

CriticSearch: Fine-Grained Credit Assignment for Search Agents via a Retrospective Critic

Yaocheng Zhang^{1,2*} Haohuan Huang^{1,3*} Zijun Song^{1,2} Yuanheng Zhu^{1,3†}
Qichao Zhang^{1,3} Zijie Zhao^{1,3} Dongbin Zhao^{1,2,3}

¹Institute of Automation, CAS, Beijing, CHINA

²School of Advanced Interdisciplinary Sciences, UCAS, Beijing, CHINA

³School of Artificial Intelligence, UCAS, Beijing, CHINA

zhangyaocheng2023@ia.ac.cn

Abstract

Tool-Integrated Reasoning (TIR) with search engines enables large language models to iteratively retrieve up-to-date external knowledge, enhancing adaptability and generalization in complex question-answering tasks. However, existing search agent pipelines typically depend on reinforcement learning based optimization, which often suffers from sparse outcome rewards, leading to inefficient and unstable training. We introduce CriticSearch, a fine-grained credit-assignment framework that supplies dense, turn-level feedback via a retrospective critic mechanism. During training, a frozen, asymmetric critique LLM retrospectively evaluates each turn using privileged information from the full trajectory and gold answers, converting these assessments into stable, dense rewards that guide policy improvement. Experimental results across diverse multi-hop reasoning benchmarks demonstrate that CriticSearch consistently outperforms existing baselines, achieving faster convergence, improved training stability, and higher performance.

1 Introduction

Recently, multi-turn Tool-Integrated Reasoning (TIR) has achieved great potential across a wide range of downstream tasks, including code generation, mathematical reasoning, and question-answering (Q&A) (Shao et al., 2024; Jin et al., 2025; Feng et al., 2025a). This paradigm draws inspiration from the interaction and iterative mechanisms in reinforcement learning (RL) (Schulman et al., 2017; Shao et al., 2024), enabling models to invoke external tools (e.g., search engines, code interpreters), access up-to-date information, and progressively refine their reasoning process through feedback-driven iterations (Jin et al., 2025; Feng et al., 2025a). In particular, in complex

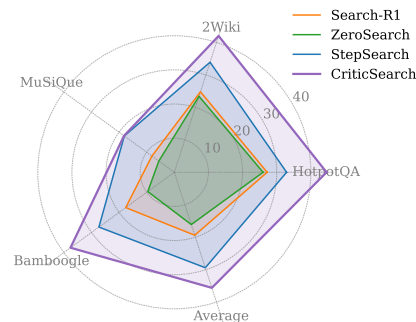


Figure 1: CriticSearch achieves leading performance on most of the datasets compared with RL methods.

Q&A and problem-solving scenarios, TIR with search engines (i.e., search agents) effectively mitigates inherent limitations of large language models (LLMs), such as outdated knowledge and limited contextual awareness, thereby achieving stronger performance and generalization.

However, the effectiveness of search agent hinges on RL-based optimization. Existing algorithms typically rely on sparse outcome rewards, where all search actions within a trajectory share the same reward regardless of quality (Jin et al., 2025; Feng et al., 2025a; Chen et al., 2025; Zheng et al., 2025b). Such coarse feedback prevents the model from distinguishing between effective and ineffective tool-calling, leading to inefficient training and slow policy improvement (Zheng et al., 2025a; Dong et al., 2026). Moreover, the reward sparsity induces high variance in returns, making training unstable (Plappert et al., 2018). These challenges become particularly severe in deep-search or other frequent tool-call settings, where a large number of externally generated tokens exacerbate instability, often resulting in reward collapse and gradient explosion (Tu et al., 2025a; Xue et al., 2025; Jin et al., 2025).

A natural solution to these challenges is to provide fine-grained, turn-level (action-level) feedback for search agent rather than relying solely

*Equal contribution.

†Corresponding author.

Method	Dense Reward	No additional Rollout	Generality
Search-R1 (Jin et al., 2025)	✗	✓	✓
StepSearch (Zheng et al., 2025a)	✓	✓	✗
ARPO (Dong et al., 2026)	✓	✗	✓
ReasonRAG (Zhang et al., 2025b)	✓	✗	✓
CriticSearch (Ours)	✓	✓	✓

Table 1: **Comparison of agentic RL with search engines.** CriticSearch satisfies three criteria: 1) dense reward, 2) no need for additional rollout during training, and 3) generality to all training data.

on sparse global rewards (Feng et al., 2025b; Dong et al., 2026; Zheng et al., 2025a). However, existing search agent paradigms struggle to perform effective credit assignment under sparse rewards (Tab. 1). Feng et al. (2025b); Dong et al. (2026) use Monte Carlo to approximate fine-grained rewards, which often suffer from high variance and require large amounts of rollout for stable estimation. Some studies have attempted to use ground-truth references for supervision, but such methods depend heavily on annotated data and lack generality (Zheng et al., 2025a).

To overcome these limitations, we propose *CriticSearch*—a fine-grained credit assignment method that leverages LLM to assist policy training. The key idea of CriticSearch is to introduce a *retrospective critic* mechanism: after the model produces a complete reasoning trajectory and obtains the gold answer, a critique model revisits each intermediate action to assess its contribution toward the final outcome and generate fine-grained turn-level rewards. Compared with forward-looking decision-making, this retrospective assessment leverages privileged information that is asymmetrically available during training (e.g., the gold answer and future interaction turns). Such hindsight access makes it easier and more reliable to generate turn-level rewards (Section 5.4) (Zhou et al., 2025a; Lee et al., 2020; Lowe et al., 2017). By integrating both turn-level and global rewards, CriticSearch assigns lower rewards to redundant or low-quality actions while amplifying the learning signals of key actions. With the aid of stable and denser feedback, CriticSearch effectively enhances training stability and learning efficiency.

Our main contributions are as follows:

- We propose a retrospective critic, a novel mechanism for search agents that leverages a frozen LLM to retrospectively evaluate trajectories with privileged hindsight, thereby producing dense and informative turn-level rewards.

- We provide valuable experimental analyses that reveal key phenomena arising from the introduction of dense rewards, including accelerated convergence and mitigation of training instability.
- CriticSearch surpasses dense-reward baselines by 16.7% and 6.7% on 3B and 7B models, respectively, across Q&A benchmarks.

Prompt 1: Training Prompt

Answer the given question. You must conduct reasoning inside `<think>` and `</think>` first every time you get new information. After reasoning, if you find you lack some knowledge, you can call a search engine by `<search>` query `</search>`, and it will return the top searched results between `<information>` and `</information>`. You can search as many times as you want. If you find no further external knowledge needed, you can directly provide the answer inside `<answer>` and `</answer>` without detailed illustrations. For example, `<answer>` Beijing `</answer>`. Question:

2 Preliminaries

2.1 Agentic RL with Search Engines

In a deep search task, the policy model is designed to interact with the search engine, producing a multi-turn reasoning trajectory $y = \{s_0, s_1, \dots, s_T\}$ to answer the question q . Here, T represents the total number of interaction turns. This process can be modeled as a partially observable Markov decision process (POMDP). As illustrated in Prompt 1, each interaction turn $s_t = (a_t, c_t)$ consists of two stages: a *decision stage* a_t and an *environment feedback stage* c_t (Jin et al., 2025). During the decision stage, the policy

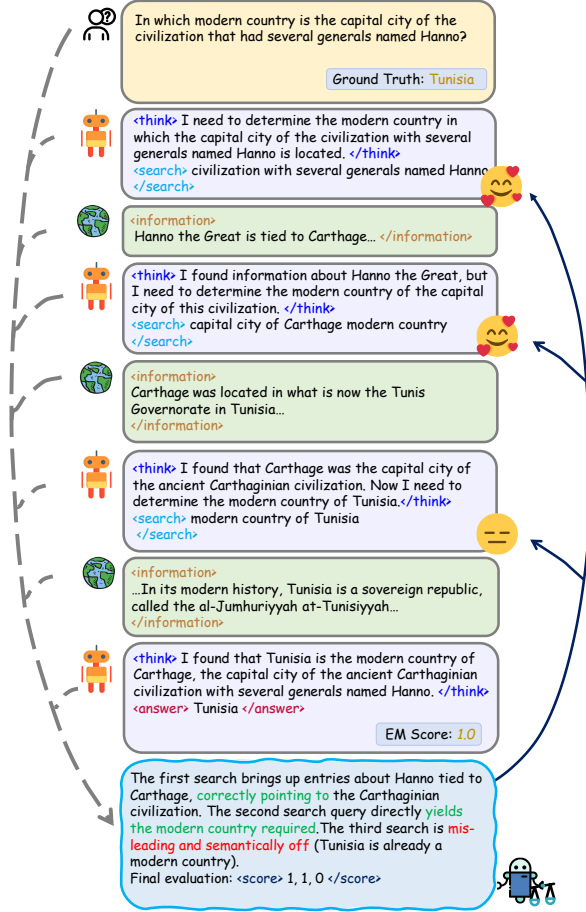


Figure 3: An example from CriticSearch illustrating the evaluation process of the Critique model. In this trajectory, the final answer is correct but the whole trajectory involves a potential mis-leading search action. A frozen critique LLM, after deliberation, provides accurate binary rewards for each turn.

rewards for advantage estimation in GRPO. The global reward, determined by the correctness of the final answer, measures the overall quality of the entire reasoning trajectory, while the turn-level rewards, assigned after each search action, assess the local quality of individual reasoning steps and search queries. The following section provides a detailed description of both reward schemes.

3.1 Global Reward Signal

Following prior work (Jin et al., 2025), in addition to employing the rule-based outcome reward function, we incorporate a format reward to construct the final reward function $r_\phi(q, y)$.

$$r_\phi = \begin{cases} 1 & \text{if } o_{\text{pred}} = o_{\text{gold}} \wedge f_{\text{format}}(y), \\ 1 - \lambda_f & \text{if } o_{\text{pred}} = o_{\text{gold}} \wedge \neg f_{\text{format}}(y), \\ \lambda_f & \text{if } o_{\text{pred}} \neq o_{\text{gold}} \wedge f_{\text{format}}(y), \\ 0 & \text{if } o_{\text{pred}} \neq o_{\text{gold}} \wedge \neg f_{\text{format}}(y), \end{cases} \quad (4)$$

where f_{format} verifies whether the response y follows the correct reasoning format and properly uses the special tokens (e.g., `<search>...</search>`).

For each question x , a set of responses $\{y_i\}_{i=1}^G$ is sampled from the policy LLM. Each response is scored by the reward function $r_\phi(q, y)$, producing a corresponding set of rewards $\mathbf{r} = \{r_1, r_2, \dots, r_G\}$. Following Shao et al. (2024), these rewards are then normalized by subtracting the group mean and dividing by the group standard deviation. The global advantages A_i^T for all tokens within the same output are then set to this normalized value:

$$A_i^T = \hat{r}_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}. \quad (5)$$

This design ensures that each trajectory shares a global and base advantage signal across all turns.

3.2 Retrospective Critic: Action Advantage Estimate

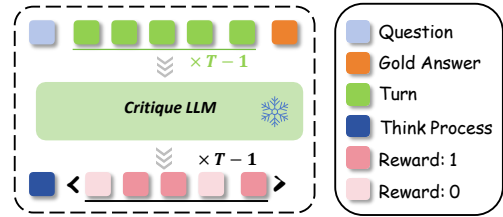


Figure 4: **Overview of the Critique LLM.** The entire reasoning trajectory, along with the gold answer, is fed into the critique LLM, which retrospectively evaluates each action as either Good or Bad and generates a corresponding turn-level reward sequence after thinking.

We next describe the design of the turn-level reward model, which provides dense feedback during multi-turn reasoning. Given a question x , the policy LLM π_θ samples a multi-turn trajectory

$$y_i = \{(a_{i,t}, c_{i,t})\}_{t=1}^T,$$

where $a_{i,t}$ is the action at the turn t and $c_{i,t}$ is the information returned by the search engines in response to $a_{i,t}$. During training, we assume access to privileged information: the ground-truth answer o_{gold} and the full trajectory. A frozen critique LLM C_ϕ takes $(x, y_i, o_{\text{gold}})$ as input and returns a per-turn judgment $\{\ell_{i,t}\}_{t=1}^T, \ell_{i,t} \in \{\text{Good}, \text{Bad}\}$ after thinking (Fig. 4). In this process, the critique LLM adopts a retrospective (hindsight) perspective to evaluate whether each search action has guided the reasoning toward useful information or, conversely,

produced misleading or redundant evidence. Based on this reflection, it judges the quality of each search action.

We map these judgment to turn-level rewards

$$r_{i,t}^a = \begin{cases} 1, & \ell_{i,t} = \text{Good}, \\ 0, & \text{otherwise}, \end{cases} \quad (6)$$

This yields a reward sequence $\mathbf{r}_i^a = \{r_{i,t}^a\}_{t=1}^{T_i}$. A complete scoring example can be found in Fig. 3. Details of the prompt for the critique LLM C_ϕ can be found in Appendix F.

To allocate credit across turns, we normalize per-turn rewards into a turn-level advantage

$$A_{i,t}^a = \frac{r_{i,t}^a}{\sum_{u=1}^{T_i} r_{i,u}^a + \varepsilon}, \quad \varepsilon > 0, \quad (7)$$

which avoids numerical issues when all turns are judged Bad.

This approach provides the model with a form of hindsight supervision, enabling the critique model to infer the contribution of each action using the asymmetric and privileged information that would not be available during forward decision-making.

3.3 Policy Optimization

Intuitively, $A_{i,t}^a$ provides a dense, hindsight credit distribution using privileged information, while $A_{i,t}^\tau$ preserves the global optimization signal. Therefore, we adopt GRPO as the backbone algorithm and replace its original advantage function with a new hybrid advantage during training:

$$\tilde{A}_{i,t} = \alpha A_{i,t}^a + (1 - \alpha) A_{i,t}^\tau, \quad (8)$$

where α is a weighting coefficient that balances the global advantage $A_{i,t}^\tau$ and the turn-level advantage $A_{i,t}^a$.

$$\mathcal{J}(\theta) = \mathbb{E}_{q, \{y_i\} \sim \pi_{\text{old}}} \left[\frac{1}{G} \sum_{i=1}^G \sum_{t=1}^{|y_i|} \tilde{\mathcal{L}}_{i,t} - \beta \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right] \quad (9)$$

where the loss $\tilde{\mathcal{L}}_{i,t}$ is defined as:

$$\tilde{\mathcal{L}}_{i,t} = \min \left(w_{i,t} \tilde{A}_{i,t}, \text{clip}(w_{i,t}, 1 - \epsilon, 1 + \epsilon) \tilde{A}_{i,t} \right) \quad (10)$$

4 Experiment

4.1 Setup

Training We implement our method based on the Search-R1 (Jin et al., 2025) framework. We conduct experiments on two models from the Qwen-2.5 series (Qwen, 2025a). The policy models are

Qwen-2.5-3B-Base and Qwen-2.5-7B-Base, while the critique models are Qwen-2.5-3B-Instruct and Qwen-2.5-7B-Instruct, respectively. The training dataset is HotpotQA (Yang et al., 2018). More details on experimental setups can be found in Appendix B.

Evaluation We evaluate our search agent primarily on four multi-hop Q&A datasets, including HotpotQA (Yang et al., 2018), 2WikiMultiHopQA (Ho et al., 2020), MuSiQue (Trivedi et al., 2022), and Bamboogle (Press et al., 2023). Since our training is conducted on the training set of HotpotQA, the test set of HotpotQA evaluates the in-domain performance, while the test or validation sets of the other three datasets evaluate the out-of-domain performance. We adopt the canonical Exact Match (EM) and word-level F1 scores as two evaluation metrics to ensure fair comparison.

Baselines To effectively evaluate the efficacy of CriticSearch, we compare it against comprehensive baselines: 1) **Direct Inference**: Naive generation (Wei et al., 2022); 2) **Inference with Retrieval**: Search-o1 (Li et al., 2025b); 3) **Agentic RL**: Existing outstanding reinforcement learning methods combined with external search engines, including Search-R1 (Jin et al., 2025), ZeroSearch (Sun et al., 2025a), ReasonRAG (Zhang et al., 2025b), and StepSearch (Zheng et al., 2025a).

4.2 Main Results

CriticSearch consistently outperforms all baseline methods. The main comparative results between CriticSearch and the baseline methods across the four datasets are summarized in Tab. 2. CriticSearch consistently outperforms all baselines on both in-domain multi-hop benchmarks (HotpotQA) and out-of-domain benchmarks (MuSiQue, 2WikiMultiHopQA, and Bamboogle), demonstrating strong robustness and generalization capability. Comparative case studies of CriticSearch versus other algorithms on the same question are provided in Appendix D.

CriticSearch surpasses methods trained with sparse outcome rewards. Compared to Search-R1 and ZeroSearch, which rely solely on sparse reward signals during training, CriticSearch achieves substantially better performance, demonstrating that incorporating dense, fine-grained rewards effectively improves training efficiency.

CriticSearch surpasses existing dense-reward methods. While existing dense-reward baselines

Method	HotpotQA [◊]		2Wiki [◊]		MuSiQue [◊]		Bamboogle [◊]		Average	
	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1
Qwen2.5-3B-Base/Instruct										
Naive Generation	14.5	23.7	24.9	35.6	1.8	7.9	3.0	8.6	11.1	19.0
Search-o1	24.0	32.6	20.7	30.9	4.5	11.7	31.6	43.6	20.2	29.7
Search-R1-base	27.2	36.1	24.8	29.6	8.1	14.6	17.6	27.0	19.4	24.9
Search-R1-instruct	30.4	40.1	29.3	35.2	12.0	18.8	24.0	34.4	23.9	32.1
ZeroSearch-base	26.0	35.4	23.4	28.1	5.6	11.6	9.6	19.3	16.1	23.6
ZeroSearch-instruct	26.5	35.5	23.3	27.8	5.9	12.1	14.4	24.3	17.5	25.0
StepSearch-base	32.9	43.4	<u>33.9</u>	<u>39.5</u>	18.1	27.3	32.8	41.9	29.4	38.0
StepSearch-instruct	<u>34.5</u>	<u>45.2</u>	32.0	38.5	17.4	26.1	<u>34.4</u>	<u>45.2</u>	<u>29.6</u>	<u>38.3</u>
CriticSearch	41.4	53.3	40.9	48.1	<u>18.0</u>	<u>26.8</u>	36.8	47.0	34.3	43.8
Qwen2.5-7B-Base/Instruct										
Naive Generation	18.7	29.1	24.6	35.2	2.7	8.3	12.3	24.2	14.6	24.2
Search-o1	19.3	28.8	18.1	28.9	5.3	12.7	30.2	42.7	18.2	28.3
Search-R1-base	<u>43.2</u>	<u>54.7</u>	35.0	41.1	20.6	29.0	43.0	54.5	35.5	44.8
Search-R1-instruct	39.4	50.2	31.2	37.6	18.1	26.2	38.4	50.1	31.8	41.0
ZeroSearch-base	29.4	39.4	27.5	32.4	10.2	17.5	25.8	37.3	23.2	31.7
ZeroSearch-instruct	32.5	43.2	30.9	37.0	12.0	20.4	26.7	40.9	25.5	35.4
ReasonRAG	38.4	48.9	43.6	50.4	12.8	20.6	36.0	45.5	32.7	41.3
StepSearch-base	38.0	49.3	38.5	45.0	<u>21.6</u>	32.4	<u>46.7</u>	<u>57.3</u>	<u>36.2</u>	<u>46.0</u>
StepSearch-instruct	38.6	50.2	36.6	43.1	22.6	<u>31.2</u>	40.0	53.4	34.5	44.5
CriticSearch	44.2	56.0	<u>42.8</u>	<u>50.1</u>	19.4	28.1	47.2	59.2	38.4	48.4

Table 2: **The main results of CriticSearch.** Bold indicates the top-performing result, while underline denotes the second-best. [◊]/_◊ represents in-domain/out-of-domain benchmarks. We employ critique models with the same parameter scale as the policy models (i.e. Qwen2.5-3B-Instruct for Qwen2.5-3B-Base and Qwen2.5-7B-Instruct for Qwen2.5-7B-Base) in CriticSearch training. CriticSearch versus other algorithms on the same question are provided in Appendix D.

occasionally outperform CriticSearch on specific datasets (e.g., StepSearch on MuSiQue, ReasonRAG on 2Wiki), their overall performance still lags behind that of CriticSearch.

5 Further Analysis

5.1 Accelerate the Convergence of Training

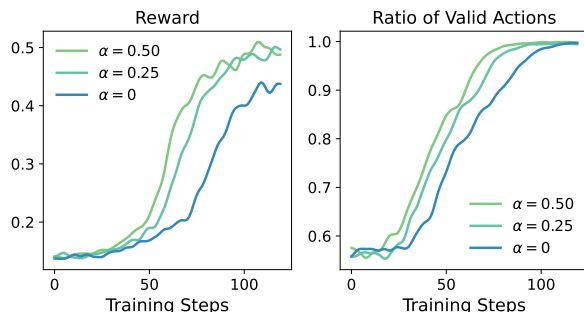


Figure 6: **Ablation study on the weight of the turn-level advantage α .** A larger α leads to faster convergence during the training phase.

As shown in Fig. 6, we illustrate the reward curve and the ratio of valid actions in the early train-

ing stage. We observe that increasing the weight of the turn-level advantage term in Eq. 8 leads to faster convergence in the training stages. This phenomenon is intuitive, as even failed reasoning trajectories may contain beneficial actions, whereas successful ones can still include suboptimal actions. By incorporating the turn-level advantage, the model can more precisely assign learning signals to individual search actions within each trajectory, thereby enhancing the LLM agent’s search capability and training efficiency, and ultimately accelerating overall convergence.

5.2 Mitigating Premature Training Collapse

As shown in Fig. 5, we illustrate the evaluation scores, gradients, and KL divergence during training. Under sparse rewards, the algorithm collapses after around 400 steps, whereas incorporating the turn-level advantage effectively mitigates this collapse and suppresses abnormal fluctuations in both the KL loss and gradients, allowing performance to continue improving.

When training with only sparse outcome re-

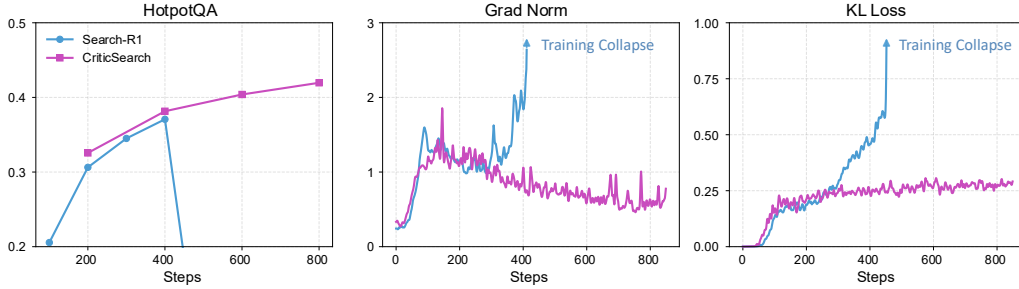


Figure 5: **Training stability and performance comparison.** CriticSearch effectively constrains both the gradient norm and the KL loss within a stable range, preventing divergence and ensuring smooth optimization. This improved stability enables CriticSearch to sustain longer training and achieve higher accuracy.

Method	Critique model	HotpotQA	2Wiki	MuSiQue	Bamboogle
<i>Qwen2.5-3B-Base</i>					
CriticSearch	<i>None</i>	38.1	39.0	14.1	32.8
CriticSearch	<i>Qwen2.5-3B-Instruct</i>	41.4	40.9	18.0	36.8
CriticSearch	<i>Qwen3-30B-A3B-Instruct-2507</i>	41.9	42.4	18.6	43.2
<i>Qwen2.5-7B-Base</i>					
CriticSearch	<i>None</i>	41.8	40.0	18.7	40.8
CriticSearch	<i>Qwen2.5-7B-Instruct</i>	44.2	42.8	19.4	47.2
CriticSearch	<i>Qwen3-30B-A3B-Instruct-2507</i>	46.1	43.2	20.5	40.8

Table 3: **Ablation study on the size of critique LLM.** A larger critique LLM yields better performance.

wards, each action within a correct trajectory receives a reward of 1, while all actions in incorrect trajectories receive 0. From the perspective of turn-level rewards, this introduces a highly noisy turn-level reward model that may misclassify poor actions as high-reward or penalize good actions with low scores. Such noise drives the LLM to optimize toward unreasonable search paths, producing redundant or inefficient behaviors misaligned with human preferences. The optimization over noisy reward signals during training causes the model to deviate from the human-aligned distribution represented by the reference model (e.g., Qwen-2.5-3B-base) (Liu et al., 2024; Wang et al., 2025b). Consequently, when suboptimal actions receive undeservedly high rewards, the model distribution diverges from that of the reference model, leading to a continuously increasing KL loss and, ultimately, training collapse. Moreover, training with sparse rewards leads to high variance in the return and advantage estimates, which further exacerbates instability during training (Appendix C.2). In contrast, turn-level rewards provide more accurate supervision for each search action. By maintaining stable and informative reward signals throughout training, they effectively alleviate the collapse phenomenon.

5.3 Ablation study on the Size of Critique

Next, we conduct an ablation to verify whether a larger critique LLM produces higher-quality dense rewards, thereby training the policy LLM more effectively (Tab. 3). We compare three settings for supervising the same policy model (Qwen2.5-3/7B-Base): 1) a small critique model (Qwen2.5-3/7B-Instruct), 2) a larger critique model (Qwen3-30B-A3B-Instruct-2507 (Qwen, 2025b)), and 3) no critique model (i.e., no turn-level reward). As shown in Tab. 3, the policy trained with the larger critique model achieves higher average performance across four multi-hop Q&A datasets. These results provide evidence that a larger critique model provides more reliable dense feedback and more effectively assists policy learning, and validating the usefulness of the critique model. A detailed analysis of the accuracy of the turn-level reward is provided in Section 5.4.

5.4 Accuracy analysis of Critique Model

We sample 20 trajectories and, for each turn within a trajectory, assigned scores using 1) different critique models, including Gemini-2.5-pro (DeepMind, 2025), Qwen3-30B-A3B-It, GPT-4o (OpenAI, 2024), and Qwen2.5-7B-It; 2) Monte Carlo estimates; 3) Outcome Reward (i.e., using the outcome reward as the turn-level reward), and 4) hu-

Method	HotpotQA	2Wiki	MuSiQue	Bamboogle	Average
<i>Qwen2.5-7B-Base</i>					
CriticSearch-Sparse	41.8	40.0	18.7	40.8	35.3
CriticSearch-NG	44.9	40.7	19.2	41.6	36.6
CriticSearch	44.2	42.8	19.4	47.2	38.4

Table 4: **Ablation study on the privileged information of the critique LLM.** CriticSearch-NG is the variant where the gold answer is removed from the critique LLM input during training. CriticSearch-Sparse is the variant trained solely with the outcome reward.

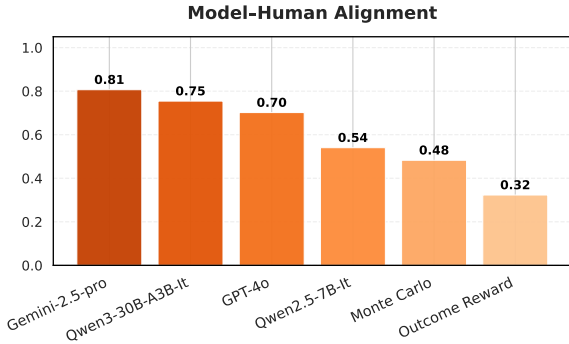


Figure 7: **Alignment Across Dense-Reward Methods.** Higher similarity to human annotations indicates more accurate turn-level rewards.

man annotations. We then analyzed the agreement among these scoring methods. Assuming human annotations as the ground truth, we observe that larger and more capable critique models exhibit higher alignment with human judgments. Specifically, Gemini-2.5-Pro achieves the highest agreement with human annotations, reaching approximately 80% similarity. Qwen2.5-7B-It achieves a level of agreement with human scores comparable to Monte Carlo estimation; however, the Monte Carlo approach requires many more rollouts to obtain reasonably accurate estimates and is therefore more resource-intensive. Notably, the larger Qwen3-30B-A3B-It model surpasses the Monte Carlo approach in its alignment with human annotations. In contrast, using the outcome reward as the turn-level score results in the lowest agreement with human judgments. Further details on this experiment are provided in Appendix C.3.

5.5 Ablation study on the privileged information of Critique LLM

To evaluate the effect of privileged information in the critique LLM (i.e., the gold answer and future trajectory), we ablated the gold answer from its input and measured the accuracy of the resulting turn-level rewards. As shown in Fig. 8, excluding the gold answer leads to a moderate decline in re-

ward accuracy. Nevertheless, with the benefit of a retrospective assessment, the critique LLM can still capture the semantics of the search actions. Therefore, its turn-level rewards remain more accurate than outcome-based turn-level rewards. Consistently, the ablation results in Table 4 indicate that removing the gold answer causes some performance degradation, yet the model still outperforms the variant without a critique LLM (i.e., using only the outcome reward).

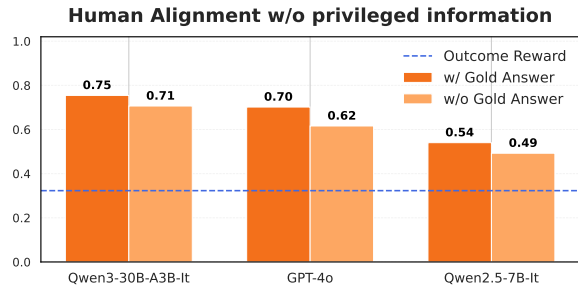


Figure 8: Removing the gold answer from the critique LLM slightly reduces its alignment with human annotations, indicating a moderate decline in reward accuracy. The blue dashed line indicates the alignment when using the outcome reward as the turn-level reward.

6 Conclusion

This paper introduces CriticSearch, a fine-grained credit assignment method that leverages large language models to assist policy training. By introducing a retrospective critic mechanism, we employ a lightweight turn-level binary evaluator that delivers more effective dense feedback while mitigating training instability. Extensive experiments and analyses demonstrate that CriticSearch consistently improves learning efficiency and stability, highlighting its effectiveness and robustness across diverse multi-hop Q&A benchmarks.

7 Limitations

The inclusion of the Critique model increases the memory footprint required for model loading,

while the additional reward computation introduces a small amount of extra training overhead compared to GRPO (Appendix C.1). Moreover, the proposed method has so far been validated only on search-related tasks with a limited number of reasoning turns. Future work may explore reusing the policy model as the critique model during training and further explore the method’s effectiveness in broader agentic RL scenarios.

8 Acknowledgements

This work was supported in part by the Strategic Priority Research Program of Chinese Academy of Sciences under Grant XDA0480302, in part by National Natural Science Foundation of China under Grants 62293541 and 62136008, and in part by Beijing Nova Program under Grant 20240484514.

References

- Mingyang Chen, Linzhuang Sun, Tianpeng Li, sunhaoze, ZhouYijie, Chenzheng Zhu, Haofen Wang, Jeff Z. Pan, Wen Zhang, Huajun Chen, Fan Yang, Zenan Zhou, and Weipeng Chen. 2025. [Research: Learning to reason with search for LLMs via reinforcement learning](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Google DeepMind. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Guanting Dong, Yifei Chen, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Yutao Zhu, Hangyu Mao, Guorui Zhou, Zhicheng Dou, and Ji-Rong Wen. 2025. 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, Guorui Zhou, Yutao Zhu, Ji-Rong Wen, and Zhicheng Dou. 2026. [Agentic reinforced policy optimization](#). In *The Fourteenth International Conference on Learning Representations*.
- Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. 2025a. Retool: Reinforcement learning for strategic tool use in llms. *arXiv preprint arXiv:2504.11536*.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025b. 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*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-rl: Training LLMs to reason and leverage search engines with reinforcement learning. In *Second Conference on Language Modeling*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick SH Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6769–6781, Online. Association for Computational Linguistics.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, and 1 others. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. 2020. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, page eabc5986.
- Weizhen Li, Jianbo Lin, Zhuosong Jiang, Jingyi Cao, Xinpeng Liu, Jiayu Zhang, Zhenqiang Huang, Qianben Chen, Weichen Sun, Qiexiang Wang, and 1 others. 2025a. Chain-of-agents: End-to-end agent foundation models via multi-agent distillation and agentic rl. *arXiv preprint arXiv:2508.13167*.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025b. Search-o1: Agentic search-enhanced large reasoning models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5420–5438.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.

- Jiacai Liu, Chaojie Wang, Chris Yuhao Liu, Liang Zeng, Rui Yan, Yiwen Sun, Yang Liu, and Yahui Zhou. 2024. Improving multi-step reasoning abilities of large language models with direct advantage policy optimization. *arXiv preprint arXiv:2412.18279*.
- Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 6382–6393, Red Hook, NY, USA. Curran Associates Inc.
- OpenAI. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, and 1 others. 2018. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A Smith, and Mike Lewis. 2023. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiushi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. Toolrl: Reward is all tool learning needs. *arXiv preprint arXiv:2504.13958*.
- Qwen. 2025a. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Qwen. 2025b. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- 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, Rotterdam, Netherlands. ACM.
- Jingwei Shi, Xinxiang Yin, Jing Huang, Jinman Zhao, and Shengyu Tao. 2026. Codehacker: Automated test case generation for detecting vulnerabilities in competitive programming solutions. *arXiv preprint arXiv:2602.20213*.
- Hao Sun, Zile Qiao, Jiayan Guo, Xuanbo Fan, Yingyan Hou, Yong Jiang, Pengjun Xie, Yan Zhang, Fei Huang, and Jingren Zhou. 2025a. Zerosearch: Incentivize the search capability of llms without searching. *arXiv preprint arXiv:2505.04588*.
- Shuang Sun, Huatong Song, Yuhao Wang, Ruiyang Ren, Jinhao Jiang, Junjie Zhang, Fei Bai, Jia Deng, Wayne Xin Zhao, Zheng Liu, and 1 others. 2025b. Simpledeepsearcher: Deep information seeking via web-powered reasoning trajectory synthesis. *arXiv preprint arXiv:2505.16834*.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, and 1 others. 2025. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*.
- Harsh Trivedi, Niranjana 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.
- Songjun Tu, Jiahao Lin, Qichao Zhang, Xiangyu Tian, Linjing Li, Xiangyuan Lan, and Dongbin Zhao. 2025a. Learning when to think: Shaping adaptive reasoning in rl-style models via multi-stage RL. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Songjun Tu, Jingbo Sun, Qichao Zhang, Yaocheng Zhang, Jia Liu, Ke Chen, and Dongbin Zhao. 2025b. In-dataset trajectory return regularization for offline preference-based reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Hanlin Wang, Chak Tou Leong, Jiashuo Wang, Jian Wang, and Wenjie Li. 2025a. Spa-rl: Reinforcing llm agents via stepwise progress attribution. *arXiv preprint arXiv:2505.20732*.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*.
- Tao Wang, Ruipeng Zhang, and Sicun Gao. 2025b. Improving value estimation critically enhances vanilla policy gradient. In *Forty-second International Conference on Machine Learning*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pages 24824–24837, Red Hook, NY, USA. Curran Associates Inc.

- Quan Wei, Siliang Zeng, Chenliang Li, William Brown, Oana Frunza, Wei Deng, Anderson Schneider, Yuriy Nevmyvaka, Yang Katie Zhao, Alfredo Garcia, and Mingyi Hong. 2025. Reinforcing multi-turn reasoning in llm agents via turn-level reward design. *arXiv preprint arXiv:2505.11821*.
- Zhenghai Xue, Longtao Zheng, Qian Liu, Yingru Li, Xiaosen Zheng, Zejun Ma, and Bo An. 2025. Simpletir: End-to-end reinforcement learning for multi-turn tool-integrated reasoning. *arXiv preprint arXiv:2509.02479*.
- 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, Brussels, Belgium. Association for Computational Linguistics.
- Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, and 1 others. 2025. Glm-4.5: Agentic, reasoning, and coding (arc) foundation models. *arXiv preprint arXiv:2508.06471*.
- Qi Zhang, Shouqing Yang, Lirong Gao, Hao Chen, Xiaomeng Hu, Jinglei Chen, Jiexiang Wang, Sheng Guo, Bo Zheng, Haobo Wang, and Junbo Zhao. 2025a. Lets: Learning to think-and-search via process-and-outcome reward hybridization. *arXiv preprint arXiv:2505.17447*.
- Shaokun Zhang, Yi Dong, Jieyu Zhang, Jan Kautz, Bryan Catanzaro, Andrew Tao, Qingyun Wu, Zhiding Yu, and Guilin Liu. 2026. [Nemotron-research-tool-n1: Exploring tool-using language models with reinforced reasoning](#). In *The Fourteenth International Conference on Learning Representations*.
- Wenlin Zhang, Xiangyang Li, Kuicai Dong, Yichao Wang, Pengyue Jia, Xiaopeng Li, Yingyi Zhang, Derong Xu, Zhaocheng Du, Huifeng Guo, Ruiming Tang, and Xiangyu Zhao. 2025b. [Process vs. outcome reward: Which is better for agentic RAG reinforcement learning](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Yaocheng Zhang, Yuanheng Zhu, Yuqian Fu, Songjun Tu, and Dongbin Zhao. 2025c. Offline goal-conditioned reinforcement learning with elastic-subgoal diffused policy learning. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, page 2336–2344, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.
- Zijie Zhao, Honglei Guo, Shengqian Chen, Kaixuan Xu, Bo Jiang, Yuanheng Zhu, and Dongbin Zhao. 2025. Empowering multi-robot cooperation via sequential world models. *arXiv preprint arXiv:2509.13095*.
- Xuhui Zheng, Kang An, Ziliang Wang, Yuhang Wang, and Yichao Wu. 2025a. Stepsearch: Igniting llms search ability via step-wise proximal policy optimization. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 21816–21841.
- Yuxiang Zheng, Dayuan Fu, Xiangkun Hu, Xiaojie Cai, Lyumanshan Ye, Pengrui Lu, and Pengfei Liu. 2025b. Deepresearcher: Scaling deep research via reinforcement learning in real-world environments. *arXiv preprint arXiv:2504.03160*.
- Yifei Zhou, Song Jiang, Yuandong Tian, Jason Weston, Sergey Levine, Sainbayar Sukhbaatar, and Xian Li. 2025a. Sweet-rl: Training multi-turn llm agents on collaborative reasoning tasks. *arXiv preprint arXiv:2503.15478*.
- Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Pu Liang. 2025b. Mem1: Learning to synergize memory and reasoning for efficient long-horizon agents. *arXiv preprint arXiv:2506.15841*.

A Related Works

A.1 Agentic RL with Search Engines

Recently, Reinforcement Learning has emerged as a promising framework for enhancing the reasoning capabilities of large language models, demonstrating strong proficiency in complex reasoning tasks (Shao et al., 2024; Guo et al., 2025). To address the limited access of LLMs to up-to-date external knowledge, several studies (Jin et al., 2025; Zheng et al., 2025a; Chen et al., 2025; Zhang et al., 2026) have explored RL-based approaches that integrate external search engines to improve information retrieval effectiveness. Beyond this, a number of works aim to further enhance search-agent performance from multiple perspectives (Gao et al., 2025; Sun et al., 2025b). Some methods refine underlying mechanisms by summarizing previous interactions to shorten the effective context length (Zhou et al., 2025b), introducing fine-grained reward functions for more efficient optimization (Zheng et al., 2025a; Wei et al., 2025; Zhang et al., 2025a; Qian et al., 2025), and leveraging multi-agent collaboration for improved coordination and reasoning (Li et al., 2025a; Dong et al., 2025). Other studies focus on optimizing interaction strategies and data utilization. For example, DeepResearch integrates search agents with real-world search environments (Zheng et al., 2025b), while ZeroSearch replaces traditional search engines with LLM-based simulated retrieval, and several works (Sun et al., 2025b; Gao et al., 2025) investigate improvements in multi-task settings and with richer data resources. Our work is similar to StepSearch, introducing a fine-grained reward mechanism for credit assignment under sparse rewards. Unlike StepSearch, which relies on curated ground-truth references for supervision, CriticSearch adopts a streamlined, general framework that leverages the existing strong reasoning ability of the LLM to evaluate intermediate actions and assign credit.

A.2 Reward Models in LLM Reinforcement Learning

The Reinforcement Learning with Verifiable Rewards (RLVR) paradigm (Lambert et al., 2024) has greatly advanced large language models’ reasoning in mathematics and code generation, as shown by recent systems such as Guo et al. (2025), Team et al. (2025), and Zeng et al. (2025). However, outcome-based rewards remain inherently sparse, offering supervision only at the end of reasoning trajectory-

ries. This limitation can reward spurious reasoning paths that coincidentally yield correct answers and cause severe credit-assignment issues (Lightman et al., 2024), especially in multi-turn agentic RL where models must coordinate sequential tool calls. To address these problems, prior research has explored complementary directions: 1) designing explicit process-level rewards that score intermediate reasoning or tool usage (Qian et al., 2025; Zheng et al., 2025a). However, such approaches are often tailored to specific tools or datasets, resulting in limited generalizability. 2) employing outcome-driven reward inference that propagates final rewards backward through trajectory grouping or branch analysis (Feng et al., 2025b; Dong et al., 2026). These approaches often suffer from high variance and require large amounts of rollout for stable estimation. 3) employing process reward model to generate dense reward, including MT-PPO and PPR. These methods primarily rely on forward-only heuristic signals and do not explicitly leverage future information or hindsight. 4) training turn-level critics that provide local supervision for each tool-calling step (Zhou et al., 2025a; Wang et al., 2025a). Similar to the fourth approach, our method introduces a lightweight turn-level binary evaluator that simply labels each tool invocation as either good or bad, thereby reducing the risk of reward hacking. Unlike training-based critics, the evaluator remains frozen during learning and requires no additional training. This not only avoids the instability and computational overhead associated with joint critic optimization but also eliminates the need for extra training data, thereby enhancing both efficiency and generalizability. This design effectively bridges the gap between sparse outcome supervision and dense step-level feedback, yielding a more stable and scalable online RL framework for agentic reasoning.

B Experiment Setups

Our training is conducted on the VeRL (Sheng et al., 2025) framework. For GRPO training, we set the group size G to 5 and the policy LLM learning rate 5×10^{-7} for all models. The rollout temperature is set to 1.0. We train for a total of 800 steps, with a warm-up ratio of 0.35. The total batch size, mini batch size, and micro batch size are set to 128, 64, and 32 respectively. The entropy regularization coefficient and the KL loss coefficient are both set to $1e-3$, and the PPO clip ratio ϵ to 0.2. To maintain

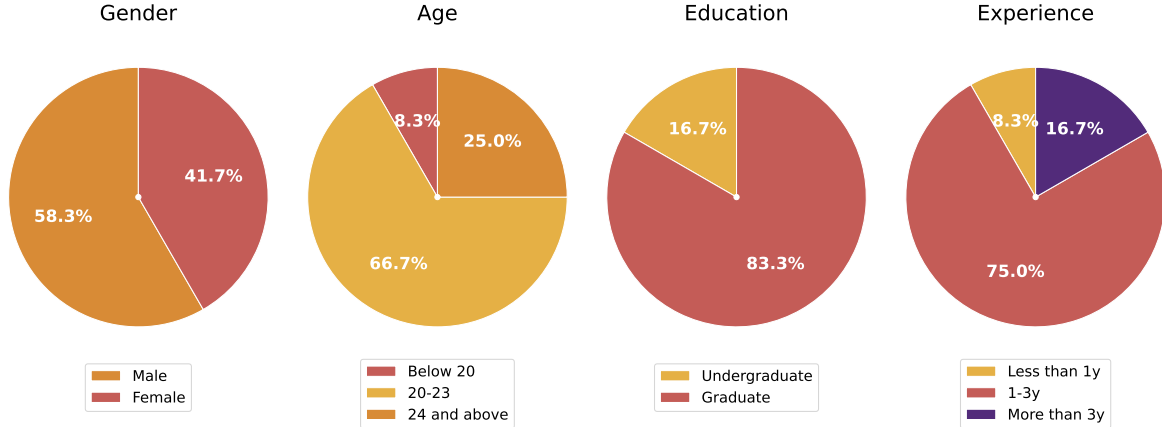


Figure 9: **Distribution of participants involved in the human evaluation experiment is presented.** Each participant was asked to provide basic information, including gender, age, education level, and experience of AI. We present the distribution of these information categories in the form of pie charts.

alignment with prior work (Jin et al., 2025; Zheng et al., 2025a), we use the 2018 Wikipedia dump (Karpukhin et al., 2020) as the corpus and employ E5 (Wang et al., 2022) as the retriever during training. At each retrieval step, we uniformly sample $k = 3$ documents. The maximum context length is set to 4K tokens, with a maximum of four turns of searching. The same configuration is adopted during evaluation.

We set λ_f as 0.2 for all experiments. We sweep the weighting coefficient $\alpha \in \{0.25, 0.5, 0.75\}$ and observe that the optimal setting is consistently $\alpha = 0.25$ for all models. All training jobs are conducted on a node with 8 NVIDIA H100 GPUs. Early stopping is applied when training collapse is detected from the reward curve. For the results in Tab. 2 and Tab. 3, we report either the final checkpoint at step 800 or the latest checkpoint before collapse.

To ensure a fair comparison, we adopt the original open-source model checkpoints and their published prompt configurations for all baselines. For consistency, they are evaluated using the same retrieval configuration as our agent.

C Additional Results

C.1 Training Cost

Tab. 5 reports the average per-step training time of different methods. Across various search agents, we observe that although our approach introduces an additional retrospective critic mechanism compared with Search-R1, it does not significantly increase the per-step training time. When using a critique model of the same size as the policy model,

Method	Critique model	Time
<i>Qwen2.5-3B-Base</i>		
Search-R1	-	68.4
CriticSearch	<i>Qwen2.5-3B-Instruct</i>	79.5
CriticSearch	<i>Qwen3-30B-A3B-Instruct-2507</i>	88.9
<i>Qwen2.5-7B-Base</i>		
Search-R1	-	114.7
CriticSearch	<i>Qwen2.5-7B-Instruct</i>	110.9
CriticSearch	<i>Qwen3-30B-A3B-Instruct-2507</i>	140.8

Table 5: **The training time.** We analyze the per-step training time (s) of different methods. When using a critique model of the same size as the policy model, CriticSearch introduces almost no additional training time compared with Search-R1.

CriticSearch introduces almost no additional training time compared with Search-R1. Especially when using Qwen2.5-7B-Base as the policy model and Qwen2.5-7B-Instruct as the critique model, CriticSearch even achieves a shorter per-step time than Search-R1. We attribute this to the fact that Search-R1 lacks fine-grained rewards to penalize low-quality actions, leading to redundant and repetitive search action during rollouts (Fig. 10), which in turn increases training time.

The reported training times were measured on a node equipped with 8 NVIDIA H100 GPUs, and we follow the same training configurations described in Appendix B.

C.2 Variance Analysis of Advantage Estimates

To analyze the variance of the advantage estimates, we used the same checkpoint (Qwen2.5-7B trained

Table 6: Variance analysis under different percentile ranges.

Method	0–20%↓	20–40%↓	40–60%↓	60–80%↓	80–100%↓
A^τ	0.381	0.253	0.250	0.265	0.245
$0.75A^\tau$	0.214	0.142	0.141	0.149	0.138
$0.75A^\tau + 0.25A^a$	0.137	0.116	0.120	0.121	0.118

with outcome-based GRPO at step 300) to perform a strictly identical rollout and compared three advantage formulations: A^τ , $0.75A^\tau$, and our hybrid advantage $0.75A^\tau + 0.25A^a$. Tokens were divided into five buckets by log-probability percentiles ($\sim 130k$ tokens each) to measure advantage variance across probability ranges. This allows us to examine how advantage variance behaves across different probability regimes, particularly in the low-probability tail. As shown in Tab. 6, CriticSearch achieves lower advantage variance than both baselines across all buckets, confirming that reduction stems from the retrospective critique, not mere scaling. The effect is most pronounced in the low-probability bucket (0–20%), where variance drops by 64% ($0.381 \rightarrow 0.137$). This shows that the process critique effectively prevents low-probability tokens from receiving extreme, noisy advantage signals. Since low-probability tokens are widely recognized as the primary cause of training collapse in agentic RL (Xue et al., 2025), CriticSearch successfully improves training stability by suppressing these fluctuations.

C.3 Accuracy analysis of Critique Model

We sampled 20 trajectories using the SearchR1-qwen2.5-7b-grpo-v0.3 model. For the Monte Carlo rearward estimates, we used the same model to roll out 10 trajectories forward from each action and took the average as the estimated action reward. For Gemini-2.5-Pro (DeepMind, 2025) and GPT-4o (OpenAI, 2024), we adopted the same prompt (Appendix F) as Qwen2.5-7B-It (Qwen, 2025b) to perform reward generation.

For human-labeled rewards, we recruited 12 human evaluators to assess the quality of each action along the trajectories. Each evaluator was informed of the purpose of the data collection. Additionally, evaluators were asked to provide basic demographic information to help us understand the evaluator distribution. The demographic fields included gender, age, education, and experience with AI. The distribution of evaluators for each task is

shown in Fig 9

D Case Study

We observe the sampled trajectories during the CriticSearch training process, and representative examples are presented in Tab. 7, Tab. 8 and Tab. 9. All cases are obtained from the model trained on Qwen2.5-7B-Base. In particular, we further compare the trajectory generated by the baseline SearchR1 and our CriticSearch model on the same question, as shown in Fig. 10. Each trajectory contains multi-turn interactions between the model and the external search engine. Our CriticSearch method introduces retrospective critic fine-grained credit assignment at the process level, where each search action is judged by an external reasoning-capable LLM according to its usefulness to the final answer. This dense supervision enables the policy to assign credit at the step level—encouraging informative searches while penalizing redundant or misleading ones. As a result, in the later stage, the model naturally learns to issue concise, logically progressive queries (e.g., decomposing multi-entity comparisons or refining layered reasoning chains), forming more efficient trajectories with minimal redundancy and stable convergence.

E Algorithm Workflow

In this section, we present a detailed description of the complete training workflow of *CriticSearch* (Alg. 1).

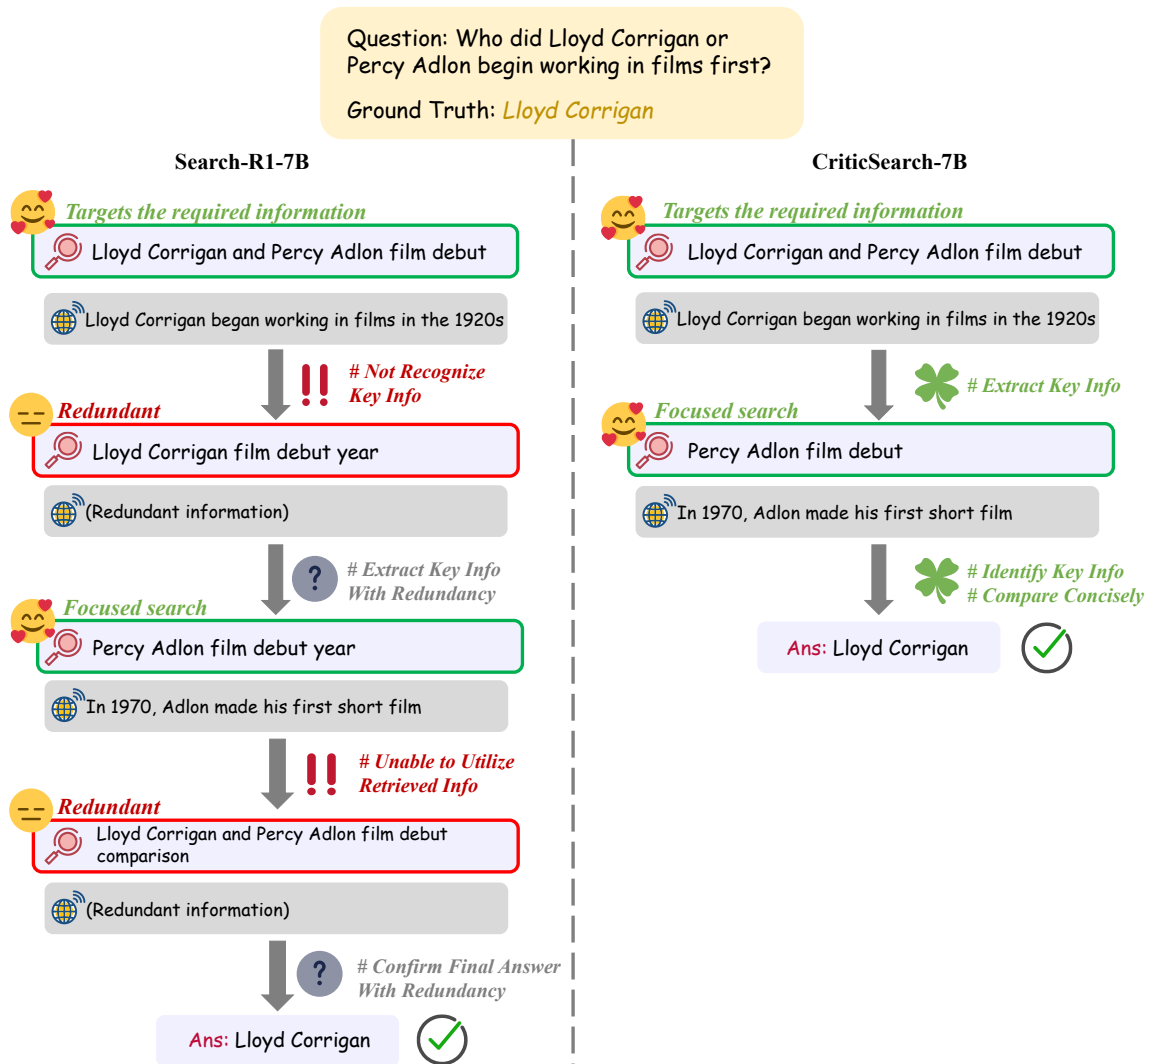


Figure 10: Side-by-side trajectories of Search-R1 (left, generated by *SearchR1-qwen2.5-7b-grpo-v0.3*) and CriticSearch (right) on the same query. Each trajectory shows multi-turn interactions with the search engine (actions, responses, and final answer). Although both Search-R1 and CriticSearch answer the query correctly, our method (right) uses fewer, more focused queries and reaches a logically structured answer with minimal redundancy.

Question: When was the college, for which Willie Fritz was head coach from 2014 to 2015, founded?

Ground Truth: ['1906']

~~<think> I need to determine when the college for which Willie Fritz was head coach from 2014 to 2015 was founded. </think>~~
~~<search> Willie Fritz head coach 2014 to 2015 </search>~~
~~<information> Doc 1(Title: "Willie Fritz") From 2014 to 2015, he was head coach at Georgia Southern University... Doc 2(Title: "Willie Fritz") From 2014 to 2015, he was head coach at Georgia Southern University... Doc 3(Title: "Willie Fritz")...On January 10, 2014, Fritz became the head football coach at Georgia Southern University... </information>~~
~~<think> I found information about Willie Fritz and his coaching positions, but I need to determine the founding year of the college he was head coach at from 2014 to 2015. </think>~~
~~<search> Georgia Southern University founded </search>~~
~~<information> Doc 1(Title: "Georgia Southern University") Founded in 1906 as a land grant college, Georgia Southern is part of the University System... Doc 2(Title: "Georgia State University") Georgia State University Georgia State University is a public research university in downtown Atlanta ... Doc 3(Title: "Savannah State University") Savannah State University was founded as a result of the Second Morrill Land Grant Act of August 30, 1890. </information>~~
~~<think> I found the founding year of Georgia Southern University. </think>~~
~~<answer> 1906 </answer>~~

Table 7: An example from the late stage of CriticSearch training. The model demonstrates the ability to perform progressive search actions in a natural manner. Red-font text denotes **valid information**; struck-through text indicates **question-independent information**; and gray text represents **redundant content**.

Question: Aftermath was based on a real life midair plane collision that happened on what date?

Ground Truth: ['1 July, 2002']

~~<think> I need to find out when the real life midair plane collision that Aftermath was based on happened. </think>~~
~~<search> Aftermath midair plane collision date </search>~~
~~<information> Doc 1(Title: "Mid-air collision") The US Defense Safety Oversight Council established the official U.S. civil-military mid-air collision prevention website in 2006... Doc 2(Title: "Mid-air collision") The primary causes of mid-air collisions include miscommunication, navigational errors, and the absence of collision-avoidance systems... Doc 3(Title: "Mid-air collision") Several historical U.S. civil-military collision incidents led to the Federal Aviation Act of 1958 and later initiatives. </information>~~
~~<think> I found information about midair collisions and their prevention, but I need to find the date of the specific collision that Aftermath was based on. </think>~~
~~<search> Aftermath movie midair collision date </search>~~
~~<information> Doc 1(Title: "Aftermath (2017 film)") The film is based on the July 2002 Überlingen mid-air collision... Doc 2(Title: "Aftermath (2017 film)") Aftermath (originally titled 478) is a 2017 American drama thriller film directed by Elliott Lester and written by Javier Gullón... Doc 3(Title: "Aftermath (2017 film)") In November 2015, Highland Film Group sold the film to different international distributors at the American Film Market. </information>~~
~~<think> I found that Aftermath is based on the July 2002 Überlingen mid-air collision. Now I need to confirm the date of this collision. </think>~~
~~<search> Überlingen midair collision date </search>~~
~~<information> Doc 1(Title: "2002 Überlingen mid-air collision") On the night of 1 July 2002, Bashkirian Airlines Flight 2937, ... Doc 2(Title: "2002 Überlingen mid-air collision") It was being flown by two Bahrain-based pilots, 47-year-old British Captain... Doc 3(Title: "2002 Überlingen mid-air collision") The Tupolev aircraft continued descending despite TCAS advisories and that the accident prompted recommendations to upgrade TCAS systems... </information>~~
~~<think> I found the date of the Überlingen mid-air collision. </think>~~
~~<answer> July 1, 2002 </answer>~~

Table 8: An example from the late stage of CriticSearch training. After obtaining invalid information in the first search, the model adjusts its query keywords and continues the subsequent search process. Red-font text denotes **valid information**; struck-through text indicates **question-independent information**; and gray text represents **redundant content**.

Question: What rock band was formed in Oxford and reached 5 in the UK Charts?

Ground Truth: ['Supergrass']

<think> I need to determine which rock band was formed in Oxford and reached number 5 in the UK Charts. *</think>*
<search> rock band formed in oxford reached number 5 in the uk charts *</search>*
<information> Doc 1(Title: "Oxford") **The most notable Oxford act is Radiohead, who all met at nearby Abingdon School, though other well known local bands include Supergrass...** Doc 2(Title: "The Oxford Circle")The Oxford Circle was an American garage rock and psychedelic rock band from Davis... Doc 3(Title: "Oxford") Bands from over 30 years of the Oxford music scene's history feature in the documentary film ""Anyone Can Play Guitar?"". *</information>*
<think> I found information about different bands from Oxford, but I need to find the specific band that reached number 5 in the UK Charts. *</think>*
<search> Radiohead reached number 5 in the uk charts *</search>*
<information> Doc 1(Title: Radiohead) ""**Street Spirit**"" placed Radiohead in the top five for the first time... Doc 2(Title: Radiohead) **Radiohead are an English rock band from Abingdon, Oxfordshire, formed in 1985...** Doc 3(Title: "Radiohead discography") ""Amnesiac"", Radiohead's fifth studio album, was released in June 2001, **topping the UK charts** and producing successful singles including... *</information>*
<think> I found that Radiohead is a rock band from Abingdon, Oxfordshire, not Oxford, but they reached number 5 in the UK Charts with the song "Street Spirit (Fade Out)". *</think>*
<search> Supergrass reached number 5 in the uk charts *</search>*
<information> Doc 1(Title: "Supergrass")**The band's debut album, ""I Should Coco"" (May 1995), which entered the UK Albums Chart at number one. The album's fourth single, the double A-side release ""Alright""/"Time"", stayed in the UK Top Three for a month, peaking at number two ...** Doc 2(Title: "Supergrass")Supergrass Supergrass were an English rock band, formed in 1993 in Oxford. ~~The band consisted of brothers Gaz...~~ Doc 3(Title: "Supergrass")**The record entered the UK albums chart at number four and has since gone gold in the UK.** *</information>*
<think> I found that Supergrass, a rock band from Oxford, reached number 5 in the UK Charts with the song "Alright". *</think>*
<answer> Supergrass *</answer>*

Table 9: An example from the late stage of CriticSearch training. The policy model demonstrates improved search reasoning and evidence localization ability. Red-font text denotes **valid information**; struck-through text indicates **question-independent information**; and gray text represents **redundant content**.

Algorithm 1 Training Stage of CriticSearch

Require: Policy model π_θ , Reference model π_{ref} , Frozen critic C_ϕ , Training dataset \mathcal{D} , Hyperparameters $\lambda_f, \alpha, \beta, \epsilon, \varepsilon$

Ensure: Optimized policy π_θ

```

1: # Rollout
2: Sample a batch  $\mathcal{D}_b$  from  $\mathcal{D}$ 
3: for each question  $q \in \mathcal{D}_b$  do
4:   Sample  $G$  responses  $\{y_i\}_{i=1}^G \sim \pi_\theta(q)$ 
5:
6:   # Estimate Global Advantage
7:   for each response  $y_i$  do
8:     Compute outcome reward  $r_i = r_\phi(q, y_i)$ 
9:   end for
10:  for each response  $y_i$  do
11:    Compute global advantage  $A_{i,t}^r$  via Eq. 5
12:  end for
13:
14:  # Estimate Turn-level Advantage
15:  for each response  $y_i$  do
16:    Obtain per-turn judgments  $\{\ell_{i,t}\}_{t=1}^T = C_\phi(q, y_i, o_{\text{gold}})$ 
17:    Map to binary rewards  $r_{i,t}^a$  by Eq. 6
18:    Compute turn-level advantage  $\{A_{i,t}^a\}_{t=1}^T$  using Eq. 7
19:  end for
20:
21:  # Hybrid Advantage
22:  Compute Hybrid Advantage  $\tilde{A}_{i,t}$  via Eq. 8.
23: end for
24:
25: # Policy Optimization
26: for step  $j \leftarrow 1$  to  $M$  do
27:   Update  $\pi_\theta$  with Eq. 9, where per-token loss  $\hat{\mathcal{L}}_{i,t}$  is computed by Eq. 10
28: end for
29: return  $\pi_\theta$ 

```

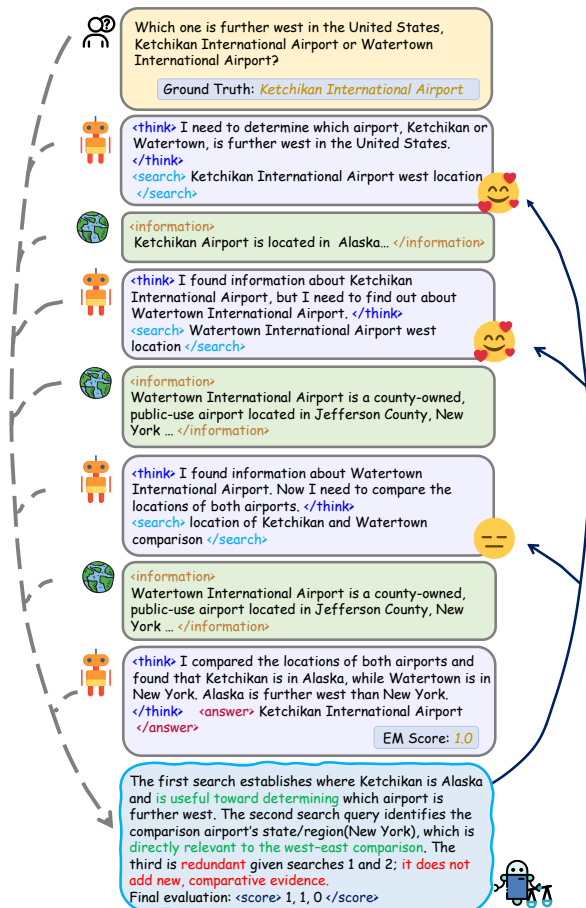


Figure 11: An example from CriticSearch illustrating the evaluation process of the Critique model. In this trajectory, the final answer is correct but the whole trajectory involves a redundant search action. A frozen critique LLM, after deliberation, provides accurate binary rewards for each turn.

F Template

Prompt 2: Reward Prompt

Please analyze whether each search process is good or bad, and include the final evaluation result within a `<score>...</score>` tag. If it is good, output 1; if it is bad, output 0. For example, if there are two actions, output `<score>1, 1</score>`; if there are four actions, output `<score>1, 0, 0, 1</score>`. When making judgments, strictly follow these rules:

1. If the "search" action contributes to the final answer or provides partially useful information, it is good.
2. If the "search" action provides redundant information, or repeats a search that has already been done, it is bad.
3. If the "search" action points to the wrong search direction or misleading information, it is bad.
4. If the "search" action results in incorrect information due to unclear expression, it is bad.
5. Do not evaluate the "answer" actions, only evaluate the "search" actions.
6. If there are no search actions in the entire trajectory, only return `<score></score>`.
7. Only evaluate valid "search" actions (queries that contain `<search>...</search>` and receive `<information>...</information>` feedback afterward).
8. The final number of scores must match the number of "search" actions.
9. Ultimately, first provide a detailed analysis of the search process, then enclose the evaluation results within a `<score>...</score>` tag.

#####

#####

Golden answers: {Ground truth}

Extracted answer: {Answer}

Solution string: {Solution Trajectory}

#####

#####

Please conduct a detailed analysis.

The final number of scores should be consistent with the number of search actions.