

# MemPO: Self-Memory Policy Optimization for Long-Horizon Agents

Ruoran Li<sup>1</sup>, Xinghua Zhang<sup>2</sup>, Haiyang Yu<sup>2</sup>, Shitong Duan<sup>2</sup>, Xiang Li<sup>2</sup>,  
Wenxin Xiang<sup>1</sup>, Chonghua Liao<sup>1</sup>, Xudong Guo<sup>2</sup>, Yongbin Li<sup>2\*</sup>, Jinli Suo<sup>1\*</sup>,

<sup>1</sup>Tsinghua University

<sup>2</sup>Tongyi Lab, Alibaba Group,

lrr24@mails.tsinghua.edu.cn, zhangxinghua.zxh@alibaba-inc.com,

shuide.lyb@alibaba-inc.com, jlsuo@tsinghua.edu.cn

## Abstract

Long-horizon agents face the challenge of growing context size during interaction with environment, which degrades the performance and stability. Existing methods typically introduce the external memory module and look up the relevant information from the stored memory, which prevents the model itself from proactively managing its memory content and aligning with the agent’s overarching task objectives. To address these limitations, we propose the self-memory policy optimization algorithm (**MemPO**), which enables the agent (policy model) to autonomously summarize and manage their memory during interaction with environment. By improving the credit assignment mechanism based on memory effectiveness, the policy model can selectively retain crucial information, significantly reducing token consumption while preserving task performance. Extensive experiments and analyses confirm that MemPO achieves absolute F1 score gains of 25.98 over the base model and 7.1 over the previous SOTA baseline, while reducing token usage by 67.58% and 73.12%. The code is released at <https://github.com/TheNewBeeKing/MemPO>.

## 1 Introduction

As large language models (LLMs) continue to evolve, LLM agents are becoming increasingly proficient in addressing more complex problems. In areas such as deep research (Zhang et al., 2025; Zheng et al., 2025), data analysis (Hong et al., 2025), and vibe coding (Zhang et al., 2024; Islam et al., 2024; Ho et al., 2025), they have showcased remarkable performance. Long-horizon decision-making has always been one of the core capabilities for agents to solve complex user queries.

Currently, the dominant method for the agent-environment interaction is ReAct paradigm (Yao et al., 2022). The feedback from the environment is attached to the previous interaction history and is

used as a prompt, which then determines the next course of action. However, this approach causes the context to grow linearly with each round of interaction, resulting in longer contexts when tackling more complex problems, and presenting several challenges. Firstly, current LLMs have relatively limited context window sizes, which impose an explicit upper bound on the number of interactions. Secondly, long contexts lead to excessively high token costs, which impedes the widespread adoption of agent systems in practical scenarios. Furthermore, excessively long contexts can lead to the “lost in the middle” phenomenon (Liu et al., 2023), which degrades the model’s ability, thereby reducing the overall performance of the agent.

To address this challenge, a growing body of research is focusing on agent memory, with the aim of providing LLMs with historical interaction records to reduce the need for the entire context. Currently, the mainstream solution involves designing a memory module as an external knowledge database to maintain the agent’s interaction history. When the memory module is accessed, relevant historical information is retrieved and integrated into the prompt (Borgeaud et al., 2022; Gao et al., 2024; Lewis et al., 2020) based on the retrieval-augmented generation technique (RAG). However, the offline memory context compression method lacks the capacity for joint optimization oriented toward the agent task execution, making it difficult to effectively align with the agent’s overarching task objectives. As a result, the model’s memory retrieval remains passive, rather than leveraging its own capabilities to proactively select and organize information, and the latter would facilitate more effective task completion.

To this end, we formalize the agent interaction paradigm as autonomously refining and organizing historical information, while simultaneously reasoning and invoking tools with three actions <mem>, <think>, and <tool\_call>. In this

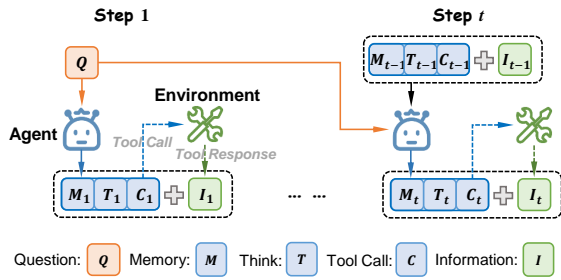


Figure 1: The self-memory inference process of our method, which only uses the previous step interaction for next step input with `<mem>` action.

paradigm, the agent itself proactively compresses and reorganizes long-horizon historical information for the next step of interaction, making memory management an intrinsic part of its capabilities, as shown in Figure 1. To further enhance this ability, we propose self-memory policy optimization (**MemPO**), which incorporates the trajectory-level and memory-level information into advantage estimation to optimize the `<mem>` action for agent with task-objective awareness. Concretely, the tokens output by the agent are assigned trajectory-level advantages, and in each step of the interaction, the tokens of the `<mem>` action additionally take into account memory-level advantages, effectively alleviating the credit assignment problem in long-horizon, multi-turn interactions. In terms of dense rewards for `<mem>` action in each step, the conditional probability of the answer given `<mem>` content is designed to measure the quality of the `<mem>` action. Our contributions are as follows:

- We render memory management an intrinsic part of the agent’s own capabilities that differs from external memory modules, achieving joint optimization of long-horizon memory, reasoning, and tool invocation.
- We propose **MemPO**, a self-memory policy optimization algorithm, which effectively addresses credit assignment and steers the `<mem>` action toward retaining the most relevant information for solving the task.
- Extensive experiments on five long-horizon benchmarks confirm the efficacy of MemPO with 25.98% and 7.1% absolute F1 gains over the base model and previous SOTA, 67.58% and 73.12% reductions in token usage.

## 2 Related Works

### 2.1 Memory for LLM agents

In recent years, researchers have introduced external memory and experience systems to address the limitations of LLM context windows (Xu et al., 2025; Chhikara et al., 2025; Zheng et al., 2024; Packer et al., 2024; Zhong et al., 2023; Zhang et al., 2026; Liao et al., 2025b). MemGPT (Packer et al., 2024) proposes an operating-system-inspired virtual memory management framework that employs multiple memory hierarchies to manage contextual information. Mem0 (Chhikara et al., 2025) enhances memory capacity through dynamic extraction, consolidation, and retrieval of conversational information. Despite their effectiveness in specific domains, most of these approaches rely on fixed workflows and limited optimization flexibility. They typically fail to support flexible cross-stage joint optimization, which constrains the adaptability and scalability of the overall system.

### 2.2 RAG in Memory System

RAG has emerged as a powerful approach for enhancing LLM by incorporating external knowledge sources to improve model performance (Borgeaud et al., 2022; Gao et al., 2024; Lewis et al., 2020). In existing memory systems, the retrieval of relevant memory fragments is predominantly implemented based on RAG. While this approach can efficiently surface relevant information in certain scenarios, its major limitation lies in the lack of flexibility and end-to-end joint optimization. Specifically, retrieval relies solely on embedding similarity between the query and chunks, which does not necessarily yield information that is most useful for solving the target problem.

### 2.3 RL for LLM Agents

The recent success of reinforcement learning methods in LLMs has established RL as a central tool to enhance LLM-based agents to solve increasingly complex tasks (Jin et al., 2025; Chen et al., 2025; Zheng et al., 2025; Kang et al., 2025; Nie et al., 2026; Liao et al., 2025a). However, relatively few studies have explored applying RL to the optimization of agent memory. Existing approaches exhibit notable limitations. For example, MEM1 (Zhou et al., 2025) integrates memory into the reasoning process and applies RL optimization for the policy model. However, it does not explicitly design objectives for memory optimization, which can lead

to suboptimal memory representations. In contrast, our method introduces a dedicated credit assignment mechanism for memory rewards, encouraging the model to retain information that is most relevant for solving the target task.

### 3 Preliminaries

#### 3.1 Task Formulation

Given a question-answer pair  $(q, a_{gt})$ , when an LLM-based agent is tasked with solving the question  $q$ , it interacts with the external environment through multiple rounds of reasoning and tool invocation to acquire the information required for problem solving. If the agent completes the task after  $T$  steps, a full trajectory can be denoted as  $\tau = \{s_1, s_2, \dots, s_T\}$ .

Each state  $s_t$  is further decomposed into  $\{s_t^{mem}, s_t^{think}, s_t^{call}, s_t^{resp}\}$ . Specifically,  $s_t^{mem}$  represents the model-generated summary of effective information from previous outputs  $s_{<t}$ , which is enclosed by `<mem>` `</mem>`.  $s_t^{think}$  corresponds to the model’s reasoning process and is wrapped by `<think>` `</think>`.  $s_t^{call}$  denotes the invocation of external tools by the model, which is represented as `<tool_call>` `</tool_call>`.  $s_t^{resp}$  captures the information returned by the tool and is enclosed by `<information>` `</information>`. Once the agent has gathered sufficient information to answer the question  $q$ , it produces a predicted answer  $a_{pred}$ , wrapped by `<answer>` `</answer>`.

#### 3.2 Group Relative RL

In reinforcement learning for LLM, a class of group-based methods, exemplified by Group Relative Policy Optimization (GRPO) (Shao et al., 2024), abandon per-trajectory value function modeling and instead performs relative comparison within a batch of candidate trajectories. Concretely, for a given task input  $q$ , the policy  $\pi_{\theta_{old}}$  generates  $N$  complete trajectories  $\{\tau_1, \tau_2, \dots, \tau_N\}$  in one shot, and each trajectory is assigned a scalar return  $R(\tau_i)$  that measures the overall quality of the generated outcome. The algorithm then relies solely on statistics within this group to construct advantages, without explicitly learning a value network:

$$A(\tau_i) = \text{GroupAgg}\left(\{R(\tau_j)\}_{j=1}^N, i\right), \quad (1)$$

where  $\text{GroupAgg}(\cdot)$  is the aggregation operator based on normalization, or pairwise comparison.

The design bypasses the instabilities of value function estimation and reduces it to modeling relative preferences among a set of candidate answers. In large-scale LLM training, group-based methods can reduce the memory overhead of extra networks, making them an efficient alternative for RL training.

#### 3.3 Behavior Cloning

To enable the model to better follow the action format, we first adopt GPT-4.1 (OpenAI et al., 2024) to perform inference on the publicly available training dataset from the work of (Tang et al., 2025), filter out trajectory with incorrect answers, and finally generate approximately 10K trajectories following the predefined action format in § 3.1. Based on these trajectory data, we fine-tune the LLM and provide a promising starting point for self-memory policy optimization.

### 4 Self-Memory Policy Optimization

As mentioned above, memory is introduced to address long contexts of agents by removing irrelevant information and retaining key details. Vanilla GRPO computes rewards based on answer correctness and uses trajectory-level advantages, where tokens within the same trajectory share the same reward. It provides the sparse rewards and limited guidance for memory generation, as the correctness of the final answer can not directly reflect the quality of each `<mem>` action during the interaction.

To address this, we propose **MemPO**, a self-memory policy optimization algorithm. We design a novel advantage computation method that, in addition to trajectory-level advantages, evaluates the information content of memory within `<mem>` `</mem>` at each step and computes an additional advantage, ensuring memory remains concise while preserving important information.

#### 4.1 Advantages of Global Trajectory

We first evaluate the trajectory format and the accuracy of the final answer to provide a coarse-grained assessment of the overall trajectory quality. Suppose that for a single training sample, we perform  $N$  rollouts and assign an overall score to each resulting trajectory, denoted as a group:

$$G^T = \{(\tau_1, R^T(\tau_1)), (\tau_2, R^T(\tau_2)), \dots, (\tau_N, R^T(\tau_N))\}. \quad (2)$$

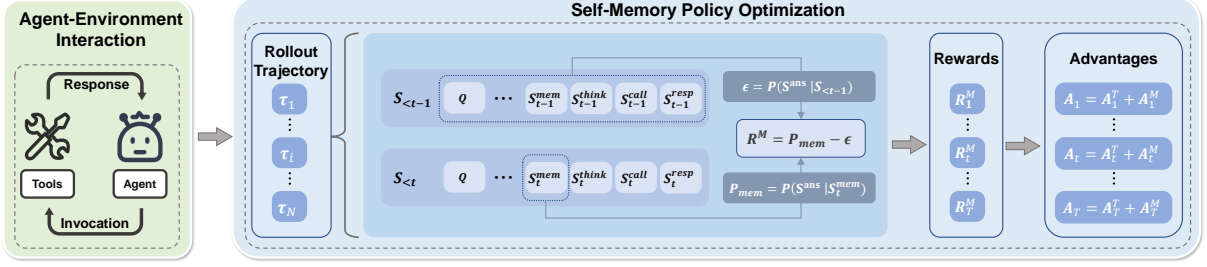


Figure 2: Overview of MemPO. At step  $t$  of any trajectory  $\tau_i$ , the context is represented as  $\{s_t^{mem}, s_t^{think}, s_t^{call}, s_t^{resp}\}$ . The memory reward  $R^M$  is calculated using conditional probabilities and contributes to the advantage  $A^M$ . The final advantage is the sum of  $A^M$  and the trajectory-level advantage  $A^T$ . During inference, only the previous step’s content is used as context, discarding earlier information.

where  $\tau_i$  denotes a trajectory, and  $R^T(\tau_i)$  represents the trajectory-level reward. In our method, the reward consists of evaluations of both the output format and the correctness of predicted answer. Concretely, the reward is set to 1 if and only if the predicted answer is correct and output format is proper; otherwise, it is set to 0.

To assess the global relative quality of each trajectory within the group, we adopt the advantage calculation strategy from GRPO (Shao et al., 2024), which normalizes the total reward using the mean and standard deviation computed over the group:

$$A^T(\tau_i) = \frac{R^T(\tau_i) - \text{mean}(\{R^T(\tau_j)\}_{j=1}^N)}{\text{std}(\{R^T(\tau_j)\}_{j=1}^N)}. \quad (3)$$

## 4.2 Advantages of Informative Memory

According to the probabilistic formulation of LLMs, the output of model is characterized as conditional probabilities given the preceding context, i.e.,  $\pi_\theta(s_t | q, s_{<t})$  (Vaswani et al., 2017). As demonstrated and exploited in previous work (Wang et al., 2025; Kim and Lee, 2024; Lewis et al., 2020), if the context  $s_{<t}$  contains sufficient information to solve the problem, the probability that  $s_t$  is sampled as the answer-generation step  $s^{\text{ans}}$  will be relatively high. Similarly, this insight suggests that for an arbitrary context  $s^{\text{any}}$ , a higher value of  $\pi_\theta(s^{\text{ans}} | q, s^{\text{any}})$  indicates that  $s^{\text{any}}$  contains more key information relevant to solving the question  $q$ , thereby increasing the model’s confidence in generating the correct answer. Consequently, conditional probability can be used as a quantitative measure of the effective information content contained in a given context.

Based on this intuition, we design a step-level reward for the memory (`<mem>` action) generated at each interaction step, which reflects the quality of

effective information retained in memory:

$$R^M(\tau_i(s_t^{mem})) = P[s^{\text{ans}} | \tau_i(s_t^{mem})] - \epsilon, \quad (4)$$

$$1 \leq i \leq N, 1 \leq t \leq T.$$

where  $\tau_i(s_t^{mem})$  denotes the memory content within `<mem>` action in step  $t$  of trajectory  $\tau_i$ , and  $s^{\text{ans}} = \{a_1, a_2, \dots, a_L\}$  represents the correct answer string, where  $a_l$  denotes the  $l$ -th token of the answer. The  $\epsilon$  represents  $P[s^{\text{ans}} | \tau_i(s_{<t})]$ , which serves as a bias term. The term  $\tau_i(s_{<t})$  corresponds to the trajectory of the first  $t-1$  steps of  $\tau_i$ . The operator  $P(\cdot)$  denotes a posterior-probability-based measure. Specifically,  $P[s^{\text{ans}} | \tau_i(s_t^{mem})]$  can be represented as:

$$\sqrt[L]{\prod_{l=1}^L \pi_\theta(a_l | q, \tau_i(s_t^{mem}), a_{<l})}. \quad (5)$$

where  $\pi_\theta(a_l | q, \tau_i(s_t^{mem}), a_{<l})$  denotes the probability of generating token  $a_l$  given the user query, the trajectory prefix up to step  $t-1$ , and the preceding answer tokens  $a_{<l}$ .

Under this formulation, a larger value of  $R^M(\tau_i(s_t^{mem}))$  indicates that the memory generated at step  $t$  provides a more effective summary of the trajectory up to the first  $t-1$  steps, and better preserves contextual information that is relevant to generating the correct answer.

Based on the reward formulation above, the resulting memory group can be expressed as:

$$G^M = \{(\tau_i(s_t^{mem}), R^M(\tau_i(s_t^{mem}))) \mid 1 \leq i \leq N, 1 \leq t \leq T\}. \quad (6)$$

We then normalize the rewards using the group-wise mean and standard deviation to obtain the corresponding advantages  $A^M(\tau_i(s_t^{mem}))$ :

$$A^M(\tau_i(s_t^{mem})) = \frac{R^M(\tau_i(s_t^{mem})) - M(\tau_i(s_t^{mem}))}{\text{std}(\{R^M(\tau_i(s_t^{mem}))\})}. \quad (7)$$

where  $M(\tau_i(s_t^{mem}))$  denotes the mean reward within the same group, defined as:

$$M(\tau_i(s_t^{mem})) = \text{mean}(\{R^M(\tau_i(s_t^{mem})) \mid (\tau_i(s_t^{mem}), R^M(\tau_i(s_t^{mem}))) \in G^M\}). \quad (8)$$

The advantage  $A^M(\tau_i(s_t^{mem}))$  provides a quantitative assessment of memory quality, enabling finer-grained supervision over the model-generated memory content.

### 4.3 Combination of Advantages

The final token-level advantage is obtained by combining the two types of advantages in § 4.1 and § 4.2. Let the  $k$ -th token of the  $i$ -th trajectory  $\tau_i$  in a group be denoted as  $\tau_{i,k}$ . The advantage assigned to this token  $A_{i,k}$  is defined as:

$$A_{i,k} = \begin{cases} A^T(\tau_i) + A^M(\tau_i(s_t^{mem})), & \tau_{i,k} \in \tau_i(s_t^{mem}) \\ A^T(\tau_i), & \text{otherwise.} \end{cases} \quad (9)$$

That is, when  $\tau_{i,k}$  corresponds to a token within the memory segment (`<mem>` action), its advantage is given by the sum of the trajectory-level advantage and the memory-level advantage; otherwise, only the trajectory-level advantage  $A^T$  is used. In this way, tokens belonging to memory receive richer and more explicit feedback signals, which more effectively guide the rollout process for memory generation.

### 4.4 Policy Optimization and Inference

**Optimization.** The policy optimization objective is to maximize  $\mathcal{J}(\theta)$ , written as:

$$\mathcal{J}(\theta) = \mathbb{E} \left[ \frac{1}{N} \sum_{i=1}^N \frac{1}{|\tau_i|} \sum_{k=1}^{|\tau_i|} \min \left( \gamma_{i,k} A_{i,k}, \text{clip}(\gamma_{i,k}, 1 - \epsilon, 1 + \epsilon) A_{i,k} \right) - \beta \mathcal{D}_{\text{KL}}(\pi_\theta \parallel \pi_{\text{ref}}) \right]. \quad (10)$$

where  $\gamma_{i,k}$  is the importance sampling ratio:

$$\frac{\pi_\theta(\tau_{i,k} \mid q, \tau_{i,<k})}{\pi_{\theta_{\text{old}}}(\tau_{i,k} \mid q, \tau_{i,<k})}. \quad (11)$$

where  $q \sim p(Q)$ ,  $\{\tau_i\}_{i=1}^N \sim \pi_{\theta_{\text{old}}}$ .  $p(Q)$  denotes the distribution of queries in the training set, and  $\beta$  controls the weight of the KL-divergence regularization term.

**Inference.** In vanilla ReAct framework, the  $t$ -th step inference is denoted as  $\pi_\theta(s_t \mid q, s_{<t})$ . In our method, since  $s_{t-1}^{mem}$  contains the effective information of  $s_{<t-2}$ , we replace  $s_{<t}$  with  $s_{t-1}^{mem}$  as the inference context, represented as  $\pi_\theta(s_t \mid q, s_{t-1}^{mem})$ .

## 5 Experiments

### 5.1 Benchmarks

To evaluate the effectiveness of our approach, following the method of MEM1 (Zhou et al., 2025), we test on multi-objective tasks, where the number of interaction rounds required for the agent to solve a problem is significantly higher compared to single-objective tasks. This allows us to better assess the performance of our method in scenarios with long contexts. Additionally, we can observe the changes in agent performance by progressively increasing the number of objectives. We created a 2-objective task test set by combining queries from the validation sets of the HotpotQA (Yang et al., 2018) and NQ (Kwiatkowski et al., 2019) QA datasets, and synthesized test sets with more objectives using the HotpotQA validation set. We conducted tests under both local wiki search engine and web search engine scenarios to enhance the credibility of the experiments.

Following previous work (Zhou et al., 2025), we use the **F1** score as a criterion for word-level matching, and Exact Match (**EM**) for exact matching. Furthermore, to evaluate the token consumption of the agent when solving a problem, we use the total number of tokens consumed to solve a question (**TT**), as well as the maximum number of tokens (peak tokens) consumed in a single step (**PT**).

### 5.2 Baselines

We compare our method with various baselines. For prompt-based baselines, we use ReAct (Yao et al., 2022). For agentic RL-based baselines, we adopt DeepResearcher (Zheng et al., 2025) and ReSearch (Chen et al., 2025). For agent memory-related baselines, we use RL-based method MEM1 (Zhou et al., 2025) and RAG-based method A-MEM (Xu et al., 2025). Additionally, we also trained a model without memory using GRPO in the exact same environment as a baseline. To ensure fairness, all methods use the 7B model from the Qwen2.5 series as the base model.

### 5.3 Implementation Details

We first performed inference using GPT-4.1 (OpenAI et al., 2024) on the dataset from the work of (Tang et al., 2025), and obtained approximately 10k trajectories containing memory. We then fine-tuned the model for one epoch using these data to enhance its ability to follow instructions related to the memory component. Simultaneously, we removed

Table 1: The accuracy and token consumption for multi-objective tasks of baselines. **Text with bold** means SOTA.

Local Wiki Search														
Model	2-objective		4-objective		6-objective		8-objective		10-objective		Avg			
	F1	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1 ↑	EM ↑	TT ↓	PT ↓
Qwen2.5 (ReAct)	33.73	25.60	10.59	7.00	5.37	4.00	5.92	4.25	2.63	1.96	11.65	8.56	3.64	0.61
ReSearch	47.40	36.00	24.13	16.70	20.84	15.93	10.80	7.70	5.08	3.56	21.65	15.98	3.29	0.71
DeepResearcher	30.94	24.70	24.52	18.10	13.88	10.73	9.12	6.85	5.07	3.52	16.71	12.78	4.29	0.77
A-MEM	33.24	25.20	13.71	10.10	9.80	7.07	6.86	5.10	5.56	3.60	13.83	10.21	2.62	0.38
MEM1	47.74	37.10	26.51	18.90	18.81	14.07	19.04	13.55	19.61	13.36	26.34	19.40	1.38	0.20
GRPO (w/o <mem>)	54.57	42.95	38.31	28.60	29.78	22.60	18.97	13.65	11.01	7.84	30.53	23.13	4.39	0.81
<b>MemPO (Ours)</b>	<b>56.47</b>	<b>46.15</b>	<b>42.75</b>	<b>31.90</b>	<b>34.32</b>	<b>26.93</b>	<b>30.48</b>	<b>23.70</b>	<b>24.15</b>	<b>18.16</b>	<b>37.63</b>	<b>29.37</b>	<b>1.18</b>	<b>0.18</b>
Online Web Search														
Model	2-objective		4-objective		6-objective		8-objective		10-objective		Avg			
	F1	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1 ↑	EM ↑	TT ↓	PT ↓
Qwen2.5 (ReAct)	45.71	33.60	16.83	11.80	12.12	9.07	9.66	6.80	7.25	4.68	18.31	13.19	3.14	0.34
ReSearch	51.17	39.20	29.92	21.10	25.21	18.73	17.09	12.50	10.37	7.44	26.75	19.79	2.17	0.40
MEM1	50.56	39.60	30.43	22.00	21.67	16.20	19.48	14.30	18.06	12.12	28.04	20.84	0.96	0.14
<b>MemPO (Ours)</b>	<b>57.40</b>	<b>45.20</b>	<b>41.42</b>	<b>30.20</b>	<b>37.92</b>	<b>28.60</b>	<b>34.30</b>	<b>25.80</b>	<b>22.92</b>	<b>16.32</b>	<b>38.79</b>	<b>29.22</b>	<b>0.86</b>	<b>0.12</b>

the memory component from the trajectories to use them for fine-tuning the baseline model, which was trained using GRPO and does not include memory, ensuring fairness in the comparison.

In the RL phase, we followed MEM1 and used the 2-objective task synthesized from HotpotQA and NQ as part of the training set. And we randomly sampled a subset from both datasets as another part of training set. The rollout group size  $N$  for group-based RL methods is set to 16, with a batch size of 128 and a learning rate of 1e-6. The maximum number of interaction rounds is set to 16. During training, we use the local wiki search engine as the search tool.

## 5.4 Experimental Results

**Multi-objective task.** The results of each baseline on the multi-objective task are shown in Table 1. We selected tasks with 4, 6, 8, and 10 objectives as progressively harder task groups and recorded F1 and EM for the answers from each baseline as precision metrics. Among the baselines, MEM1, A-MEM, and our method use truncated contexts, meaning the model only has access to the previous step of interactions, while the other baselines use the complete context. Additionally, we also recorded the total number of tokens required to solve a single problem (TT) and the peak token consumption per step (PT) during the model’s execution. The more detailed results are presented in Appendix Table 2.

**Conditional probability analysis.** To investigate the impact of our reward design on the model, we performed a statistical analysis of the true values of  $P[s^{\text{ans}} | s^{\text{mem}}]$  during inference on the 10-objective task. We compared the results of models

trained with vanilla GRPO and our method. Specifically, let the dataset size be  $M$ , and denote the memory component at step  $t$  of the  $m$ -th trajectory as  $\tau_m(s_t^{\text{mem}})$ , with the corresponding ground truth answer string being  $\tau_m(s^{\text{ans}})$ . Figure 5 shows the grouped results of  $P[\tau_m(s^{\text{ans}}) | \tau_m(s_t^{\text{mem}})]$ , where the x-axis represents the group values of the conditional probability, and the y-axis shows the proportion of memory samples in that range relative to the total memory samples. This distribution illustrates the conditional probability distribution of the memory produced by the model. The line graph’s y-axis represents the average accuracy of trajectories whose memory falls within each group, providing insight into the relationship between accuracy and the conditional probability.

Figure 4 displays the aggregated results of  $P[\tau_m(s^{\text{ans}}) | \tau_m(s_t^{\text{mem}})]$  by step. The x-axis represents the step  $t$  of the memory, and the y-axis of the line graph shows the average conditional probability of the memory at step  $t$  across all  $M$  trajectories. This provides insight into how conditional probability evolves with the step. The histogram shows the proportion of memories at each step relative to the total number of trajectories.

**Ablation study.** To validate the effectiveness of our design, we compared the performance of the memory-enabled model trained with vanilla GRPO and our model, keeping all other conditions identical. The only difference between the two models is the inclusion of a reward specifically for memory. The results are shown in the left panel of Figure 3. Additionally, we evaluated our method under different context retention settings: full context, retaining 1 or 3 interaction rounds. The results




Figure 3: The results of ablation study.

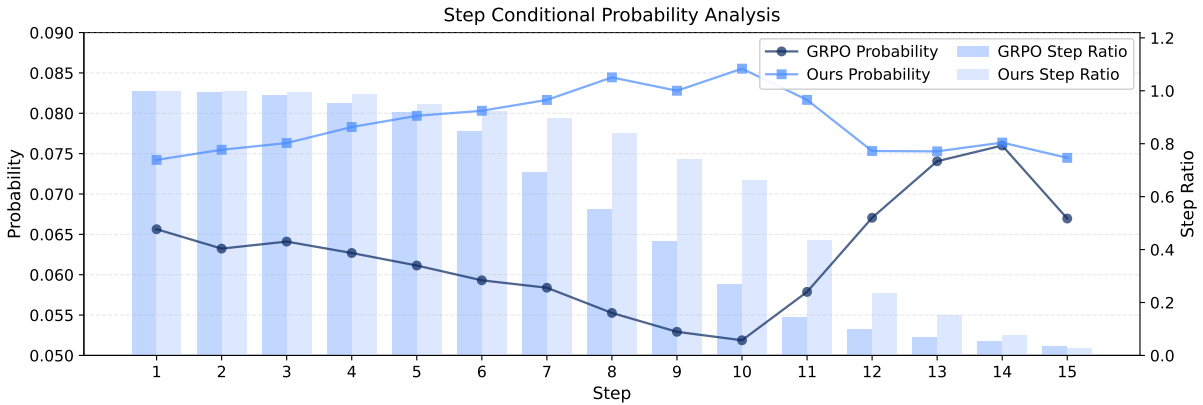


Figure 4: The result of step conditional probability analysis.

of these experiments are shown in the right panel of Figure 3.

## 5.5 Experimental Analyses

**Our method demonstrates remarkable performance and generalization.** As shown in Table 1, we achieve SOTA performance on tasks that are much more difficult than those in the training set, and our model continues to maintain leading performance when switched to a real-world web search environment that differs from the training setup. Additionally, as presented in Table 1, our model achieves SOTA performance while minimizing token consumption, thereby achieving the highest performance with the least resource usage. Furthermore, as seen in Figure 3, whether using full context, retaining 1 or 3 interaction rounds for inference, our method’s performance remains consistently stable, showcasing strong generalization capabilities.

**The memory mechanism significantly reduces token consumption.** As shown in Appendix Table 2, token consumption for tasks solved by MEM1, A-MEM, and our method, which all incorporate the memory mechanism, is noticeably

lower than that of other baselines. Taking ReSearch as an example and comparing it with our method, when the task is relatively simple, such as a 2-objective task, the token consumption is only slightly higher than ours. However, as the complexity of the task increases, the gap between the two methods becomes more pronounced. By the time the task reaches 10 objectives, the number of tokens required by our method to solve a problem is approximately 1/3 of ReSearch’s, with the token peak being 1/5. This is comparable to ReSearch’s token consumption on a 4-objective task. Moreover, our method not only uses truncated contexts but also provides effective guidance on the memory content, resulting in even more compact contexts compared to other memory-related methods. As a result, token consumption in our method is the lowest among all baselines.

**The ability of memory to provide strongly relevant information is crucial for task success.** In the baseline methods, A-MEM generates memory based on a RAG approach. As mentioned in the introduction, the memory obtained by this method is not necessarily the most relevant for solving the task and contains a significant amount of redun-

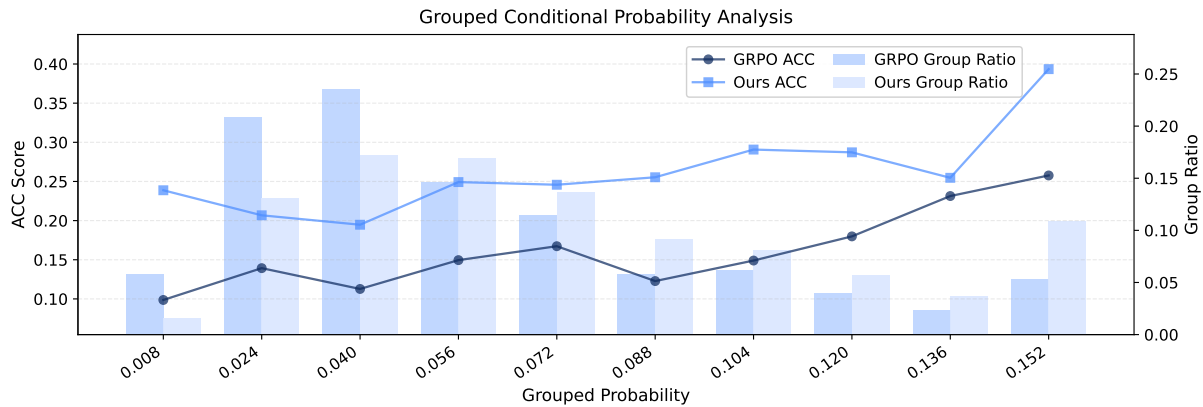


Figure 5: The result of grouped conditional probability analysis.

dancy. As a result, while A-MEM reduces token consumption compared to ReAct, its performance does not show a significant improvement. On the other hand, MEM1 generates memory by combining the model’s summary with reasoning, creating a stronger link between memory generation and the task-solving process. This allows MEM1 to show a considerable improvement on long-horizon tasks. Furthermore, our method explicitly guides the model to retain the context that most strongly contributes to solving the problem, outperforming other memory-related baselines in all datasets.

**The number of context steps impacts performance on long-horizon tasks.** As shown in the right panel of Figure 3, we present the results of inference using complete context, truncated in 1 step and in 3 steps. Overall, the performance differences between these methods fluctuate within an acceptable range, with the trend showing that the more context steps used, the better the performance on short-horizon tasks, but the weaker the performance on long-horizon tasks. This effect is particularly noticeable on long-horizon tasks. We believe that this aligns with the phenomenon of attention dilution caused by long contexts, which leads to performance degradation.

**Our reward design positively contributes to improving the effective information content.** Figure 3 presents a performance comparison between vanilla GRPO and our method. The results show that our reward design leads to an improvement in the model’s performance. We also quantitatively analyzed the information content of the memory in the trajectories of both vanilla GRPO and our method. The bar graph in Figure 5 indicates that, compared to the baseline, our method’s probability distribution is more skewed toward higher values,

which contributes to greater precision of responses, as confirmed by the line graph.

Additionally, the line graph in Figure 4 shows that, for the first 10 steps of 10-objective task, the mean probability of our method increases as the steps progress, whereas the baseline shows a decreasing trend. We believe this reflects the more effective organization of memory by our method compared to the baseline. After 10 steps, our method’s probability starts to decrease, which is reasonable given that the typical number of search steps in a 10-objective task is around 10. If the task is not completed by this point, it suggests that some information is difficult to find and is still being searched for. In contrast, the baseline experiences more difficulty in the first 10 steps, and as seen in the bar chart, only 20% of the search examples continue after the 10th step. We hypothesize that the few remaining examples that did not abandon exploration likely achieved relatively higher accuracy, which explains the continued increase in probability after the 10th step.

Overall, both the final performance and the probability analysis validate that our reward design is effective and aligns with expectations.

## 6 Conclusion

Our method optimizes memory management for agents by introducing a novel reward design that retains only relevant information, improving task performance and reducing computational costs. By integrating memory, reasoning, and tool invocation via reinforcement learning, we achieve superior performance, especially on long-horizon tasks, with efficient token consumption. Future work will focus on optimizing memory reward design and enhancing the scalability of self-memory method in more applications.

## Limitations

Although our method shows promising performance, one potential limitation is that, due to varying tool invocation at different steps, the information content in memory naturally differs and the states are not completely equivalent across all steps in all rollout trajectories, which may introduce bias when calculating group-based advantages. While we alleviate this issue by introducing  $\epsilon$  in Equation 4, more refined solutions may reduce this bias in more complex environments, and understanding the generalization to diverse real-world settings needs further investigation.

## LLM Usage Statement

In this paper, LLMs are used to assist in polishing the grammar of the article and translating related content, without any non-compliant usage.

## Acknowledgment

This work was jointly supported by the National Key R&D Program of China (Grant No. 2024YFF0505703) and Beijing Municipal Natural Science Foundation (Grant No. L257009).

## References

- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego De Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, and 9 others. 2022. [Improving Language Models by Retrieving from Trillions of Tokens](#). In *Proceedings of the 39th International Conference on Machine Learning*, pages 2206–2240. PMLR. ISSN: 2640-3498.
- Mingyang Chen, Linzhuang Sun, Tianpeng Li, Haoze Sun, Yijie Zhou, 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](#). *Preprint*, arXiv:2503.19470.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. [Mem0: Building Production-Ready AI Agents with Scalable Long-Term Memory](#). *Preprint*, arXiv:2504.19413.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-Augmented Generation for Large Language Models: A Survey](#). *arXiv preprint*. ArXiv:2312.10997 [cs].
- Chia-Tung Ho, Haoxing Ren, and Brucek Khailany. 2025. [Verilogcoder: Autonomous verilog coding agents with graph-based planning and abstract syntax tree \(ast\)-based waveform tracing tool](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 300–307. Issue: 1.
- 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 [cs].
- Sirui Hong, Yizhang Lin, Bang Liu, Bangbang Liu, Binhao Wu, Ceyao Zhang, Danyang Li, Jiaqi Chen, Jiayi Zhang, and Jinlin Wang. 2025. [Data interpreter: An Llm agent for data science](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19796–19821.
- Md Ashraful Islam, Mohammed Eunus Ali, and Md Rizwan Parvez. 2024. [MapCoder: Multi-Agent Code Generation for Competitive Problem Solving](#). *arXiv preprint*. ArXiv:2405.11403 [cs].
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. [Search-R1: Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning](#). *Preprint*, arXiv:2503.09516.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension](#). *arXiv preprint*. ArXiv:1705.03551 [cs].
- Zilin Kang, Chonghua Liao, Tingqiang Xu, and Huazhe Xu. 2025. [Entropy regularizing activation: Boosting continuous control, large language models, and image classification with activation as entropy constraints](#). *arXiv preprint arXiv:2510.08549*.
- Kiseung Kim and Jay-Yoon Lee. 2024. [RE-RAG: Improving Open-Domain QA Performance and Interpretability with Relevance Estimator in Retrieval-Augmented Generation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22149–22161, Miami, Florida, USA. Association for Computational Linguistics.
- Satyapriya Krishna, Kalpesh Krishna, Anhad Mohanney, Steven Schwarcz, Adam Stambler, Shyam Upadhyay, and Manaal Faruqui. 2025. [Fact, Fetch, and Reason: A Unified Evaluation of Retrieval-Augmented Generation](#). *arXiv preprint*. ArXiv:2409.12941 [cs] version: 3.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural Questions: A Benchmark for Question Answering](#)

- Research.** *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. **Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.** In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Chonghua Liao, Ke Wang, Yuchuan Wu, Fei Huang, and Yongbin Li. 2025a. **Moa: Multi-objective alignment for role-playing agents.** *arXiv preprint arXiv:2512.09756*.
- Chonghua Liao, Ruobing Xie, Xingwu Sun, Haowen Sun, and Zhanhui Kang. 2025b. **Exploring forgetting in large language model pre-training.** In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2112–2127.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. **Lost in the Middle: How Language Models Use Long Contexts.** *arXiv preprint ArXiv:2307.03172* [cs].
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khoshdel, and Hannaneh Hajishirzi. 2023. **When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories.** *arXiv preprint ArXiv:2212.10511* [cs].
- Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2023. **GAIA: a benchmark for General AI Assistants.** *arXiv preprint ArXiv:2311.12983* [cs].
- Shuaiyi Nie, Siyu Ding, Wenyuan Zhang, Linhao Yu, Tianmeng Yang, Yao Chen, Tingwen Liu, Weichong Yin, Yu Sun, and Hua Wu. 2026. **Attnpo: Attention-guided process supervision for efficient reasoning.** *Preprint, arXiv:2602.09953*.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. **GPT-4 Technical Report.** *arXiv preprint ArXiv:2303.08774* [cs].
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. 2024. **MemGPT: Towards LLMs as Operating Systems.** *Preprint, arXiv:2310.08560*.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah 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.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. **DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models.** *Preprint, arXiv:2402.03300*.
- Qiaoyu Tang, Hao Xiang, Le Yu, Bowen Yu, Yaojie Lu, Xianpei Han, Le Sun, WenJuan Zhang, Pengbo Wang, Shixuan Liu, Zhenru Zhang, Jianhong Tu, Hongyu Lin, and Junyang Lin. 2025. **Beyond Turn Limits: Training Deep Search Agents with Dynamic Context Window.** *arXiv preprint ArXiv:2510.08276* [cs].
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. **MuSiQue: Multi-hop Questions via Single-hop Question Composition.** *arXiv preprint ArXiv:2108.00573* [cs].
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. **Attention is All you Need.** In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Guoqing Wang, Sunhao Dai, Guangze Ye, Zeyu Gan, Wei Yao, Yong Deng, Xiaofeng Wu, and Zhenzhe Ying. 2025. **Information Gain-based Policy Optimization: A Simple and Effective Approach for Multi-Turn LLM Agents.** *arXiv preprint ArXiv:2510.14967* [cs].
- Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and Fei Huang. 2025. **WebWalker: Benchmarking LLMs in Web Traversal.** *arXiv preprint ArXiv:2501.07572* [cs].
- Wujiang Xu, Kai Mei, Hang Gao, Juntao Tan, Zujie Liang, and Yongfeng Zhang. 2025. **A-MEM: Agentic Memory for LLM Agents.** *Preprint, arXiv:2502.12110*.
- 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.
- 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*.
- Kechi Zhang, Jia Li, Ge Li, Xianjie Shi, and Zhi Jin. 2024. **CodeAgent: Enhancing Code Generation with Tool-Integrated Agent Systems for Real-World Repo-level Coding Challenges.** *arXiv preprint ArXiv:2401.07339* [cs].

Wenlin Zhang, Xiaopeng Li, Yingyi Zhang, Pengyue Jia, Yichao Wang, Huifeng Guo, Yong Liu, and Xiangyu Zhao. 2025. *Deep Research: A Survey of Autonomous Research Agents*. *arXiv preprint ArXiv:2508.12752* [cs].

Wenyuan Zhang, Xinghua Zhang, Haiyang Yu, Shuaiyi Nie, Bingli Wu, Juwei Yue, Tingwen Liu, and Yongbin Li. 2026. *Expseek: Self-triggered experience seeking for web agents*. *arXiv preprint arXiv:2601.08605*.

Longtao Zheng, Rundong Wang, Xinrun Wang, and Bo An. 2024. *Synapse: Trajectory-as-Exemplar Prompting with Memory for Computer Control*. *Preprint*, arXiv:2306.07863.

Yuxiang Zheng, Dayuan Fu, Xiangkun Hu, Xiaojie Cai, Lyumanshan Ye, Pengrui Lu, and Pengfei Liu. 2025. *DeepResearcher: Scaling Deep Research via Reinforcement Learning in Real-world Environments*. *arXiv preprint*. ArXiv:2504.03160 [cs].

Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2023. *MemoryBank: Enhancing Large Language Models with Long-Term Memory*. *Preprint*, arXiv:2305.10250.

Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Pu Liang. 2025. *MEM1: Learning to Synergize Memory and Reasoning for Efficient Long-Horizon Agents*. *Preprint*, arXiv:2506.15841.

## A Appendix

### A.1 Single-Objective Tasks

To assess the effectiveness of our approach on single-objective settings, we conduct experiments on seven question answering benchmarks: 2WikiMultiHopQA, HotpotQA, Bamboogle, Musique, Natural Questions (NQ), TriviaQA, and PopQA (Ho et al., 2020; Yang et al., 2018; Press et al., 2023; Trivedi et al., 2022; Kwiatkowski et al., 2019; Joshi et al., 2017; Mallen et al., 2023). These datasets span diverse domains and are widely used in prior agent-oriented research. For datasets with more than 1k samples, we randomly sample 1k samples for evaluation.

**Datasets.** The following are details of these datasets:

- **Natural Questions (NQ):** a QA dataset whose questions are derived from real anonymized and aggregated queries issued to the Google Search engine.
- **TriviaQA:** a large-scale dataset with compositional questions that often require non-trivial reasoning.

- **PopQA:** 14K questions focusing on long-tail factual knowledge.
- **Bamboogle:** a manually constructed multi-hop QA benchmark in which questions are designed to be difficult to answer with a single search engine call.
- **Musique:** a 25K-question multi-hop QA dataset requiring evidence composition across multiple facts.
- **HotpotQA:** a Wikipedia-based multi-hop dataset where answering requires retrieving and reasoning over multiple supporting documents.
- **2WikiMultiHopQA:** a multi-hop QA dataset combining structured and unstructured evidence, explicitly constructed to necessitate multi-hop reasoning.

**Results.** As shown in Table 3 and Table 4, our method achieves strong performance across all benchmarks. On several datasets (e.g., TriviaQA), it surpasses all baselines and reaches SOTA performance, while substantially reducing token consumption. These results suggest that our method can significantly improve performance on long-horizon tasks while maintaining competitiveness on short-horizon tasks, matching or even exceeding agent models trained specifically for single-objective settings.

### A.2 Deep Research Tasks

We further evaluate our approach on three deep research benchmarks: GAIA (Mialon et al., 2023), Frames (Krishna et al., 2025), and WebWalkerQA (Wu et al., 2025). In contrast to the above datasets, deep research questions are often constructed from real web search results, and thus are fully out-of-domain (OOD) relative to our training setting. This evaluation tests whether our method generalizes to complex OOD tasks. To reduce evaluation cost, for the larger benchmarks (Frames and WebWalkerQA), we randomly sample 250 instances.

**Datasets.**

- **GAIA:** a collection of 165 tasks spanning three difficulty levels (53 Level-1, 86 Level-2, and 26 Level-3), designed to measure tool use and multi-step reasoning.

Table 2: The token consumption for multi-objective tasks of baselines. **Text with bold** means SOTA.

Local Wiki Search												
Model	2-objective		4-objective		6-objective		8-objective		10-objective		Avg	
	TT	PT	TT	PT	TT	PT	TT	PT	TT	PT	TT	PT
Qwen2.5 (ReAct)	1.90	0.38	2.80	0.51	3.97	0.63	4.89	0.76	4.66	0.77	3.64	0.61
Research	0.45	0.26	1.58	0.52	3.18	0.76	4.63	0.91	6.62	1.11	3.29	0.71
DeepResearcher	0.96	0.31	2.60	0.61	4.13	0.79	5.48	0.93	8.26	1.19	4.29	0.77
A-MEM	1.14	0.34	2.16	0.38	2.61	0.38	3.49	0.41	3.69	0.40	2.62	0.38
MEM1	0.50	0.16	0.88	0.18	1.31	0.20	1.81	0.22	2.40	0.24	1.38	0.20
GRPO (no mem)	0.45	0.26	2.29	0.62	3.91	0.84	6.32	1.07	8.96	1.28	4.39	0.81
Ours	<b>0.32</b>	<b>0.14</b>	<b>0.80</b>	<b>0.17</b>	<b>1.22</b>	<b>0.19</b>	<b>1.61</b>	<b>0.20</b>	<b>1.94</b>	<b>0.21</b>	<b>1.18</b>	<b>0.18</b>

Online Web Search												
Model	2-objective		4-objective		6-objective		8-objective		10-objective		Avg	
	TT	PT	TT	PT	TT	PT	TT	PT	TT	PT	TT	PT
Qwen2.5 (ReAct)	0.24	0.13	4.53	0.36	2.63	0.34	3.21	0.40	5.08	0.48	3.14	0.34
Research	0.27	0.14	1.04	0.29	1.94	0.41	3.13	0.52	4.48	0.62	2.17	0.40
MEM1	0.31	0.10	0.57	0.12	0.89	0.14	1.30	0.16	1.74	0.19	0.96	0.14
Ours	<b>0.19</b>	<b>0.08</b>	<b>0.56</b>	<b>0.11</b>	<b>0.85</b>	<b>0.12</b>	<b>1.17</b>	<b>0.14</b>	<b>1.53</b>	<b>0.15</b>	<b>0.86</b>	<b>0.12</b>

Table 3: The results of multi-hop QA datasets. They are the first part of single-objective QA datasets.

Multi-hop QA																				
Model	2WikiMultiHopQA				Bamboogle				HotpotQA				Musique				Avg			
	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT
qwen2.5 (ReAct)	45.21	35.90	1.12	0.29	43.70	32.80	0.95	0.24	47.12	34.30	1.07	0.26	23.86	14.40	2.06	0.37	39.97	29.35	1.30	0.29
ReSearch	50.07	41.90	0.55	0.28	53.61	40.80	0.41	0.24	50.26	33.90	0.40	0.23	29.63	17.80	0.58	0.29	45.89	33.60	0.48	0.26
DeepResearcher	51.44	43.90	1.24	0.34	48.48	37.60	1.22	0.30	51.96	38.40	1.01	0.29	26.55	17.00	2.54	0.44	44.61	34.23	1.50	0.34
GRPO (no mem)	62.35	53.80	0.93	0.32	53.18	40.80	0.79	0.28	57.29	42.30	0.62	0.27	33.78	21.50	0.86	0.33	51.65	39.60	0.80	0.30
Ours	59.17	50.20	0.37	0.15	52.90	36.80	0.29	0.14	57.64	42.90	0.34	0.15	33.48	22.10	0.39	0.15	50.80	38.00	0.35	0.15

Table 4: The results of single-hop QA datasets. They are the second part of single-objective QA datasets.

Single-hop QA																
Model	NQ				PopQA				TriviaQA				Avg			
	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT
qwen2.5 (ReAct)	49.18	36.30	1.87	0.30	47.77	40.50	1.43	0.26	62.73	52.60	1.02	0.22	53.23	43.13	1.44	0.26
ReSearch	52.49	38.00	0.23	0.16	52.41	43.80	0.23	0.16	62.04	50.10	0.24	0.17	55.65	43.97	0.23	0.16
DeepResearcher	50.15	39.50	1.19	0.25	47.90	40.80	0.60	0.22	61.89	52.30	0.95	0.23	53.32	44.20	0.91	0.23
GRPO (no mem)	56.04	43.90	0.63	0.23	53.77	46.50	1.00	0.26	64.57	54.50	0.52	0.23	58.13	48.30	0.72	0.24
Ours	57.46	46.10	0.22	0.14	53.53	46.70	0.25	0.14	67.83	57.10	0.25	0.14	59.61	49.97	0.24	0.14

Table 5: The results of deep research datasets.

Deep Research																
Model	GAIA				Frames				WebWalker				Avg			
	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT	F1	EM	TT	PT
qwen2.5 (ReAct)	14.28	8.74	1.54	0.22	29.37	20.00	0.61	0.15	29.93	8.50	0.42	0.14	24.52	12.41	0.86	0.17
ReSearch	16.00	8.74	0.41	0.17	35.04	22.40	0.37	0.16	33.02	7.29	0.25	0.13	28.02	12.81	0.34	0.15
DeepResearcher	22.44	16.50	4.94	0.33	37.69	24.40	0.44	0.17	32.10	9.31	1.16	0.19	30.75	16.74	2.18	0.23
Ours	25.07	17.48	0.62	0.12	37.56	24.00	0.33	0.10	31.85	8.91	0.29	0.09	31.49	16.79	0.41	0.10

- **Frames:** measures multi-perspective reasoning and role-conditioned information synthesis, requiring consistent integration of evidence across different contextual frames.
- **WebWalkerQA:** evaluates complex, multi-turn web interaction, consisting of 680 real-world queries across four domains and over 1,373 webpages.

**Results.** As reported in Table 5, in terms of average accuracy, our method achieves performance comparable to DeepResearcher while significantly reducing token usage. Notably, DeepResearcher is trained specifically in real web search environments. Moreover, on longer-horizon tasks with higher token demands (e.g., GAIA), our method delivers relatively strong performance. Overall, these findings indicate that our approach generalizes well to out-of-domain settings while still demonstrating solid capability on long-horizon reasoning tasks.

Table 6: Training prompt text with cumulative summary memory.

---

---

**Prompt:**

You will answer multiple complex questions using iterative reasoning, summarization, and web search. At each step, you will see the questions, a cumulative summary of relevant information, the current search query, and search results (except in the first step, where only the questions are provided). Your task is to:

1. Update a cumulative, concise summary and perform reasoning within `<mem>\n...\n</mem>\n<think>\n...\n</think>`, respectively. `<mem>\n...\n</mem>` must include all essential information from previous `<think>` and `<tool_response>` tags.

2. Then choose one of the following actions:

- If any question remains unanswered, issue a single query for one question inside `<search>\n...\n</search>`.
- If all questions are answered, provide the final answers separated by semicolons within `<answer>\n answer1; answer2; ... \n</answer>`. The answers must be concise, usually short phrases or words, and avoid any explanations.

Important:

- The summaries only contain concise and essential information related to the user questions from previous trajectories, but do not include any your own internal knowledge or reasoning content.
  - Must strictly follow one of these two structures: `<mem>\n...\n</mem>\n<think>\n...\n</think>\n<search>\n...\n</search>` or `<mem>\n...\n</mem>\n<think>\n...\n</think>\n<answer>\n...\n</answer>`.
  - Do not search multiple queries or questions simultaneously. Don't give up searching for information until you find clear information that provides the answer.
- 
-

Table 7: Case study of our method.

<p><b>Query</b></p> <p>As of Aug 3, 2024, the artist who released the album "Father of Asahd" went to the same high school as an Olympic diver. How many Olympic teams did this diver participate on?</p>
<p><b>Answer</b></p> <p>2</p>
<p><b>Step 1</b></p> <p><b>Assistant:</b></p> <p>&lt;mem&gt;Need to identify the artist, the diver, and the number of Olympic teams.&lt;/mem&gt;          &lt;think&gt;To answer the question, I first need to find the artist of "Father of Asahd" ...&lt;/think&gt;          &lt;search&gt;Who released the album "Father of Asahd"&lt;/search&gt;</p> <p><b>User:</b></p> <p>&lt;information&gt; [Doc 1-5] ... The album "Father of Asahd" is the 11th studio album by American DJ and producer DJ Khaled, released in May 2019 ... &lt;/information&gt;</p>
<p><b>Step 2</b></p> <p><b>Assistant:</b></p> <p>&lt;mem&gt;DJ Khaled released "Father of Asahd". He attended the same high school as an Olympic diver.&lt;/mem&gt;          &lt;think&gt;Next, I need to confirm DJ Khaled's high school to find the related Olympic diver ...&lt;/think&gt;          &lt;search&gt;DJ Khaled high school&lt;/search&gt;</p> <p><b>User:</b></p> <p>&lt;information&gt; [Doc 1-5] ... DJ Khaled graduated from Dr. Phillips High School ... &lt;/information&gt;</p>
<p><b>Step 3</b></p> <p><b>Assistant:</b></p> <p>&lt;mem&gt;DJ Khaled graduated from Dr. Phillips High School. Need to find the Olympic diver from this school.&lt;/mem&gt;          &lt;think&gt;I will search for Olympic divers who attended Dr. Phillips High School ...&lt;/think&gt;          &lt;search&gt;Olympic diver who attended Dr. Phillips High School&lt;/search&gt;</p> <p><b>User:</b></p> <p>&lt;information&gt; [Doc 1-5] ... Mark Ruiz, a graduate of Dr. Phillips High School, was a member of the US Olympic diving team at the 2000 and 2004 Summer Olympics ... &lt;/information&gt;</p>
<p><b>Step 4</b></p> <p><b>Assistant:</b></p> <p>&lt;mem&gt;Olympic diver Mark Ruiz attended Dr. Phillips High School with DJ Khaled. He joined the 2000 and 2004 US Olympic diving teams.&lt;/mem&gt;          &lt;think&gt;Mark Ruiz participated in two Olympic Games (2000 and 2004), so he was on 2 Olympic teams ...&lt;/think&gt;          &lt;answer&gt;2&lt;/answer&gt;</p>

Table 8: Case study of ReAct. It was failed because of the Hallucination in Step 3. Greg Louganis is also a diver, but he is not mentioned anywhere in the context.

<b>Query</b>
As of Aug 3, 2024, the artist who released the album "Father of Asahd" went to the same high school as an Olympic diver. How many Olympic teams did this diver participate on?
<b>Answer</b>
2
<b>Step 1</b>
<b>Assistant:</b>
<think>First, identify the artist who released the album "Father of Asahd" ...</think>
<search> Who is the artist that released the album "Father of Asahd"? </search>
<b>User:</b>
<information> [Doc 1-5] ... Father of Asahd is the eleventh studio album by American DJ and producer DJ Khaled, released on May 17, 2019 ... </information>
<b>Step 2</b>
<b>Assistant:</b>
<think>The artist who released the album "Father of Asahd" is DJ Khaled ...</think>
<search> Which high school did DJ Khaled go to? </search>
<b>User:</b>
<information> [Doc 1-5] ... DJ Khaled graduated from Dr. Phillips High School and also visited New Rochelle High School for a graduation event ... </information>
<b>Step 3</b>
<b>Assistant:</b>
<think>DJ Khaled graduated from Dr. Phillips High School ...</think>
<search> Which high school did Greg Louganis go to? </search>
<b>User:</b>
<information> [Doc 1-5] ... Greg Louganis attended Santa Ana High School, Valhalla High School, and Mission Viejo High School in California ... </information>
<b>Step 4</b>
<b>Assistant:</b>
<think>Greg Louganis attended Santa Ana High School, Valhalla High School, and Mission Viejo High School. He went to Valhalla High School ...</think>
<search> How many Olympic teams did Greg Louganis participate in? </search>
<b>User:</b>
<information> ... [Doc 1-5] Greg Louganis is a four-time Olympian, five-time Olympic medalist and four-time Olympic gold medalist diver ... </information>
<b>Step 5</b>
<b>Assistant:</b>
<think> Greg Louganis was a four-time Olympian ... </think>
<answer> 4 </answer>