures a stable property of tool-use quality.

## C.3 Hyperparameter Sensitivity Analysis

We further study the hyperparameter sensitivity of  $\text{TEPO}_{\text{dense}}$ . Unlike  $\text{TEPO}_{\text{sparse}}$ , which does not introduce additional tunable hyperparameters,  $\text{TEPO}_{\text{dense}}$  includes a single parameter,  $\alpha$ , to control the fine-grained entropy shaping term. In the main experiments, we set  $\alpha$  to 0.1, whereas setting  $\alpha$  to 0 reduces  $\text{TEPO}_{\text{dense}}$  to GRPO. We addition-

$\alpha$	aime24	aime25	HPQA	2Wiki	Musi.	Avg
-0.1	23.33	16.67	59.80	55.70	28.30	36.76
0	23.33	20.00	60.30	59.30	31.70	38.93
0.1	30.67	30.00	62.10	57.00	32.20	42.39

Table 8: Sensitivity analysis of the entropy-shaping hyperparameter  $\alpha$  on Qwen2.5-7B-Instruct. The setting with  $\alpha$  set to 0 corresponds to GRPO, while  $\alpha$  set to -0.1 serves as a counterfactual negative perturbation.

ally evaluate a counterfactual setting in which  $\alpha$  is set to -0.1, which can be viewed as a “negative” perturbation of the entropy-based shaping signal.

The results are shown in Table 8. Overall, setting  $\alpha$  to 0.1 achieves the best average performance across the five reasoning benchmarks, improving the average score from 38.93 when  $\alpha$  is set to 0 to 42.39. In contrast, setting  $\alpha$  to -0.1 leads to a clear degradation, reducing the average score to 36.76. This trend suggests that the proposed entropy-shaping design is meaningful rather than arbitrary. Although the setting with  $\alpha$  at 0 performs slightly better on 2Wiki, the setting with  $\alpha$  at 0.1 yields the best results on four out of the five datasets and gives the strongest overall average, supporting our choice in the main experiments.