

# WebClipper: Efficient Evolution of Web Agents with Graph-based Trajectory Pruning

Junjie Wang, Zequn Xie, Dan Yang, Jie Feng, Yue Shen, Duolin Sun, Meixiu Long, Yihan Jiao, Zhehao Tan, Jian Wang, Peng Wei, Jinjie Gu\*

Ant Group

Correspondence: wjj417805@antgroup.com, wangjj2018@zju.edu.cn, jinjie.gujj@antgroup.com

## Abstract

Deep Research systems based on web agents have shown strong potential in solving complex information-seeking tasks, yet their search efficiency remains underexplored. We observe that many state-of-the-art open-source web agents rely on long tool-call trajectories with cyclic reasoning loops and exploration of unproductive branches. To address this, we propose WebClipper, a framework that compresses web agent trajectories via graph-based pruning. Concretely, we model the agent’s search process as a state graph and cast trajectory optimization as a minimum-necessary Directed Acyclic Graph (DAG) mining problem, yielding pruned trajectories that preserve essential reasoning while eliminating redundant steps. Continued training on these refined trajectories enables the agent to evolve toward more efficient search patterns and reduces tool-call rounds by about 20% while improving accuracy. Furthermore, we introduce a new metric called F-AE Score to measure the model’s overall performance in balancing accuracy and efficiency. Experiments demonstrate that WebClipper compresses tool-call rounds under excellent performance, providing practical insight into balancing effectiveness and efficiency in web agent design. Code is available at <https://github.com/AQ-MedAI/AntAFu-DeepResearch>.

## 1 Introduction

With the continuous evolution of Large Language Models (LLMs), artificial intelligence systems have transformed from static text-based models into sophisticated agents capable of utilizing tools and interacting with environments (Bai et al., 2025b; Zeng et al., 2025; Xiong et al., 2026). Among these, web agents have demonstrated remarkable capabilities in complex information-seeking, completing challenging tasks in tens

\*Corresponding author

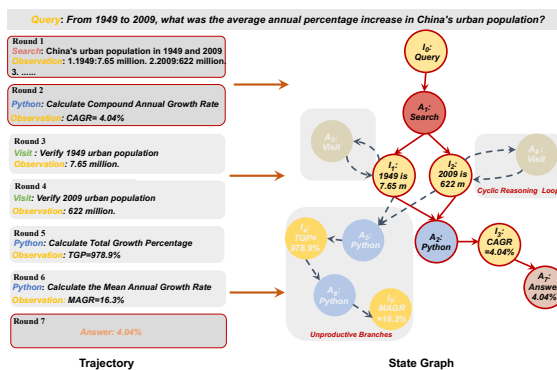


Figure 1: The trajectory of a web agent can be built as a graph. The minimum number of steps to solve the problem is the minimum necessary DAG from the Query node (I0) to the final Answer Action node (A7).

of minutes that would typically require humans several hours. Representative examples include commercial systems such as OpenAI’s Deep Research (OpenAI, 2025a), Gemini (Gemini Team, 2025), and Claude (Claude Team, 2025), alongside emerging open-source alternatives like Tongyi-DeepResearch (Li et al., 2025a) and MiroThinker (Bai et al., 2025a).

However, current open-sourced web agents primarily focus on the final problem-solving accuracy while paying little attention to efficiency during the search process. In pursuit of higher accuracy, these agents continuously scale up search depth and context length (Li et al., 2025a), leading to extremely long contexts and excessive tool usage. For example, Tongyi-DeepResearch uses a 128K context length and up to 100 tool-call rounds, while MiroThinker sets a maximum context length of 256K and allows up to 600 tool-call rounds. Considering the long inference time and the high costs of commercial search tools (e.g., Google Search and Jina Reader), the user experience in practice is far from ideal.

To understand what causes such inefficiency, we conduct a deeper analysis of the agent’s search

behavior. Prior work (Yen et al., 2025; Tao et al., 2025a) highlighted that effective actions are sparsely distributed across long trajectories. For many failure cases, the agent repeatedly re-searches information it has already obtained or over-focuses on noisy signals (Yen et al., 2025), causing it to drift away from the correct direction, which should ideally be avoided. To systematically identify such inefficiency patterns, we model the trajectory of the agent as a state graph. As is illustrated in Figure 1, the agent’s action and environmental observation can be abstracted as nodes in the graph. This formalization reveals two major inefficient patterns: **cyclic reasoning loops** and **unproductive branches** that diverge from the correct solution, while the ideal path should be the minimum DAG from the original query to the final answer.

The above observation motivates us to prune these inefficient patterns to construct a more robust web agent. However, training a robust web agent from scratch remains both costly and challenging due to complex data synthesis pipelines and multi-stage training paradigms (Li et al., 2025a; Hu et al., 2025) that range from agentic mid-training to SFT and RL. This leads us to explore a different direction: **Instead of building a new agent from scratch, can we evolve pre-existing, high-performance but low-efficiency web agents into more efficient ones by pruning their inefficient patterns?**

To achieve this, we introduce WebClipper, a novel framework designed to optimize the search behavior of web agents toward a better accuracy-efficiency balance. Specifically, our framework consists of: 1) Trajectory to State-Graph transformation: transforming raw trajectories into state graphs by abstracting agent actions and environment information. 2) Pruning via an approximate minimal necessary DAG (MNDAG): mining an approximate MNDAG that connects initial information nodes to final action nodes, thereby pruning redundant steps. 3) Coherence-aware thought rewriting: rewriting the agent’s thoughts on the pruned trajectories to ensure semantic consistency and usability. 4) Agent Evolution: training existing agents to improve efficiency based on collected trajectories combined with a hybrid evolution strategy. To quantify the accuracy-efficiency trade-off, we further propose a new evaluation metric, **F-AE Score**. Instead of separately reporting performance and resource usage, the F-AE Score reflects how well a web agent balances these two aspects, pro-

viding a direct view for comparing different optimization strategies and guiding the design of more practical web agents.

Experiments on multiple benchmarks show that WebClipper reduces tool-call rounds and token usage by about 20% while maintaining or even improving accuracy. Our contributions are summarized as follows:

- 1) We propose WebClipper, a novel pruning method for existing Deep Research-style web agents, enabling them to evolve toward a more efficient search behavior.
- 2) Our methods explicitly target the accuracy-efficiency trade-off, together with the F-AE score as a unified metric to evaluate this balance.
- 3) We evaluate WebClipper on multiple benchmarks and empirically demonstrate its good balance between accuracy and efficiency.

## 2 Related Work

**Deep Research Agents.** Methods for web agents can be broadly divided into two categories. The first is training-free approaches, which solve tasks by designing multi-agent collaborative architectures, such as OpenDeepResearch (Research, 2025b), GPT Researcher (Research, 2025a), and WebWeaver (Li et al., 2025d). These works typically focus on how to structure the agent state space, using context engineering to compress and share context across agents so that they perform better on long-horizon, complex tasks. The second category is training-based approaches, which aim to train a single, powerful core agent that can flexibly use various tools within a constructed environment. To obtain such agents, a large body of work focuses on synthesizing training data for web agents, generating complex multi-hop questions from open webpages or knowledge graphs (Li et al., 2025b; Tao et al., 2025b; Wang et al., 2024), and then applying SFT or RL to improve the agent’s capability on challenging tasks (Liu et al., 2025; Li et al., 2025c; Chen et al., 2025). However, these methods almost exclusively target end-to-end task success rates, while paying very little attention to the efficiency of web agents.

**Efficient Reasoning in LLMs.** With the emergence of reasoning models such as OpenAI-o1 (OpenAI, 2024) and DeepSeek-R1 (Guo et al., 2025), there has been growing interest in efficient reasoning for single LLMs. A simple yet effective line of work is prompt-based, where explicit

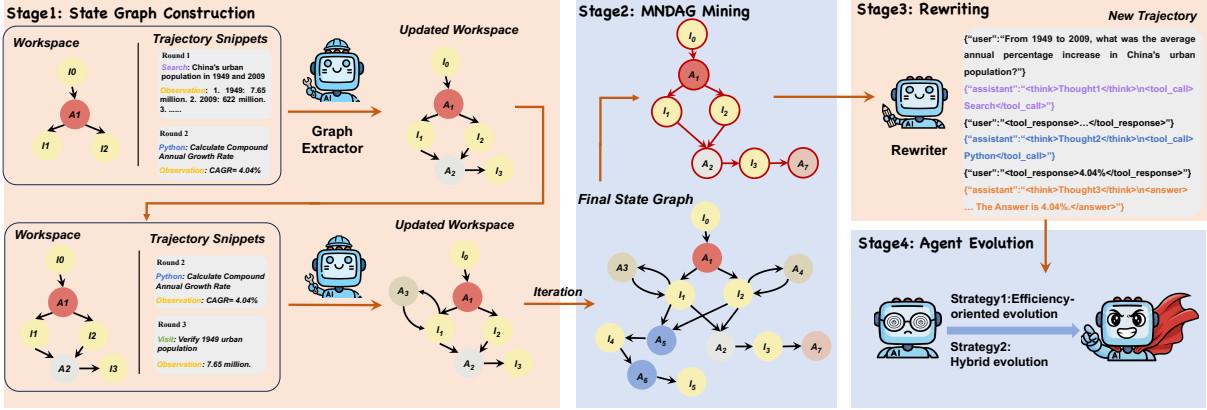


Figure 2: The overview of WebClipper

instructions are added to the prompt to encourage the model to reason in a more efficient manner (Han et al., 2025; Xu et al., 2025; Poddar et al., 2025). Beyond prompting, many methods rely on training-based strategies: for example, compressing the long chain-of-thoughts (CoT) into shorter ones to train a model that acquires short-thinking capabilities and maintains performance under low-resource settings (Ma et al., 2025; Munkhbat et al., 2025; Cui et al., 2025); or incorporating length-related rewards into RL training so that the model learns to discover more efficient reasoning paths (Luo et al., 2025; Aggarwal and Welleck, 2025; Dumitru et al., 2025). These compression techniques for single models inspire our design of methods to improve the search efficiency of web agents.

### 3 Methodology

In this section, we present **WebClipper**, a framework to evolve an existing Deep Research-style web agent into a more efficient one. As shown in Figure 2, our framework consists of four main components: (1) constructing state graphs from raw trajectories, (2) mining an approximate MNDAG for pruning, (3) coherence-aware thought rewriting, followed by (4) agent evolution based on the pruned trajectories.

#### 3.1 Preliminaries and Notation

Let a query be denoted by  $q$ . Given  $q$ , a web agent interacts with the environment through a trajectory:

$$\tau = (o_0, r_1, a_1, o_1, \dots, r_T, a_T),$$

where  $o_t$  is the observation from the environment at round  $k$  (with  $o_0 = q$ ),  $r_t$  is the agent’s thought, and  $a_t$  is the agent’s action. Actions include tool invocations (e.g., SEARCH, VISIT, PYTHON) and

the final answer ANSWER. Our goal is to transform each raw trajectory  $\tau$  samples from raw agent  $\mathcal{M}$  into an accurate and efficient trajectory  $\tilde{\tau}$ , and then use a collection of such trajectories to train a model  $\mathcal{M}'$  that achieves comparable accuracy with fewer action steps.

#### 3.2 Initial Trajectory Collection and Filtering

We first collect question–answer (QA) pairs from public datasets such as WebShaper (Tao et al., 2025b), WebDancer (Wu et al., 2025), WebExplorer (Liu et al., 2025), TaskCraft (Shi et al., 2025), and Voyager (Bai et al., 2025a). Using the pre-built environment, we distill trajectories from the existing web agent  $\mathcal{M}$ , which follows a ReAct-style (Yao et al., 2023) loop of observation–think–action.

For each  $q$ , we first sample  $K$  distinct trajectories  $\{\tau^{(k)}\}_{k=1}^K$  from  $\mathcal{M}$ . We then employ a rejection sampling strategy: let  $\text{PR}(q) \in [0, 1]$  denote the pass rate of  $q$  on the target task, we retain only trajectories of queries satisfying:  $0 < \text{PR}(q) \leq 0.5$ , which keeps the task challenging. These trajectories constitute the input to our pruning pipeline.

#### 3.3 From Trajectory to State Graph

##### 3.3.1 State Graph Definition

Given a trajectory  $\tau$ , we construct a directed graph  $\mathcal{G} = (\mathcal{V}^A \cup \mathcal{V}^I, \mathcal{E})$ , where  $\mathcal{V}^A = \{A_1, \dots, A_T\}$  is the set of **Action nodes**. Each  $A_t$  abstracts the agent’s thought and action at step  $t$ .  $\mathcal{V}^I = \{I_0, I_1, \dots\}$  is the set of **Information nodes**, representing atomic pieces of information obtained from the environment, including the initial query.

We denote the initial query node as  $I_0$ , and the final answer node as  $A_T$ . Edges  $\mathcal{E}$  capture the dependency between actions and information:  $I \rightarrow A$  if action  $A$  is taken based on information  $I$ ;  $A \rightarrow I$

if information  $I$  is produced as a result of action  $A$ . This yields a bipartite, directed structure between  $\mathcal{V}^A$  and  $\mathcal{V}^I$ .

### 3.3.2 State Graph Construction

We construct  $\mathcal{G}$  from  $\tau$  with an LLM-based extractor. First, for each step  $t$  with internal thought  $r_t$  and action  $a_t$ , the extractor summarizes  $(r_t, a_t)$  into a compact Action node  $A_t$  (recording action type and goal), yielding  $\{A_t\}_{t=1}^T$ . We then build information nodes  $\mathcal{V}^I$  and edges  $\mathcal{E}$  iteratively using a workspace  $\mathcal{W}$  that stores current information nodes and links. Initially,  $\mathcal{W} = \{I_0\}$ , where  $I_0$  encodes the original query. For each step  $t = 0, \dots, T-1$ , we feed the snippet  $(A_t, o_t, A_{t+1})$  and  $\mathcal{W}$  to the extractor, prompting it to:

1) **Decompose observation into atomic information.**  $o_t$  is decomposed into atomic units  $\{I^*\}$ . Each  $I^*$  is matched against existing nodes in  $\mathcal{W}$ ; on a semantic match, we add  $A_t \rightarrow I$ ; otherwise we create a new information node  $I^*$ , insert it into  $\mathcal{V}^I$  and  $\mathcal{W}$ , and add  $A_t \rightarrow I^*$ .

2) **Link new action to supporting information.** The extractor analyzes  $A_{t+1}$  to identify a set of information nodes  $\mathcal{S}_k \subseteq \mathcal{V}^I$  in  $\mathcal{W}$  that the agent relies on when executing  $A_{t+1}$ . For each  $I \in \mathcal{S}_k$ , we add an edge  $I \rightarrow A_{t+1}$ .

This process continues until the final answer action  $A_T$  is reached. The result is a state graph  $\mathcal{G}$  that explicitly encodes the dependency between all actions and information along the trajectory.

### 3.4 Pruning via MNDAG

Given the state graph  $\mathcal{G}$ , we aim to identify the approximate minimal subgraph that is necessary and sufficient to support the final answer. Intuitively, actions that do not contribute any information used (directly or indirectly) by the answer are deemed redundant and should be pruned.

We treat the initial query node  $I_0$  as the source and the final answer node  $A_T$  as the sink. Each action node  $A_t$  is assigned a unit cost  $c(A_t) = 1$ , and each information node cost is set to zero, i.e.,  $c(I_t) = 0$ . Our objective is to find an approximate minimal-cost directed acyclic subgraph  $\mathcal{G}^*$  that connects  $I_0$  to  $A_T$  and preserves all necessary dependencies.

We approximate this by:

1) **Shortest-path forward search.** We run a Dijkstra-style shortest-path algorithm on  $\mathcal{G}$  from  $I_0$  to  $A_T$ , using node costs  $c(\cdot)$  aggregated along the path. This yields the shortest path  $P = (I_0 \rightarrow$

$\dots \rightarrow A_T)$ , which captures one minimal-cost path from query to answer.

2) **Backward closure of necessary predecessors.** Starting from  $A_T$ , we perform a reverse traversal on  $\mathcal{G}$ , recursively adding predecessor nodes that are on some shortest path contributing to the answer. This ensures that we do not miss necessary branching dependencies. The resulting set of nodes  $\mathcal{V}^* \subseteq \mathcal{V}^A \cup \mathcal{V}^I$  and edges  $\mathcal{E}^*$  form a MNDAG:  $\mathcal{G}^* = (\mathcal{V}^*, \mathcal{E}^*)$ . A detailed algorithm description of the MNDAG is expanded at Algorithm 1 in Appendix B.

All action nodes  $A_t \notin \mathcal{V}^*$  are considered redundant and will be removed from the trajectory, thus obtaining a necessary actions set  $\mathcal{A}^*$ . To improve robustness, we repeat the graph construction and MNDAG mining process three times for the same raw trajectory  $\tau$ , obtaining three candidate sets of necessary actions:  $\mathcal{A}^{*(1)}, \mathcal{A}^{*(2)}, \mathcal{A}^{*(3)}$ . We then perform a majority vote at the action set level. The final set of necessary actions,  $\mathcal{A}_{final}^*$ , is determined only if at least two of the three candidate sets are identical.

### 3.5 Coherence-aware Thought Rewriting

Directly removing intermediate steps from a trajectory may break the coherence of the ReAct loop. We therefore perform coherence-aware rewriting over the pruned trajectory via context-aware selective rewriting and perplexity-based selection.

Given  $\mathcal{A}_{final}^*$ , we map it back to a pruned trajectory by retrieving each selected thought-action pair and its following observation from  $\tau$ , yielding

$$\tilde{\tau} = (o_0^{\text{new}}, r_1^{\text{new}}, a_1^{\text{new}}, o_1^{\text{new}}, \dots, r_L^{\text{new}}, a_L^{\text{new}}),$$

where  $L \leq T$  and all actions  $a_t^{\text{new}}$  and thought  $r_t^{\text{new}}$  correspond to nodes in  $\mathcal{A}_{final}^*$ .

1) **Context-aware selective rewriting.** For consecutive snippets  $(r_t^{\text{new}}, a_t^{\text{new}}, o_t^{\text{new}}, r_{t+1}^{\text{new}}, a_{t+1}^{\text{new}})$ , if  $a_t^{\text{new}}$  and  $a_{t+1}^{\text{new}}$  were adjacent in the original trajectory, we keep them unchanged. Otherwise, we rewrite  $r_{t+1}^{\text{new}}$  with a rewriter LLM based on the full context, including the pruned intermediate steps, prompting the rewriter to maintain logical continuity and remove references to pruned observations in the  $r_{t+1}^{\text{new}}$ , obtaining the rewritten thought  $\hat{r}_{t+1}^{\text{new}}$ .

2) **Perplexity-based selection.** To align the rewritten thoughts  $\hat{r}_{t+1}^{\text{new}}$  with the base model’s intrinsic reasoning style, we generate three candidate rewrites and select the one with the lowest perplexity (PPL) as calculated by the base model  $\mathcal{M}$  itself.

This process ensures alignment with the model’s intrinsic reasoning style as much as possible. Finally, we obtain a set of high-quality pruned trajectories  $\mathcal{D}_{pruned} = \{\tilde{\tau}\}$ .

### 3.6 Agent Evolution via Efficient and Hybrid Training

After obtaining  $\mathcal{D}_{pruned} = \{\tilde{\tau}\}$ , we use them to further train the base  $\mathcal{M}$ , evolving it into more efficient search behavior.

We propose two evolution paradigms:

1) **Efficiency-oriented evolution:** Fine-tune  $\mathcal{M}$  solely on  $\mathcal{D}_{pruned}$  to maximize search efficiency:

$$\mathcal{L}_{eff} = - \sum_{\tilde{\tau} \in \mathcal{D}_{pruned}} \log P_{\mathcal{M}}(\tilde{\tau})$$

2) **Hybrid evolution:** To balance efficiency and accuracy, we construct a hybrid dataset  $\mathcal{D}_{hybrid} = \mathcal{D}_{pruned} \cup \mathcal{D}_{unpruned}$ , where  $\mathcal{D}_{unpruned}$  contains unpruned trajectories with different queries (non-overlapping with  $\mathcal{D}_{pruned}$ ) and similar difficulty ( $0 < PR(q) \leq 0.5$ ). Trajectories in  $\mathcal{D}_{unpruned}$  are those where our MNDAG extraction finds no redundant rounds to prune. They have average longer steps than  $\mathcal{D}_{pruned}$ , but still provide valuable training signals for improving accuracy on complex queries. The training objective is:

$$\mathcal{L}_{hybrid} = - \sum_{\tau^* \in \mathcal{D}_{hybrid}} \log P_{\mathcal{M}}(\tau^*)$$

This strategy allows the model to learn efficient search patterns while retaining the capability to handle complex queries requiring longer but necessary reasoning chains, achieving an optimal trade-off between efficiency and accuracy.

## 4 Experiments

### 4.1 Experimental Settings

**Evaluation Metrics.** We evaluate web agents from three perspectives:

1) Accuracy: Accuracy (Acc) measured using LLM-as-Judge with o3-mini (OpenAI, 2025b) as the evaluator.

2) Efficiency: Tool-call rounds and token consumption during inference.

3) F-AE Score: Inspired by the F1 score (Hand et al., 2021), we propose **F-AE Score** to measure an agent’s ability to balance accuracy and efficiency:

$$F-AE = 2 \times \frac{Acc \times \left(1 - \frac{Rounds}{Max\_Rounds}\right)}{Acc + \left(1 - \frac{Rounds}{Max\_Rounds}\right)},$$

where Max\_Rounds is the maximum number of tool calls allowed in the experiment. Following common practice (Li et al., 2025a), we set Max\_Rounds = 100. F-AE penalizes both low accuracy and excessive tool usage, thereby avoiding over-optimization of either dimension alone. More explanations can be found in Appendix A.

**Datasets.** We conduct evaluations on four widely-used web agent benchmarks: xbench-deepsearch (Xbench Team, 2025), Browsecomp (Wei et al., 2025), GAIA (Mialon et al., 2023), and HLE (Phan et al., 2025). For GAIA, we use the 103 text-only subset from its development set. For HLE, we follow the setup of previous studies (Li et al., 2025c) and use a 500 text-only subset.

**Baselines.** Our comparison includes both closed-source and open-source agents. Closed-source systems include OpenAI o3 (OpenAI, 2025b), OpenAI DeepResearch (OpenAI, 2025a), and Claude-4-Sonnet (anthropic, 2025); test results are cited from their official reports. The open-source agents include Kimi-K2-Instruct-0905 (Bai et al., 2025b), Qwen3-235B-A22B-Instruct-2507 (Yang et al., 2025), WebExplorer (Liu et al., 2025) and Tongyi-DeepResearch. As trajectory pruning is underexplored, we design two baselines: 1) *Prompt Control*: We add instructions to the agent’s system prompt, explicitly asking it to avoid irrelevant information and repetitive validation, and to control the number of tool calls. 2) *Coarse Prune*: We use Qwen3-235B-A22B-Instruct-2507 to directly identify and remove turns from the trajectory that it deems redundant. The resulting coarsely pruned trajectories are then used for SFT.

**Implementation.** We use Tongyi-DeepResearch (30B-A3B) (Li et al., 2025a) as the base web agent  $\mathcal{M}$ . Trajectories are distilled from public QA datasets, including WebShaper, WebDancer, WebExplorer, TaskCraft, and Voyager. We adopt Qwen3-235B-A22B-Instruct-2507 as the extractor and rewriting model for state graph construction and thought rewriting. The extractor and rewriting model is deployed on 8×H800 (80GB) GPUs. Under this setup, the overall extraction and rewriting can be completed within 1 day. Training is conducted on 32×H800 GPUs with a learning rate of 5e-6 and a cosine decay schedule. For WebExplorer, we reproduce its results ourselves. For other open-source models that do not report tool and token usage, we reproduce them by deploying on H800 GPUs within the Tongyi-DeepResearch environment. For web content retrieval, we use

Method	xbench-deepsearch				Browsecomp			
	Acc ↑	F-AE ↑	Rounds ↓	Token ↓	Acc ↑	F-AE ↑	Rounds ↓	Token ↓
<i>Close-sourced System</i>								
OpenAI o3	0.670	-	-	-	0.497	-	-	-
OpenAI DeepResearch	-	-	-	-	0.515	-	-	-
Claude-4-Sonnet	0.646	-	-	-	0.122	-	-	-
<i>Open-sourced Agent</i>								
Kimi-K2-Instruct-0905*	0.540	0.686	5.98	1316	0.094	0.169	16.65	3426
Qwen3-235B-A22B-Instruct-2507*	0.490	0.637	8.84	938	0.046	0.087	13.70	1837
WebExplorer-8B	0.517	0.659	9.05	2246	0.137	0.229	29.43	6289
Tongyi-DeepResearch*	<u>0.713</u>	0.779	14.26	6918	0.410	0.385	63.70	12014
<i>Pruning Method(vs. Tongyi-DeepResearch)</i>								
Prompt Control	0.676	0.763	12.50	6321	0.373	0.372	62.80	12222
Coarse Prune	0.603	0.725	8.85	4774	0.220	0.326	37.10	8365
WebClipper (Eff)	<u>0.713</u>	<u>0.792</u>	10.81	5931	<u>0.427</u>	<b>0.431</b>	56.50	10599
WebClipper (Hybrid)	<b>0.733</b>	<b>0.797</b>	12.57	6205	<b>0.467</b>	<u>0.428</u>	60.42	11507
Method	GAIA				HLE			
	Acc ↑	F-AE ↑	Rounds ↓	Token ↓	Acc ↑	F-AE ↑	Rounds ↓	Token ↓
<i>Close-sourced System</i>								
OpenAI o3	-	-	-	-	0.249	-	-	-
OpenAI DeepResearch	0.674	-	-	-	0.266	-	-	-
Claude-4-Sonnet	0.683	-	-	-	0.203	-	-	-
<i>Open-sourced Agent</i>								
Kimi-K2-Instruct-0905*	0.469	0.625	6.45	1281	0.146	0.253	5.17	2349
Qwen3-235B-A22B-Instruct-2507*	0.456	0.612	7.14	1128	0.199	0.327	7.45	2960
WebExplorer-8B	0.372	0.521	12.88	3560	0.116	0.203	15.52	6579
Tongyi-DeepResearch*	0.682	0.733	20.56	7378	<u>0.358</u>	0.487	23.92	13664
<i>Pruning Method (vs. Tongyi-DeepResearch)</i>								
Prompt Control	0.663	0.730	18.70	6752	0.349	0.479	23.91	14107
Coarse Prune	0.514	0.638	15.60	4068	0.327	0.467	18.03	11851
WebClipper (Eff)	<u>0.684</u>	<b>0.760</b>	14.44	4756	0.353	<u>0.492</u>	18.60	11458
WebClipper (Hybrid)	<b>0.695</b>	<u>0.744</u>	19.92	6635	<b>0.361</b>	<b>0.495</b>	21.07	13532

Table 1: Performance comparison across various web agent benchmarks. The comparison for best (**bold**) and second-best (underline) results is conducted between the base model (Tongyi-DeepResearch) and the Pruning Methods, highlighted in light blue. The ↑ arrow indicates that higher values are better, while ↓ indicates lower values are better. \* denotes the result is conducted by ourselves in a unified environment.

the Serper API (SerpAPI, 2025) for search and Jina Reader (Jina.ai, 2025) for URL parsing. To reduce evaluation variance, each model is run three times with different random seeds, and we report the average Pass@1 and corresponding efficiency metrics.

## 4.2 Main Results

We organize our experimental investigation around four research questions (RQ):

**RQ1:** Is WebClipper an effective pruning strategy?

**RQ2:** How does WebClipper compare with direct pruning approaches?

**RQ3:** How well does the F-AE Score balance the accuracy-efficiency trade-off in web agents?

**RQ4:** Are the key components of WebClipper effective?

**Overall Performance (RQ1).** Tables 1 present the main results. We highlight several **key observations:** **1)** WebClipper(Eff) achieves leading performance among open-source models while reducing resource consumption. Compared to the Tongyi-DeepResearch baseline, it reduces token usage by **19.4%** and tool-call rounds by **21%** on average across all benchmarks, while maintaining comparable or even superior accuracy. This demonstrates the effectiveness of efficiency-oriented training in preserving task accuracy while significantly improving search efficiency. **2)** WebClipper(Hybrid) further improves accuracy with acceptable resource consumption. It achieves the best accuracy among all open-source models, with an average improvement of 4.8% over the base model, while simultaneously reducing tool-call rounds by 7%. This validates our hybrid evolution strategy’s ability to




Figure 3: Comparison of tool-call distribution and cumulative accuracy.

Method	xbench-deepresearch				Browsecomp				GAIA				HLE			
	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Token $\downarrow$	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Token $\downarrow$	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Token $\downarrow$	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Token $\downarrow$
Tongyi-DeepResearch	0.713	0.779	14.26	6918	0.410	0.385	63.70	12014	0.682	0.733	20.56	7378	0.358	0.487	23.92	13664
Unpruned-Distill	<b>0.746</b>	0.785	17.13	6317	<b>0.467</b>	0.408	63.80	11703	0.683	0.722	23.51	6992	<b>0.363</b>	0.492	23.70	14099
WebClipper (Eff)	0.713	0.792	10.81	5931	0.427	<b>0.431</b>	56.50	10599	0.684	<b>0.760</b>	14.44	4756	0.353	0.492	18.60	11458
WebClipper (Hybrid)	0.733	<b>0.797</b>	12.57	6205	<b>0.467</b>	0.428	60.42	11507	<b>0.695</b>	0.744	19.92	6635	0.361	<b>0.495</b>	21.07	13532

Table 2: Performance comparison of different training strategies.

balance efficiency and accuracy optimization. **3)** Further analysis in Figure 3 (a) shows that WebClipper(Eff)’s tool-call distribution is concentrated in lower-round buckets compared to the baseline, and Figure 3 (b) indicates WebClipper(Eff)’s accuracy curve converges much earlier, indicating superior performance in resource-constrained (low-round) scenarios. These results confirm that WebClipper effectively evolves agents to be more efficient without sacrificing, and sometimes even improving, their information-seeking capabilities.

**Comparison with Pruning Baselines (RQ2).** Results in Table 1 demonstrate WebClipper’s superiority over naive pruning strategies: **1)** Prompt-based pruning is insufficient. Compared to WebClipper(Eff), Prompt Control achieves only a marginal reduction in tool calls while suffering noticeable accuracy degradation. This suggests that directly prompting pre-trained web agents for efficiency is ineffective. **2)** Coarse-grained pruning causes severe performance drops. The Coarse Prune baseline, which relies on a single LLM to construct training samples through directly identifying redundant rounds, leads to a substantial accuracy drop. This indicates that trajectory optimization requires fine-grained, structured analysis rather than coarse judgment. In contrast, WebClipper’s structured, graph-based distillation process

allows for precise and reliable identification of redundancies, making it a far more effective pruning strategy.

**Validity of F-AE Score (RQ3).** The F-AE Score proves to be a balanced metric that avoids bias toward either dimension. As shown in Table 1: **1)** Despite using shorter rounds, Kimi-K2-Instruct-0905 scores low on F-AE due to their inferior accuracy, preventing the metric from rewarding efficiency alone. **2)** Although the accuracy of Tongyi-DeepResearch is close to WebClipper(Eff), its longer tool-call rounds result in lower F-AE scores, demonstrating the metric’s sensitivity to efficiency. **3)** WebClipper(Eff) achieves leading F-AE scores by maintaining high accuracy without excessive tool usage, reflecting its superior efficiency-accuracy balance. These patterns show that F-AE does not over-favor either accuracy or efficiency alone; instead, it rewards models that achieve a balanced performance. This supports F-AE as a reasonable and practically useful metric for evaluating web agents. Further explanation can be found in Appendix A.

### 4.3 Ablation Study

We now investigate **RQ4:** Are the key components of WebClipper effective? We conduct ablations on three aspects: the graph-based pruning method, the




Figure 4: Ablation Study of the key components of WebClipper.

coherence-aware rewriting strategy, and the agent evolution strategy.

**Ablation on Pruning Method & Rewriting Strategy.** We evaluate three variants: (1) *w/o GP*, replacing graph-based pruning with Coarse Prune but retaining the rewriting strategy; (2) *w/o PPL-S*, removing PPL-based selection and using the first generated rewriting as the final thought in trajectories; (3) *w/o CSR*, replacing context-aware selective rewriting with unconditional rewriting of all thoughts without providing the historical context. As shown in Figure 4, removing any component causes performance degradation. The decline of *w/o GP* can be attributed to the fact that single-pass LLM comprehension struggles with long trajectories. The drop in *w/o PPL-S* validates PPL-based filtering in maintaining alignment with the base model’s reasoning style. Most critically, *w/o CSR* leads to catastrophic collapse, confirming that naive rewriting without understanding context breaks reasoning coherence.

**Ablation on Evolution Strategy.** We compare three training strategies: WebClipper(Eff), WebClipper(Hybrid), and “Unpruned-Distill”. “Unpruned-Distill” follows the commonly adopted self-evolve paradigm (Aksitov et al., 2023), where original unpruned trajectories with  $0 < PR(q) \leq 0.5$  are directly used for SFT, and data is directly obtained from Section 3.2. As shown in Table 2, Unpruned-Distill improves accuracy over the base model but increases tool-call rounds, amplifying both strengths and inefficiencies. WebClipper(Eff)

achieves the lowest resource usage while maintaining accuracy comparable to the base model, making it preferable when efficiency is the primary concern. WebClipper(Hybrid) provides a more balanced option: relative to both Unpruned-Distill and the base model, it uses fewer rounds, attains accuracy clearly above the base model and close to Unpruned-Distill, and achieves the better F-AE scores. In practice, WebClipper(Eff) suits cost-sensitive deployments, whereas WebClipper(Hybrid) delivers a more comprehensive improvement in both efficiency and accuracy.

#### 4.4 Analysis and Discussion

Beyond efficiency gains, WebClipper also improves accuracy. We attribute this to the reasoning patterns induced by our pruned data, which trains the agent to focus on critical-path information. Existing web agents often fall into failure modes where they become stuck in unproductive branches, drift from the core objective, or enter cyclic reasoning loops. As shown in our case studies in Appendix C.2, over-focusing on trivial details can make the agent lose sight of the main goal. This not only reduces efficiency but also harms accuracy by inflating context length: an overly long context increases the risk that useful clues are drowned out by a mass of irrelevant, more recent tool interactions. Our pruning method counteracts this by constructing training samples in which irrelevant or repetitive tool-calling rounds are removed.

We also find that WebClipper’s efficiency gains are particularly notable on the GAIA dataset, where tool-call rounds are reduced by about 30%. We attribute this to the dataset’s characteristics: around 15% of its questions are brain teasers or logical puzzles that rely on the model’s intrinsic abstract reasoning and instruction-following, rather than long-horizon tool use. Excessive emphasis on external tools during training can not improve performance on such problems. Our method prevents the model from over-relying on external tools in these cases, substantially reducing unnecessary tool calls.

## 5 Conclusion

In this paper, we propose WebClipper, an innovative trajectory pruning method for web agents. We model web agent trajectories as state graphs and perform pruning on them. We further introduce two agent evolution strategies, which significantly reduce the number of tool calls while maintaining

or even improving the agent’s accuracy. In addition, we propose the F-AE score to better evaluate the overall capability of web agents in terms of both accuracy and efficiency. Extensive experiments demonstrate that WebClipper is an effective approach for balancing accuracy and efficiency.

## Limitations

WebClipper has achieved significant improvements in the efficiency of web agents, but there remain several limitations that point to future directions.

First, WebClipper inherits the planning and reasoning capabilities of the base model it distills from—if the base model’s performance is poor, the pruning process can only remove redundancy within those suboptimal trajectories rather than fundamentally improving the search strategy. Future work could explore integrating WebClipper with reinforcement learning or online learning mechanisms to enable the agent to discover novel, more efficient search patterns beyond those present in the base model’s behavior. Second, our pruning method is trained and evaluated on trajectories from specific web agent benchmarks that primarily involve search, web browsing, and code execution, leaving the generalization to emerging tool types (e.g., multimodal tools, database queries, or API integrations) unexplored. Extending WebClipper’s graph-based framework to accommodate diverse action spaces and information modalities represents a valuable direction for building more versatile and efficient agents across broader application domains.

## References

- Pranjal Aggarwal and Sean Welleck. 2025. [L1: Controlling how long a reasoning model thinks with reinforcement learning](#). *Preprint*, arXiv:2503.04697.
- Renat Aksitov, Sobhan Miryoosefi, Zonglin Li, Daliang Li, Sheila Babayan, Kavya Kopparapu, Zachary Fisher, Ruiqi Guo, Sushant Prakash, Pranesh Srinivasan, and 1 others. 2023. Rest meets react: Self-improvement for multi-step reasoning llm agent. *arXiv preprint arXiv:2312.10003*.
- anthropic. 2025. [Introducing claude 4](#).
- Song Bai, Lidong Bing, Carson Chen, Guanzheng Chen, Yuntao Chen, Zhe Chen, Ziyi Chen, Jifeng Dai, Xuan Dong, Wenhan Dou, Yue Deng, Yunjie Fu, Junqi Ge, Chenxia Han, Tammy Huang, Zhenhang Huang, Jerry Jiao, Shilei Jiang, Tianyu Jiao, and 35 others. 2025a. [Mirothinker: Pushing the performance boundaries of open-source research agents via model, context, and interactive scaling](#). *Preprint*, arXiv:2511.11793.
- Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, Yichen Feng, and 149 others. 2025b. [Kimi k2: Open agentic intelligence](#). *Preprint*, arXiv:2507.20534.
- 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.
- Claude Team. 2025. [Claude research](#).
- Yingqian Cui, Pengfei He, Jingying Zeng, Hui Liu, Xianfeng Tang, Zhenwei Dai, Yan Han, Chen Luo, Jing Huang, Zhen Li, Suhang Wang, Yue Xing, Jiliang Tang, and Qi He. 2025. [Stepwise perplexity-guided refinement for efficient chain-of-thought reasoning in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 18581–18597, Vienna, Austria. Association for Computational Linguistics.
- Razvan-Gabriel Dumitru, Darius Peteleaza, Vikas Yadav, and Liangming Pan. 2025. [ConciseRL: Conciseness-guided reinforcement learning for efficient reasoning models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 17099–17123, Suzhou, China. Association for Computational Linguistics.
- Gemini Team. 2025. [Gemini deep research](#).
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. 2025. [Token-budget-aware LLM reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24842–24855, Vienna, Austria. Association for Computational Linguistics.
- David J Hand, Peter Christen, and Nishadi Kirielle. 2021. F\*: an interpretable transformation of the f-measure. *Machine learning*, 110(3):451–456.
- Chen Hu, Haikuo Du, Heng Wang, Lin Lin, Mingrui Chen, Peng Liu, Ruihang Miao, Tianchi Yue, Wang You, Wei Ji, Wei Yuan, Wenjin Deng, Xiaojian Yuan, Xiaoyun Zhang, Xiangyu Liu, Xikai Liu, Yanming Xu, Yicheng Cao, Yifei Zhang, and 48 others. 2025. [Step-deepresearch technical report](#). *Preprint*, arXiv:2512.20491.

- Jina.ai. 2025. [Jina](#).
- Baixuan Li, Bo Zhang, Dingchu Zhang, Fei Huang, Guangyu Li, Guoxin Chen, Huifeng Yin, Jialong Wu, Jingren Zhou, Kuan Li, Liangcai Su, Litu Ou, Liwen Zhang, Pengjun Xie, Rui Ye, Wenbiao Yin, Xinmiao Yu, Xinyu Wang, Xixi Wu, and 37 others. 2025a. [Tongyi deepresearch technical report](#). *Preprint*, arXiv:2510.24701.
- Kuan Li, Zhongwang Zhang, Huifeng Yin, Liwen Zhang, Litu Ou, Jialong Wu, Wenbiao Yin, Baixuan Li, Zhengwei Tao, Xinyu Wang, Weizhou Shen, Junkai Zhang, Dingchu Zhang, Xixi Wu, Yong Jiang, Ming Yan, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025b. [Websailor: Navigating super-human reasoning for web agent](#). *Preprint*, arXiv:2507.02592.
- Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yongkang Wu, Ji-Rong Wen, Yutao Zhu, and Zhicheng Dou. 2025c. [Webthinker: Empowering large reasoning models with deep research capability](#). *Preprint*, arXiv:2504.21776.
- Zijian Li, Xin Guan, Bo Zhang, Shen Huang, Houquan Zhou, Shaopeng Lai, Ming Yan, Yong Jiang, Pengjun Xie, Fei Huang, Jun Zhang, and Jingren Zhou. 2025d. [Webweaver: Structuring web-scale evidence with dynamic outlines for open-ended deep research](#). *Preprint*, arXiv:2509.13312.
- Junteng Liu, Yunji Li, Chi Zhang, Jingyang Li, Aili Chen, Ke Ji, Weiyu Cheng, Zijia Wu, Chengyu Du, Qidi Xu, Jiayuan Song, Zhengmao Zhu, Wenhui Chen, Pengyu Zhao, and Junxian He. 2025. [Webexplorer: Explore and evolve for training long-horizon web agents](#). *Preprint*, arXiv:2509.06501.
- Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. 2025. [O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning](#). *Preprint*, arXiv:2501.12570.
- Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. 2025. [CoT-valve: Length-compressible chain-of-thought tuning](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6025–6035, Vienna, Austria. Association for Computational Linguistics.
- Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2023. [Gaia: a benchmark for general ai assistants](#). In *The Twelfth International Conference on Learning Representations*.
- Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim, and Se-Young Yun. 2025. [Self-training elicits concise reasoning in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 25127–25152, Vienna, Austria. Association for Computational Linguistics.
- OpenAI. 2024. [Learning to reason with LLMs](#).
- OpenAI. 2025a. [Deep research system card](#).
- OpenAI. 2025b. [Introducing openai o3 and o4-mini](#).
- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, and 1 others. 2025. [Humanity’s last exam](#). *arXiv preprint arXiv:2501.14249*.
- Soham Poddar, Paramita Koley, Janardan Misra, Niloy Ganguly, and Saptarshi Ghosh. 2025. [Brevity is the soul of sustainability: Characterizing LLM response lengths](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 21848–21864, Vienna, Austria. Association for Computational Linguistics.
- GPT Research. 2025a. [Gpt research](#).
- Open Deep Research. 2025b. [Open deep research](#).
- SerpAPI. 2025. [Serpapi: Google search api](#).
- Dingfeng Shi, Jingyi Cao, Qianben Chen, Weichen Sun, Weizhen Li, Hongxuan Lu, Fangchen Dong, Tianrui Qin, King Zhu, Minghao Liu, Jian Yang, Ge Zhang, Jiaheng Liu, Changwang Zhang, Jun Wang, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2025. [Taskcraft: Automated generation of agentic tasks](#). *Preprint*, arXiv:2506.10055.
- Zhengwei Tao, Haiyang Shen, Baixuan Li, Wenbiao Yin, Jialong Wu, Kuan Li, Zhongwang Zhang, Huifeng Yin, Rui Ye, Liwen Zhang, Xinyu Wang, Pengjun Xie, Jingren Zhou, and Yong Jiang. 2025a. [Webbleaper: Empowering efficiency and efficacy in webagent via enabling info-rich seeking](#). *Preprint*, arXiv:2510.24697.
- Zhengwei Tao, Jialong Wu, Wenbiao Yin, Junkai Zhang, Baixuan Li, Haiyang Shen, Kuan Li, Liwen Zhang, Xinyu Wang, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025b. [Webshaper: Agentic data synthesizing via information-seeking formalization](#). *Preprint*, arXiv:2507.15061.
- Junjie Wang, Mingyang Chen, Binbin Hu, Dan Yang, Ziqi Liu, Yue Shen, Peng Wei, Zhiqiang Zhang, Jinjie Gu, Jun Zhou, Jeff Z. Pan, Wen Zhang, and Huajun Chen. 2024. [Learning to plan for retrieval-augmented large language models from knowledge graphs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7813–7835, Miami, Florida, USA. Association for Computational Linguistics.
- Jason Wei, Zhiqing Sun, Spencer Papay, Scott McKinney, Jeffrey Han, Isa Fulford, Hyung Won Chung, Alex Tachard Passos, William Fedus, and Amelia Glaese. 2025. [Browsecomp: A simple yet challenging benchmark for browsing agents](#). *arXiv preprint arXiv:2504.12516*.

- Jialong Wu, Baixuan Li, Runnan Fang, Wenbiao Yin, Liwen Zhang, Zhengwei Tao, Dingchu Zhang, Zekun Xi, Gang Fu, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025. [Webdancer: Towards autonomous information seeking agency](#). *Preprint*, arXiv:2505.22648.
- Xbench Team. 2025. [Xbench-deepsearch](#).
- Siheng Xiong, Oguzhan Gungordu, Blair Johnson, James C Kerce, and Faramarz Fekri. 2026. Scaling search-augmented llm reasoning via adaptive information control. *arXiv preprint arXiv:2602.01672*.
- Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. 2025. [Chain of draft: Thinking faster by writing less](#). *Preprint*, arXiv:2502.18600.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Howard Yen, Ashwin Paranjape, Mengzhou Xia, Thejas Venkatesh, Jack Hessel, Danqi Chen, and Yuhao Zhang. 2025. Lost in the maze: Overcoming context limitations in long-horizon agentic search. *arXiv preprint arXiv:2510.18939*.
- Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, Kedong Wang, Lucen Zhong, Mingdao Liu, Rui Lu, Shulin Cao, Xiaohan Zhang, Xuancheng Huang, Yao Wei, Yean Cheng, and 151 others. 2025. [Glm-4.5: Agentic, reasoning, and coding \(arc\) foundation models](#). *Preprint*, arXiv:2508.06471.

## A Design of F-AE Score

In this section, we provide a more detailed explanation of the F-AE Score, drawing an analogy to the classic F1-Score in information retrieval, and clarifying the design choices behind its formulation.

### A.1 Background: The F1 Score

The F1 score is a widely-used metric in classification tasks that harmonizes precision and recall through their harmonic mean:

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

The key insight of F1 is that it balances two competing objectives—precision (quality of positive predictions) and recall (coverage of actual positives)—in a way that penalizes extreme imbalance. Unlike the arithmetic mean, which would be  $\frac{\text{Precision} + \text{Recall}}{2}$ , the harmonic mean is more sensitive to low values. For instance, if Precision = 1.0 but Recall = 0.1, the arithmetic mean yields 0.55, while F1 yields only 0.18, reflecting that a model excelling in only one dimension is suboptimal.

### A.2 Motivation for F-AE Score

When evaluating Web Agents, we face an analogous trade-off between two competing objectives:

- Accuracy ( $Acc$ ): how often the agent produces a correct answer.
- Efficiency ( $E$ ): how economical the agent is in its use of tool-calling rounds.

Existing evaluation paradigms often optimize these metrics in isolation. To holistically assess agent quality, we need a metric that captures their joint optimization.

### A.3 Design of F-AE Score

Our F-AE Score follows exactly the same philosophy as F1 Score, but replaces precision/recall with  $Acc$  and ( $E$ ):

We first normalize the number of tool-calling rounds to an “efficiency score” in  $[0, 1]$ :

$$E = 1 - \frac{\text{Rounds}}{\text{Max\_Rounds}},$$

where  $\text{Rounds}$  is the average number of tool-call turns used by the agent, and  $\text{Max\_Rounds}$  is the maximum allowable rounds in the deployment scenario (set to 100 in our experiments). Intuitively,

if an agent uses  $\text{Rounds} = \text{Max\_Rounds}$ , then  $E = 0$ , i.e., “maximally inefficient”; if an agent uses very few rounds, say  $\text{Rounds} \approx 0$ , then  $E \approx 1$ , i.e., “highly efficient”.

We then define F-AE Score as the harmonic mean of accuracy and efficiency:

$$F\text{-AE} = 2 \times \frac{Acc \times E}{Acc + E} = 2 \times \frac{Acc \times \left(1 - \frac{\text{Rounds}}{\text{Max\_Rounds}}\right)}{Acc + \left(1 - \frac{\text{Rounds}}{\text{Max\_Rounds}}\right)}$$

where both  $Acc$  and  $E$  are normalized to  $[0, 1]$ , ensuring  $F\text{-AE} \in [0, 1]$  for interpretability. A higher F-AE means better overall performance, taking both dimensions into account. This makes it easy to compare different Web Agents or training strategies.

Using the harmonic mean between  $Acc$  and  $E$  has several desirable properties:

#### 1. Balance between accuracy and efficiency.

If either accuracy or efficiency is low, F-AE will be low. For example:

- A model with high accuracy but extremely long trajectories ( $E \approx 0$ ) will receive a low F-AE.
- A model with very short trajectories but poor accuracy ( $Acc \approx 0$ ) will also receive a low F-AE. This matches our intuitive requirement that a “good” Web Agent must be both effective and efficient.

#### 2. No arbitrary dominance of one dimension.

Unlike a simple weighted sum (e.g.,  $\alpha \cdot Acc + (1 - \alpha) \cdot E$ ), the harmonic mean is far less tolerant of one dimension being much smaller than the other. This prevents scenarios where:

- Slight gains in accuracy justify arbitrarily large increases in rounds
- Slight savings in rounds justify large accuracy drops.

In other words, F-AE inherently discourages extreme trade-offs.

### A.4 Effect of Max\_Rounds and Scaling

The parameter  $\text{Max\_Rounds}$  controls how aggressively we penalize tool usage:

$$E = 1 - \frac{\text{Rounds}}{\text{Max\_Rounds}}.$$

When  $\text{Rounds} \ll \text{Max\_Rounds}$ ,  $E$  is close to 1, so efficiency is considered good and F-AE is

mainly determined by accuracy. When Rounds approaches Max\_Rounds,  $E$  decreases toward 0, pulling F-AE down even if accuracy remains high.

In our experiments, Max\_Rounds = 100 is chosen to reflect the typical upper bound used in Deep Research-style Web Agents. In principle, Max\_Rounds can be adjusted to match different deployment constraints (e.g., stricter limits in latency-critical settings).

An important point is that F-AE is relative to the chosen budget: if all methods are evaluated with the same Max\_Rounds, F-AE provides a fair way to compare them under that shared resource regime.

## B Implementation Details

This appendix elaborates on the implementation of our trajectory pruning and rewriting pipeline, providing conceptual descriptions and the specific prompts used.

### B.1 Details of Rejection Sampling

We use the public QA datasets to distill trajectories. The used dataset includes WebDancer (200 samples), WebShaper (500 samples), WebExplorer (100 samples), Voyager (a subset consisting of 5k samples), and TaskCraft (a subset consisting of 4k samples). In the Tongyi-DeepResearch environment, we ran all samples four times, keeping those with a pass rate  $0 < \text{PR}(q) \leq 0.5$ . This data was then used for subsequent pruning.

### B.2 State Graph Construction

The construction of the state graph  $\mathcal{G}$  from a raw trajectory  $\tau$  is a two-phase process orchestrated by an LLM extractor.

**Phase 1: Action Node Extraction** First, we process the trajectory to identify each assistant turn uniquely. Each turn, consisting of a thought-action pair  $(t_k, a_k)$ , is mapped to a corresponding Action Node  $A_k$ . We employ an LLM extractor (example prompt is shown in Figure 5) that receives the conversational history up to step  $k - 1$  and the current turn’s content. The extractor’s task is to summarize this turn into a compact JSON object with two fields: an “Action” type (e.g., Search, PythonInterpreter, Answer) and a “Goal” description. This process is parallelized across all turns in the trajectory for efficiency, yielding the complete set of action vertices,  $\mathcal{V}^A$ .

**Phase 2: Iterative Information and Edge Construction** With the action nodes  $\mathcal{V}^A$  established, we iteratively build the information nodes  $\mathcal{V}^I$  and the dependency edges  $\mathcal{E}$ . The process is initialized with a graph containing only the initial query node  $I_0$  and the first action node  $A_1$ , connected by an edge  $(I_0, A_1)$ .

We then iterate from  $k = 1$  to  $T - 1$ . In each iteration, the LLM extractor is prompted (example prompt is shown in Figure 6) with the current graph state and a snippet of the trajectory:  $(A_k, o_k, A_{k+1})$ , where  $o_k$  is the observation received after action  $A_k$ . The LLM performs two functions:

1. **Decomposing Observations:** It analyzes  $o_k$  to extract atomic units of information. For each unit, it checks for semantic equivalence with existing nodes in  $\mathcal{V}^I$ . If a match is found, an edge  $A_k \rightarrow I_{\text{existing}}$  is added. Otherwise, a new information node  $I_{\text{new}}$  is created and added to  $\mathcal{V}^I$ , along with an edge  $A_k \rightarrow I_{\text{new}}$ .
2. **Linking Actions:** It analyzes  $A_{k+1}$  to identify which information nodes in the current graph (including any newly created ones) served as its basis. For each identified supporting node  $I'$ , an edge  $I' \rightarrow A_{k+1}$  is added.

This iterative process continues until all actions and observations have been incorporated, resulting in the final state graph  $\mathcal{G}$ .

### B.3 Pruning via MNDAG and Majority Vote

**MNDAG Identification** Given the state graph  $\mathcal{G}$ , we identify the minimal necessary subgraph using a two-stage algorithm, detailed in Algorithm 1. This algorithm approximates the Minimal-cost Necessary Directed Acyclic Graph (MNDAG).

**Robustness via Majority Vote** A single LLM-driven graph construction can be prone to inconsistencies. To enhance robustness, we repeat the entire process—from graph construction to MNDAG mining—three times for the same trajectory. This yields three candidate sets of necessary actions:  $\mathcal{A}^{*(1)}$ ,  $\mathcal{A}^{*(2)}$ ,  $\mathcal{A}^{*(3)}$ . A final set  $\mathcal{A}_{\text{final}}^*$  is accepted only if at least two of the three candidate sets are identical. If no majority is reached, the pruning for that trajectory is considered unreliable and is discarded, ensuring we only proceed with high-confidence results.

---

**Algorithm 1** MNDAG Identification Algorithm

---

```
1: Input: State Graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , source  $I_0$ , sink  $A_T$ 
2: Output: Set of necessary action nodes  $\mathcal{A}^*$ 
   // Step 1: Forward search for shortest path costs
3: Define cost function:  $c(v) = 1$  if  $v \in \mathcal{V}^A$ , else  $c(v) = 0$ .
4: Run Dijkstra’s algorithm starting from  $I_0$  to compute the shortest distance  $d(v)$  and predecessor  $p(v)$  for every node  $v \in \mathcal{V}$ .
   // Step 2: Backward traversal to identify necessary nodes
5: Initialize necessary node set  $\mathcal{V}^* \leftarrow \{A_T\}$ .
6: Initialize a queue for traversal  $Q \leftarrow [A_T]$ .
7: while  $Q$  is not empty do
8:   Dequeue a node  $v$ .
9:   if  $v \in \mathcal{V}^A$  then  $\triangleright$  If node is an Action
10:    for each predecessor  $u$  of  $v$  in  $\mathcal{G}$  do
11:      if  $u \notin \mathcal{V}^*$  then
12:        Enqueue  $u$  and add to  $\mathcal{V}^*$ .
13:      end if
14:    end for
15:   else if  $v \in \mathcal{V}^I$  and  $p(v)$  exists then  $\triangleright$  If node is Information
16:     Let  $u \leftarrow p(v)$   $\triangleright$  Get predecessor from Dijkstra’s path tree
17:     if  $u \notin \mathcal{V}^*$  then
18:       Enqueue  $u$  and add to  $\mathcal{V}^*$ .
19:     end if
20:   end if
21: end while
   // Step 3: Extract final action set
22:  $\mathcal{A}^* \leftarrow \{A \mid A \in \mathcal{V}^* \cap \mathcal{V}^A\}$ 
23: return  $\mathcal{A}^*$ 
```

---

#### B.4 Coherence-aware Thought Rewriting

Simply deleting steps can create logical gaps. When intermediate rounds are pruned, the remaining reasoning steps may still refer to observations that have already been removed, thereby creating incoherent training signals. Consider the following example.

**Before pruning (Rounds 1  $\rightarrow$  2  $\rightarrow$  3  $\rightarrow$  4).**

- **Round 1:** Search for “Nobel Prize in Physics 2020.” *Observation:* The search engine returns relevant results.
- **Round 2:** Visit a tangentially related page about the 2019 Nobel Prize. *Observation:*

The page contains information about the 2019 laureates, which turns out to be an unproductive branch.

- **Round 3 Thought:** “*The previous page was about the 2019 prize, which is not what I need. Let me go back and focus on the 2020 laureates instead.*”
- **Round 4:** Visit the correct page for the 2020 Nobel Prize. *Observation:* The page returns the correct answer.

**After pruning Round 2 (Rounds 1  $\rightarrow$  3  $\rightarrow$  4).** Round 3 still contains the thought:

“*The previous page was about the 2019 prize, which is not what I need.*”

However, the visit to the 2019 page in Round 2 has already been removed. If this pruned trajectory is used directly for training, the model may learn to hallucinate a failed attempt that never occurred in the retained context. As a result, the model could be encouraged to fabricate non-existent observations.

**After rewriting.** The thought in Round 3 can be revised as follows:

“*I found relevant search results. Let me visit the page for the 2020 Nobel Prize to get the details.*”

This rewriting step ensures that the pruned trajectory remains logically self-consistent and free from hallucinated references, making it safer and more effective as training data.

**Context-aware Selective Rewriting** We only rewrite thoughts that become disconnected from their new predecessors after pruning. A thought  $t_{k+1}^{\text{new}}$  is rewritten if the action  $a_k^{\text{new}}$  preceding it in the pruned trajectory was not its direct predecessor in the original trajectory. The rewriting LLM (example prompt is shown in Figure 7) is conditioned on a comprehensive context:

- **Dialogue History:** The sequence of necessary messages generated so far.
- **Skipped Messages:** The raw content of all intermediate steps that were pruned between the new adjacent steps. This provides the LLM with the knowledge of what occurred in the gap.

Method	xbench-deepsearch			GAIA			Browsecomp			HLE		
	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$	Acc $\uparrow$	F-AE $\uparrow$	Rounds $\downarrow$
WebExplorer	0.517	0.659	9.05	0.372	0.512	12.88	0.137	0.229	29.43	0.116	0.203	15.52
+WebClipper(Eff)	0.560	0.696	7.82	0.407	0.564	8.22	0.135	0.230	22.95	0.146	0.252	7.4

Table 3: Performance of WebClipper based on WebExplorer-8B.

- **Current Action to Refine:** The original thought from the step being rewritten.

This rich context enables the LLM to generate a new thought that smoothly bridges the logical gap while avoiding hallucinations by not referencing pruned observations directly.

**Perplexity-based Selection** To ensure the rewritten thought aligns with the base model’s intrinsic reasoning style, we generate three candidate rewrites for each required modification. We then use the base model itself to calculate the perplexity (PPL) of each candidate. The PPL is computed over the rewritten thought, conditioned on the preceding dialogue history. The candidate with the lowest PPL is selected, which ensures maximal fluency and stylistic consistency with the model’s own text distribution.

## C Additional Experiment Results

### C.1 Training based on WebExplorer-8B

To further demonstrate the generalizability of our method, we use the trajectories generated by our WebClipper framework to fine-tune a different and existing web-agent, WebExplorer-8B.

As shown in Table 3, our method can perform well on WebExplorer-8B, maintaining model performance while reducing tool-call rounds. This demonstrates that our method is not model-specific and possesses generalization ability.

### C.2 Case Study

#### C.2.1 Case 1

As shown in Table 4, the task requires identifying the nano-compound studied in a specific 2012 Scientific Reports article that does not mention “plasmons” or “plasmonics”. The initial agent correctly identifies the target article, “Diamond photonic crystal slab: Leaky modes and modified photoluminescence emission of surface-deposited quantum dots”. However, it then engages in “excessively divergent exploration”. It shifts its focus from the primary subject of the paper (the diamond slab)

to a trivial detail—the “surface-deposited quantum dots” mentioned in the title. This leads to a long chain of tool calls (10+ rounds) to identify the quantum dots’ material (“silicon nanocrystals”) and repeatedly validate the initial conditions. This over-exploration, which significantly inflates the context length, causes the model to **lose sight of the core objective**, mistaking the experimental probe for the main subject of study. In contrast, WebClipper demonstrates a more pruned and effective reasoning path. After identifying the correct article in just two rounds, it directly infers the main subject, “diamond”, from the title and a concise tool-provided summary. By avoiding the irrelevant deep-dive into the quantum dot material, it prevents context dilution and the risk of forgetting critical initial information.

#### C.2.2 Case 2

We further demonstrate WebClipper’s efficiency gains with a second case (Table 5). The task is to calculate the time for Eliud Kipchoge to run to the Moon’s perigee. This requires finding two constants (Kipchoge’s pace and the Moon’s minimum perigee) and performing a calculation. The initial model exhibits clear hallmarks of an unpruned, exhaustive search. It gets bogged down by the slight ambiguity of the term “record-making marathon pace,” which could refer to several of Kipchoge’s historic runs. This uncertainty triggers multiple, redundant search and visit cycles to verify and re-verify his latest record time, as well as the Moon’s perigee distance. Furthermore, it engages in superfluous exploration by calculating the final answer for three different marathon times, even though the question implies a single, definitive pace. This inefficient, cyclical trajectory significantly increases the number of tool calls (over 15 rounds). In contrast, WebClipper adopts a more decisive and linear strategy. It makes a reasonable initial assumption for the record time and proceeds along a direct path: find the two required constants, then compute the result. This pruned approach, consisting of just four rounds, entirely avoids the redundant validation loops and superfluous computations seen

in the baseline. This case demonstrates that our training method teaches the model to commit to a reasonable and efficient path, improving performance by eliminating unnecessary and costly over-exploration.

## User Prompt

## Objective: Extract an Action Node from a Single Assistant Turn

Task Overview: You will receive a separate conversation history and a "current Assistant turn" to be analyzed. Your task is to focus solely on analyzing the "current Assistant turn", summarize it into a single Action Node, and output it in the format of a single JSON object.

### 1. Conversation Content

#### 1.1 Conversation History

<history>

[CHAT\_HISTORY]

</history>

#### 1.2 Current Assistant Turn to be Analyzed

<current\_turn>

[CURRENT\_TURN]

</current\_turn>

### 2. JSON Output Format Requirements

The final output must strictly follow the structure of the single JSON object below. Do not include any lists or outer wrappers for the `Action Node`.

For example:

```
{"Action": "Search", "Goal": "Get the statistical data for China's urban population in 1949 and 2009"}
```

```
{"Action": "PythonInterpreter", "Goal": "Calculate the Compound Annual Growth Rate (CAGR) based on the population data"}
```

```
{"Action": "Visit", "Goal": "Verify the official source for the 1949 urban population data"}
```

```
{"Action": "Answer", "Goal": "Provide the answer: approximately 4.04%"}
```

### 3. Action Node (A-Node) Extraction Rules

1. Action (Core Rule):

\* If the `content` field of the "Current Assistant Turn to be Analyzed" contains the special tag `<answer>`, the `Action` must be assigned the value `Answer`.

\* Otherwise, assign the `Action` as one of `Search`, `Visit`, or `PythonInterpreter` based on the tool it calls.

2. Goal:

\* For non-`Answer` actions: Summarize the purpose of the tool call in a concise description. This summary should cover the search queries for a Search, the goal for a Visit, or the calculation's purpose for a PythonInterpreter.

\* For an `Answer` action: The `Goal` field must be "Provide the answer: " + the extracted core short answer (a word or a short phrase), which is the minimal piece of information that directly answers the original user's question.

Now, based on the requirements above, process the provided "Current Assistant Turn to be Analyzed":

Output:

Figure 5: The Prompt of Action Node Extraction

### User Prompt

## Task: Incremental Graph Construction

You are a graph construction assistant. Your task is to update a state graph based on the provided context.

### Task Definition

Strictly follow the requirements below to generate a JSON object containing `new\_information\_nodes` and `new\_edges`.

#### Step A. Determine whether to construct new information nodes (new\_information\_nodes)

[Instruction Omitted]

#### Step B. Construct new edges (new\_edges)

[Instruction Omitted]

#### C. Output Format

Strictly output in the following JSON format. Do not include any other explanations.

### Task Execution Examples

.....

### Actual Current Task

Now, generate output for the following actual task.

#### Context

```
<previous_action>  
{previous_action_json}  
</previous_action>
```

```
<observation>  
{current_observation_content}  
</observation>
```

```
<current_action>  
{current_action_json}  
</current_action>
```

#### Graph Workspace

```
<graph_workspace>  
{workspace_json}  
</graph_workspace>
```

Output:

Figure 6: The Prompt of Iterative Information and Edge Construction

### User Prompt

You are an expert AI assistant content editor. Your task is to refine and condense the thinking process provided to you, based on the dialogue history.

I will give you three parts:

1. Dialogue History: The necessary messages that have been stored so far
2. Skipped Messages: Messages between the previous necessary message and current necessary message that are judged to provide no information gain
3. Current Action to Refine: The current assistant's thought process and actions that need to be refined

Your task is to modify the current assistant's thought process so that it connects smoothly with the previous necessary message (in Dialogue History), while skipping the intermediate messages that provide no information gain.

Follow these rules strictly:

[Instruction Omitted]

Your output should ONLY be the refined content, with no extra explanations from you.

## Dialogue History (Already Stored Necessary Messages):

[PREVIOUS\_CONTEXT]

## Skipped Messages (Between Previous Necessary Message and Current One):

[SKIPPED\_MESSAGES\_STR]

## Current Action to Refine (Thinking Process Only):

[CONTENT\_TO\_REFINE]

## Important Instructions:

1. The skipped messages have been identified as providing no information gain. Merge the thinking process seamlessly.
2. Maintain the original first-person thinking style with self-questioning and logical reasoning.
3. Ensure perfect coherence between the previous necessary message's tool calls and the current one.
4. If skipped messages include failed tool attempts, incorporate that reasoning into the refined think.
5. Do not introduce any content not present in Dialogue History.

You are only allowed to refine the <think> section.

Your output should be the refined <think> section only, without the <tool\_call>.

## Output Format:

<refined\_think>

YOUR REFINED THINK

</refined\_think>

## Your Response:

Figure 7: The Prompt of Message Refine

Table 4: Case 1 comparison between WebClipper and Tongyi-DeepResearch.

<b>Query:</b> In Nature journal’s Scientific Reports conference proceedings from 2012, in the article that did not mention plasmons” or plasmonics”, what nano-compound is studied? Don’t use the prefix nano in your answer if there is one.	
<b>WebClipper</b>	<b>Baseline Model</b>
<i>1. Assistant calls “search(query=[“Scientific Reports 2012 conference...”])”</i>	<i>1. Assistant calls “search(query=[“Scientific Reports 2012 conference...”])”</i>
<b>Tool Response:</b> Provides a link to the 2012 conference proceedings page on nature.com. [Conference Proceedings in 2012   Scientific Reports](url:...)	<b>Tool Response:</b> Provides a link to the 2012 conference proceedings page on nature.com. [Conference Proceedings in 2012   Scientific Reports](url:...)
<i>2. Assistant calls “visit(url=[“...”], goal=“List articles... and identify which one does not mention plasmons...”)”</i>	<i>2. Assistant calls “visit(url=[“...”], goal=“List conference proceeding articles...”)”</i>
<b>Tool Response:</b> The 2012 conference proceedings... contain four articles. Three articles explicitly mention plasmon-related terms... The article that does not mention plasmons... is “Diamond photonic crystal slab...”. The studied nano-compound in this article is diamond.	<b>Tool Response:</b> The webpage lists four conference proceeding articles: “Diamond photonic crystal slab...”, “2D quasiperiodic plasmonic crystals”, “Layered plasmonic cloaks...”, etc.
<i>3. Assistant concludes and provides the final answer. “&lt;answer&gt;diamond&lt;/answer&gt;”</i>	<i>3. Assistant calls “visit(url=[“.../srep00914”], goal=“...find the nano-compound studied...”)”</i>
	<b>Tool Response:</b> ...The research investigates the leaky modes of this diamond PhC slab and their effect on the photoluminescence (PL) emission of surface-deposited quantum dots. The specific nano-compound studied is silicon nanocrystals (SiNCs)...
	<i>4-10. A series of subsequent “visit” and “search” calls.</i> The agent repeatedly visits “srep00914” to confirm the absence of “plasmon,” visits other articles (e.g., “srep00915”) to confirm the presence of “plasmon,” and accesses the PDF version of “srep00914” for deeper analysis.
	<b>Tool Response:</b> Responses confirm that “srep00914” does not mention “plasmon”, while other articles do. This reinforces the agent’s focus on “srep00914” and its contents, particularly the SiNCs.
	<i>11. Assistant concludes and provides the final answer. “&lt;answer&gt;...the nano-compound studied in the only Scientific Reports conference proceeding article from 2012 that does not mention plasmons or plasmonics is silicon nanocrystals...&lt;/answer&gt;”</i>

Table 5: Case 2 comparison between WebClipper and Tongyi-DeepResearch.

<p><b>Query:</b> If Eliud Kipchoge could maintain his record-making marathon pace indefinitely, how many thousand hours would it take him to run the distance between the Earth and the Moon its closest approach? Please use the minimum perigee value on the Wikipedia page for the Moon when carrying out your calculation. Round your result to the nearest 1000 hours and do not use any comma separators if necessary.</p>	
WebClipper	Baseline Model
<p>1. Assistant calls “PythonInterpreter” to calculate speed based on a known record (2h 1m 39s).</p>	<p>1. Assistant calls “search(query=[“Eliud Kipchoge record making...”])”</p>
<p><b>Tool Response:</b> “stdout: 20.811...”</p>	<p><b>Tool Response:</b> Provides multiple results, including his 2:01:39 record, his sub-2-hour 1:59:40 attempt, and his more recent 2:01:09 record, creating ambiguity.</p>
<p>2. Assistant calls “search(query=[“Moon perigee minimum distance...”])”</p>	<p>2. Assistant calls “search(query=[“Moon perigee closest approach...”])”</p>
<p><b>Tool Response:</b> Provides multiple values for the Moon’s perigee.</p>	<p><b>Tool Response:</b> Provides multiple values for the Moon’s perigee.</p>
<p>3. Assistant calls “visit” on the Wikipedia page for the Moon to find the precise minimum perigee.</p>	<p>3. Assistant calls “visit” on the Wikipedia page to find the precise minimum perigee.</p>
<p><b>Tool Response:</b> ...the minimum perigee distance value listed is 356,400 kilometres.</p>	<p><b>Tool Response:</b> ...minimum perigee distance... is 356,400 km...</p>
<p>4. Assistant combines the constants and provides the final answer. “&lt;answer&gt;...it would take approximately 17 thousand hours.&lt;/answer&gt;”</p>	<p>4. Assistant calls “search” again to re-verify the latest world record.</p>
	<p><b>Tool Response:</b> Confirms the 2022 record is 2:01:09.</p>
	<p>5-8. A series of “PythonInterpreter” calls. The agent calculates the speed for the 2:01:09 record and then the total time.</p>
	<p>9-14. A series of subsequent “search”, “visit”, and “PythonInterpreter” calls. The agent engages in redundant validation (re-visiting Wikipedia) and superfluous exploration (calculating the final time for the other two marathon paces, 1:59:40 and 2:01:39).</p>
	<p><b>Tool Response:</b> The calculations for all three marathon paces round to the same final answer.</p>
	<p>15. Assistant compiles all findings and provides the final answer. “&lt;answer&gt;...Eliud Kipchoge would require approximately 17000 thousand hours...&lt;/answer&gt;”</p>