

长程深度研究智能体的 数据构建与上下文管理范式

上海交通大学

叶锐

2026 年 7 月

个人简介

叶锐, 上海交通大学博四

- 智能体模型
 - Agentic LLMs
 - Context Management
 - Agent Data
- 智能体系统
- Vibe Writing 写作产品
- 通用科研助手

<https://rui-ye.github.io/>



I am a fourth-year PhD candidate at [Shanghai Jiao Tong University \(SJTU\)](#), advised by [Prof. Siheng Chen](#). I was selected as the **student representative speaker** at **SJTU's 130th Anniversary**. To date, I have published **14 (co-)first-author papers** in top-tier conferences and journals. Prior to my PhD, I received my Bachelor degree from SJTU with a rank of **1st out of 150**.

My research generally centers around **AI Agents and Collaborative AI**, including:

- **Agentic LLMs (Reasoning, Tool, Context, etc.):**
 - [Tongyi DeepResearch](#), an open-source SOTA agentic LLM with 30B size
 - [AgentFold](#), a new agentic paradigm with proactive context management
 - [OpenSeeker](#), a frontier search agent with open-source data & model
 - [OpenSeeker-v2](#), a ReAct-based search agent (10k SFT, yet, SOTA)
 - [SciMaster](#), a general-purpose scientific agent (achieves >30% on HLE)
 - [WebSailor-V2](#), a complete pipeline from data to training for LLM agents
- **LLM-based Multi-Agent Systems (MAS, Workflow):**
 - [MAS-GPT](#), an LLM that can generate executable query-specific MAS
 - [MASLab](#), a comprehensive codebase for 20+ LLM-based MAS methods
- Collaborative Training (Federated Learning):
 - [OpenFedLLM](#), a research codebase for decentralized LLM training



I am also passionate about turning research into products that serve humans. For example, I have worked on:

- [Muse](#) (Lead, end-to-end project from UI/UX to engineering to algorithm). The 'Cursor' for general writing where LLM agents tutor, write, edit, insert citations, compile, fix... Built and shipped by a two-person (@Haochen Zhao) core team in 2 months.
- [SciMaster](#). An agentic platform for researchers that includes deep research, automated survey writing, etc.

I have gained internship experience at Microsoft Research Asia (MSRA), Shanghai AI Laboratory, and Alibaba Tongyi Lab. I am seeking opportunities as a **research intern** or a **full-time researcher** and am open to collaborations. Feel free to reach out!

目录

00 研究背景介绍

01 高质量数据合成方案

02 主动上下文管理范式

03 总结

研究背景：通往 AGI 的新阶段

随着技术爆发式发展，智能体（Agents）已跃升为全球学术界与工业界的核心焦点。OpenAI 将其定义为通往通用人工智能（AGI）的**第三阶段**。2025年被广泛视作智能体元年，而今年是智能体爆发的一年。



01 智能体范式跃迁

从早期的对话机器人（Chatbots），进化为具备逻辑推理的Reasoners，最终走向能自主决策的Agents。AI正突破工具属性，向拥有闭环执行与长期规划能力的高阶智能形态演进，成为AGI的核心载体。



02 全球行业领袖共识

英伟达CEO黄仁勋等科技巨头均断言，智能体是AI的**下一个颠覆性浪潮**。从学术顶会到产业峰会，智能体已成为定义未来十年AI格局的关键词，全球资本与研发资源正加速向这一方向倾斜与汇聚。



03 前沿技术落地示范

以OpenAI Deep Research为标杆，智能体已实现**自主检索信息、跨源整合数据、生成深度洞察报告的能力**。这标志着智能体不再停留在理论层面，而是已具备解决复杂认知任务的实际生产力。

研究路径：更智能、更长程

2024

把 Chatbot 变成 Agent

单 Chatbot 使用场景有限，通过多智能体系统，多个 Chatbots 协作完成复杂多变的现实世界任务。

2025/08

训练原生 Agent

核心参与 Tongyi DeepResearch Agent 训练，30B 模型在前沿榜单上与 OpenAI 持平

2025/01

把 Reasoner 变成 Agent

把 Python 作为交互语言，把不会工具调用的 DeepSeek-R1 变成会思考、会使用工具的 Agent

2026

优化数据与范式

研究并开源高质量数据合成方案；探索主动式上下文管理，有效解决长程任务协作难题。

2024: 把 Chatbot 变成 Agent


单个 Chatbot

只能聊天



多个 Chatbot 组成系统

调用工具，解决复杂任务



多智能体协同 + 工具调用 = 解决复杂任务

2024: 把 Chatbot 变成 Agent

原先：人工手工设计多智能体系统
灵活性低，成本高



MAS-GPT：多智能体系统生成器
不同用户 Query 进入，自动生成不同的智能体系统



最终效果




从人工设计到自动生成，让多智能体系统更灵活、更低成本，并显著提升推理表现

2025: 把 Reasoner 变成 Agent


问题：一次性生成完整系统，缺乏灵活性和适应性

解法：一边与环境交互，一边决定下一步要交互什么



2025: 把 Reasoner 变成 Agent

Humanity's Last Exam 人类最后的考试 (500+机构, 1000+主题)



超越 OpenAI Deep Research

首次踏上30分门槛

2025: 把 Reasoner 变成 Agent

SciMaster: 通用科研智能体助手

“边思考边交互”



AI 科学家

从一个问题开始，SciMaster带你探索科学答案

通用助手

深度调研

SciMaster S

人类是否有机会实现永生?

帮我做一份新能源电池的市场+科研双视角情报简报

如果黑洞会唱歌，它的声音会是什么样子?

用五岁小孩能听懂的方式解释一下量子纠缠,但要有诗意

SciMaster-你身边的AI科学家朋友 查看SciMaster产出的 [精选调研报告](#)

科研工具
+知识库




上传文件

工具

知识库

AI 工具（AlphaFold等）



2025年世界人工智能大会发布

链接科研工具数: **1000+** / 首月注册用户数: **1.6万+** / 已完成科研问答: **15万+**

2025: 把 Reasoner 变成 Agent

SciMaster Vibe Writing 产品：搜paper、想idea、做实验、分析、画图、写作全流程平台

The screenshot displays the SciMaster Vibe Writing interface, which is designed for managing research projects and writing documents. The interface is divided into several panels:

- Left Panel (Project Files and Outline):** Contains a file explorer showing project files like `checklist.tex`, `main.pdf`, `main.tex`, `main.tex.tmp`, `mindmap.md`, `neurips_2026.pdf`, `neurips_2026.sty`, `neurips_2026.tex`, `p1_v1.pdf`, and `p1_v2.pdf`. Below this is an outline of the document structure, including sections like "LongSeeker: Learning Elastic Context...", "Introduction", "Related Work", "LLM-Based Agents", "Web Agents for Long-Horiz...", "Context Management for LL...", "Sliding window and thr...", "Periodic summarization.", "Proactive curation.", "Our approach.", "Method", "Problem Formulation", "Standard ReAct.", "Context-ReAct.", and "Atomic Meta-Operations".
- Center Panel (Document Editor):** Displays the LaTeX source code for `main.tex`. The code includes package declarations, document class settings, and the beginning of the document structure, such as `\documentclass{article}`, `\usepackage{neurips_2026}`, and `\begin{document}`.
- Right Panel (Task List and Paper Preview):** Contains a task list for the document, including steps like "1. 直接读取完整文件, 将整个文件内容写成干净版本", "2. 但文件有1547行, 最高效的办法是: 仅修改第一个版本 (L109-L115) 的两个子章节, 然后删除 L503 之后的所有重复内容", and "3. 读取 L1-L462 全文 (第一版本)". Below the task list is a preview of the document, showing the title "LongSeeker: Learning Elastic Context Orchestration for Long-Horizon Search" and the abstract.

2025: 训练原生的 Agent

BrowseComp: Agent 深度调研能力的试金石，考察模型执行几十上百步搜索的能力

April 10, 2025 Publication Release Research

BrowseComp: a benchmark for browsing agents

A simple and challenging benchmark that measures the ability of AI agents to locate hard-to-find information.

Model	Accuracy (%)
GPT-4o	0.6
GPT-4o w/ browsing	1.9
GPT-4.5	0.9
OpenAI o1	9.9
Deep research*	51.5

*Note that the Deep Research model is trained on data that specifically teaches the model to be good at BrowseComp tasks.

条件1
Between 1990 and 1994 inclusive, what teams played in a soccer match with a Brazilian referee had
条件2
four yellow cards, two for each team where three of the total four were not issued during the first
条件3
half, and four substitutions, one of which was for an injury in the first 25 minutes of the match. 条件4
条件5



 Answer: Ireland v Romania

当时最强模型o1 得分 9.9分

最强专项系统 Deep Research 51.5分

2025: 训练原生的 Agent

研发并开源阿里通义深度调研模型Tong DeepResearch，智能体能力与OpenAI最强模型持平



2025: 训练原生的 Agent

研发并开源阿里通义深度调研模型Tong DeepResearch，智能体能力与OpenAI最强模型持平



Tongyi Lab

@Ali_TongyiLab · Sep 17

1/7 We're launching Tongyi DeepResearch, the first fully open-source Web Agent to achieve performance on par with OpenAI's Deep Research with only 30B (Activated 3B) parameters! Tongyi DeepResearch agent demonstrates state-of-the-art results, scoring 32.9 on Humanity's Last Exam,

Show more

Humanity's Last Exam

Tongyi DeepResearch (30B)	32.9
OpenAI o1	29.8
OpenAI o1-mini	26.9
OpenAI o3-mini	26.9
OpenAI o3	26.6
OpenAI o1.5	21.2
OpenAI o1.5	21.1

BrowseComp

OpenAI o1	51.5
Tongyi DeepResearch (30B)	43.4
OpenAI o1-mini	30.0
OpenAI o3-mini	28.3
OpenAI o3	26.4
OpenAI o1.5	14.1
OpenAI o1.5	12.2

BrowseComp-ZH

Tongyi DeepResearch (30B)	46.7
OpenAI o1	42.9
OpenAI o1-mini	37.5
OpenAI o3-mini	29.1
OpenAI o3	28.8

WebWalkerQA

Tongyi DeepResearch (30B)	72.2
OpenAI o1	71.7
OpenAI o1-mini	65.6
OpenAI o3-mini	63.0
OpenAI o3	61.7
OpenAI o1.5	61.2

GAIA

Tongyi DeepResearch (30B)	70.9
OpenAI o1	68.8
OpenAI o1-mini	67.4
OpenAI o3-mini	66.9
OpenAI o3	63.1
OpenAI o1.5	57.7

xbench-DeepSearch

Tongyi DeepResearch (30B)	75.0
OpenAI o1	71.0
OpenAI o1-mini	70.0
OpenAI o3-mini	69.0
OpenAI o3	67.0
OpenAI o1.5	65.0

FRAMES

Tongyi DeepResearch (30B)	90.6
OpenAI o1	84.9
OpenAI o1-mini	83.7
OpenAI o3-mini	80.7
OpenAI o3	78.9
OpenAI o1.5	72.0

SimpleQA

Tongyi DeepResearch (30B)	98.6
OpenAI o1	93.6
OpenAI o1-mini	93.5
OpenAI o3-mini	88.3
OpenAI o3	55.1
OpenAI o1.5	50.5
OpenAI o1.5	50.3

118

605

3.6K

1.7M

03:51

帖子

Emad

@EMostaque

This is a bigger deal than DeepSeek R1 imo, dig in and check it out

翻译帖子

Tongyi Lab

@Ali_TongyiLab · 3小时

1/7 We're launching Tongyi DeepResearch, the first fully open-source Web Agent to achieve performance on par with OpenAI's Deep Research with only 30B (Activated 3B) parameters! Tongyi DeepResearch agent demonstrates state-of-the-art results, scoring 32.9 on Humanity's Last Exam,...

显示更多

03:49 · 2025/9/17 位于 Earth · 472 次查看

2 个转帖 7 喜欢 3 个书签

发布你的回复

11:16

Stanford NLP Group

Stanford NLP Group

@stanfordnlp

Computational Linguists—Natural Language—Machine Learning @chrmmanning @jurafrsky @percyliang @ChrisGPotts @tatsu_hashimoto @MonicaSLam @Diyi_Yang @StanfordAILab

Stanford, CA, USA nlp.stanford.edu

2010年2月 加入

295 正在关注 171,699 关注者

Litu Ou, Lei Li 和 Junyang Lin 关注了此账号

帖子

回复

视频

照片

Stanford NLP Group 已转帖

Tongyi Lab

@Ali_TongyiLab · 10小时

1/7 We're launching Tongyi DeepResearch, the first fully open-source Web Agent to achieve performance on par with OpenAI's Deep Research with only 30B (Activated 3B) parameters! Tongyi DeepResearch agent demonstrates state-of-the-art results, scoring 32.9 on Humanity's Last Exam,...

显示更多

118



605

3.6K

1.7M

2026: 智能体爆发

头部厂商争抢 Agent 高地, Agentic Benchmarks 几乎占据 Model Release 的全部测评



推理、搜索、工具调用、代码、终端、GUI、视觉...

个人经验：数据很重要/范式有缺陷

数据是燃料，是最重要的核心资产之一



高质量数据驱动更强的智能能力上限

当前 ReAct 上下文不断累积的范式存在缺陷


用户 Query



上下文 (不断累积)

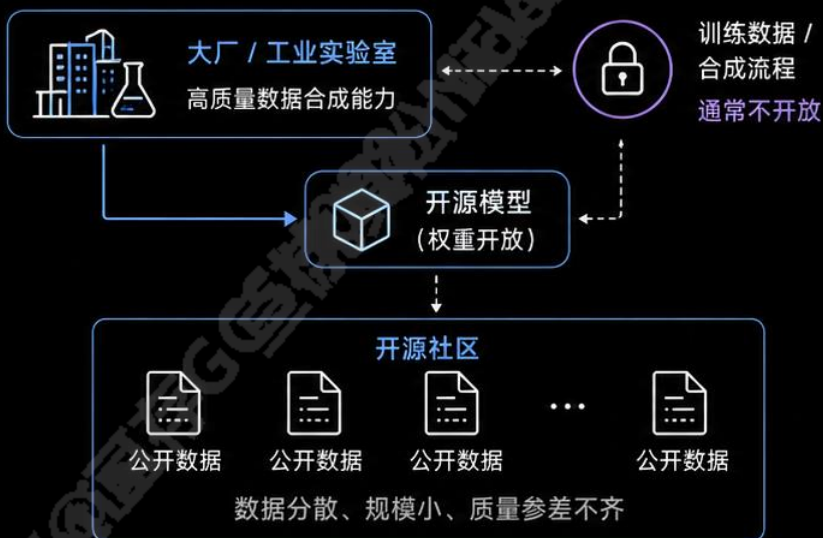


- 上下文长度迅速增长
- 关键信息被稀释，信噪比降低
- 噪声和冗余信息干扰推理
- 超长上下文带来更高的计算成本
- ...




01 高质量深度搜索数据合成

开源现状：模型开源，数据常常不开源



- 高质量数据合成技术往往掌握在大厂手里
 - 模型权重可以开源，但训练数据通常不公开
 - 社区虽有部分开源数据，但整体质量偏低
- 因此，数据成为限制开源进展的关键瓶颈。

BrowseComp 对比（越高越好）



公开数据方案与闭源高质量数据方案之间存在明显性能差距。




开源社区一直缺少一份高质量、可复用的数据，导致这方面的开源进展不多。

没有高质量数据，开源 deep research / agent 系统的发展会明显受限。

01 高质量深度搜索数据合成


高质量、高难度的数据合成方案



- 基于事实：基于真实网页信息构造
- 可控：难度/信息丰富度可由算法控制
- 可拓展：网络信息量庞大

01 高质量深度搜索数据合成

产生高质量、低噪声、长程轨迹



Teacher :

每步工具调用的同时，压缩上一步的工具返回

Student :

基于未压缩的上下文，预测 Teacher （在压缩上下文的情况下）产生的工具调用

01 高质量深度搜索数据合成


OpenSeeker：首个在深度搜索上达到前沿性能且同时开源模型和数据的工作

OpenSeeker: Democratizing Frontier Search Agents by Fully Open-Sourcing Training Data

Yuwen Du^{1,*}, Rui Ye^{1,*,#,†}, Shuo Tang¹, Xinyu Zhu¹, Yijun Lu¹, Yuzhu Cai¹, Siheng Chen^{1,†}

¹Shanghai Jiao Tong University, *Equal Core Contributions, #Project Lead

[†]Corresponding Authors: yr991129@sjtu.edu.cn, sihengc@sjtu.edu.cn



BrowseComp-ZH 超过
阿里 Tongyi DeepResearch

01 高质量深度搜索数据合成

OpenSeeker：首个在深度搜索上达到前沿性能且同时开源模型和数据的工作

Table 1: Comparisons among our OpenSeeker and other search agents. ‘# Samples’ denotes the number of total training data samples; ‘# OS Samples’ denotes the number of open-source data samples; ‘Training’ denotes training techniques (CPT: continual pre-training, SFT: supervised fine-tuning, RL: reinforcement learning); ‘Academic’ denotes whether conducted by pure academic team (✓: Yes, ×: No); ‘BC-ZH’ denotes BrowseComp-ZH; ‘WideSearch’ denotes item F1 result on the English subset. With simple SFT, OpenSeeker-v1-30B-SFT even surpasses Tongyi DeepResearch on BrowseComp-ZH which is trained via CPT, SFT, and RL. **Among all SFT-based agents, OpenSeeker performs significantly best.**

Model Name	# Samples	# OS Samples	Training	Academic	BrowseComp	BC-ZH	xbench	WideSearch
Closed-Source Proprietary Models								
Claude-4-Sonnet	?	0	?	×	14.7	22.5	-	62.0
Claude-4.5-Sonnet	?	0	?	×	24.1	42.4	-	-
Claude-4-Opus	?	0	?	×	18.8	37.4	-	-
OpenAI-o3	?	0	?	×	49.1	68.7	-	60.0
OpenAI Deep Research	?	0	?	×	51.5	42.9	-	-
GPT-5-High	?	0	?	×	54.9	63.0	-	-
Open-Source Models > 30B								
Kimi-K2-Instruct-1T	?	0	?	×	14.1	28.8	-	59.9
DeepSeek-V3.1-671B	?	0	?	×	30.0	49.2	71.2	-
DeepSeek-V3.2-671B	?	0	?	×	51.4	65.0	-	-
GLM-4.6-357B	?	0	?	×	45.1	49.5	-	-
GLM-4.7-357B	?	0	?	×	52.0	66.6	-	-
Minimax-M2-230B	?	0	?	×	44.0	48.5	-	-
~30B Models								
WebDancer-32B	?	0	SFT + RL	×	3.8	18.0	-	-
MiroThinker-32B-v0.1	147 k	147 k	SFT	×	10.6	13.8	-	-
MiroThinker-32B-v0.1	147 k	147 k	SFT + RL	×	13.0	17.0	-	-
DeepDive-32B	4.1 k	4.1 k	SFT + RL	×	15.3	29.7	51.8	-
WebSailor-32B	?	0	SFT + RL	×	10.5	25.5	53.3	-
WebSailor-V2-30B	?	0	SFT	×	24.4	28.3	61.7	-
WebSailor-V2-30B	?	0	SFT + RL	×	35.3	44.1	73.7	-
WebLeaper-30B	15 k	0	SFT	×	27.7	-	66.0	44.1
Tongyi DeepResearch	?	0	CPT + SFT + RL	×	43.4	46.7	75.0	-
OpenSeeker-v1-30B-SFT	11.7 k	11.7 k	SFT	✓	29.5	48.4	74.0	59.4

仅靠 11.7k 样本

仅靠 SFT

资源极其有限 只训了一把

唯一一个纯学术团队

部分榜单与经过持续预训练、SFT、RL的

Tongyi DeepResearch 持平

公平对比

Tongyi 方案

24.4 vs 29.5

28.3 vs 48.4

01 高质量深度搜索数据合成


OpenSeeker-v2：在深度搜索场景上全面超过大厂团队的模型

OpenSeeker-v2: Pushing the Limits of Search Agents with Informative and High-Difficulty Trajectories

Yuwen Du^{1,*}, Rui Ye^{1,*,#,†}, Shuo Tang¹, Keduan Huang¹, Xinyu Zhu¹, Yuzhu Cai¹, Siheng Chen^{1,†}

¹Shanghai Jiao Tong University, *Equal Core Contributions, #Project Lead

[†]Corresponding Authors: yr991129@sjtu.edu.cn, sihengc@sjtu.edu.cn



仅靠 10k 样本

仅靠 SFT

全面超过

CPT+SFT+RL Tongyi DR

01 高质量深度搜索数据合成

打造数据迭代进化飞轮

OpenSeeker-v1

OpenSeeker: Democratizing Frontier Search Agents by Fully Open-Sourcing Training Data

Yuwen Du^{1,2}, Rui Ye^{1,2,*}, Shuo Tang¹, Xinyu Zhu¹, Yijun Lu¹, Yuzhu Cai¹, Siheng Chen^{1,†}
¹Shanghai Jiao Tong University, ^{*}Equal Core Contributions, [†]Project Lead
[†]Corresponding Authors: yr991129@sjtu.edu.cn, sihengc@sjtu.edu.cn

Abstract

Deep search capabilities have become an indispensable competency for frontier Large Language Model (LLM) agents, yet the development of high-performance search agents remains dominated by industrial giants due to a lack of transparent, high-quality training data. This persistent data scarcity has fundamentally hindered the progress of the broader research community in developing and innovating within this domain. To bridge this gap, we introduce **OpenSeeker**, the first fully open-source search agent (i.e., model and data) that achieves frontier-level performance through two core technical innovations: (1) Fact-grounded scalable controllable QA synthesis, which reverse-engineers the web graph via topological expansion and entity obfuscation to generate complex, multi-hop reasoning tasks with controllable coverage and complexity. (2) Denoised trajectory synthesis, which employs a retrospective summarization mechanism to denoise the trajectory, therefore promoting the teacher LLMs to generate high-quality actions. Experimental results demonstrate that OpenSeeker, trained (a single training run) on only 11.7k synthesized samples, achieves state-of-the-art performance across multiple benchmarks including BrowseComp, BrowseComp-ZH, xbench-DeepSearch, and WideSearch. Notably, trained with simple SFT, OpenSeeker significantly outperforms the second-best fully open-source agent DeepDive (e.g., 29.5% v.s. 15.3% on BrowseComp), and even surpasses industrial competitors such as Tongyi DeepResearch (trained via extensive continual pre-training, SFT, and RL) on BrowseComp-ZH (48.4% v.s. 46.7%). We fully open-source the complete training dataset and the model weights to democratize frontier search agent research and foster a more transparent, collaborative ecosystem.

Code <https://github.com/rui-ye/OpenSeeker>
Data <https://huggingface.co/datasets/OpenSeeker/OpenSeeker-v1-Data>
Model <https://huggingface.co/OpenSeeker/OpenSeeker-v1-30B-SFT>




Figure 1: OpenSeeker stands out as the only fully open-source agent that achieves competitive performance on four search benchmarks, remarkably accomplishing this via simple SFT in a single training trial.

OpenSeeker-v2

OpenSeeker-v2: Pushing the Limits of Search Agents with Informative and High-Difficulty Trajectories

Yuwen Du^{1,2}, Rui Ye^{1,2,*}, Shuo Tang¹, Keduan Huang¹, Xinyu Zhu¹, Yuzhu Cai¹, Siheng Chen^{1,†}
¹Shanghai Jiao Tong University, ^{*}Equal Core Contributions, [†]Project Lead
[†]Corresponding Authors: yr991129@sjtu.edu.cn, sihengc@sjtu.edu.cn

Abstract

Deep search capabilities have become an indispensable competency for frontier Large Language Model (LLM) agents, yet their development remains dominated by industrial giants. The typical industry recipe involves a highly resource-intensive pipeline spanning pre-training, continual pre-training (CPT), supervised fine-tuning (SFT), and reinforcement learning (RL). In this report, we show that when fueled with informative and high-difficulty trajectories, a simple SFT approach could be surprisingly powerful for training frontier search agents. By introducing three simple data synthesis modifications: scaling knowledge graph size for richer exploration, expanding the tool set size for broader functionality, and strict low-step filtering, we establish a stronger baseline. Trained on merely 10.6k data points, our OpenSeeker-v2 achieves state-of-the-art performance across 4 benchmarks (30B-sized agents with ReAct paradigm): **46.0% on BrowseComp**, **58.1% on BrowseComp-ZH**, **34.6% on Humanity's Last Exam**, and **78.0% on xbench**, surpassing even Tongyi DeepResearch trained with heavy CPT+SFT+RL pipeline, which achieves 43.4%, 46.7%, 32.9%, and 75.0%, respectively. Notably, OpenSeeker-v2 represents the first state-of-the-art search agent within its model scale and paradigm to be developed by a purely academic team using only SFT. We are excited to open-source the OpenSeeker-v2 model weights and share our simple yet effective findings to make frontier search agent research more accessible to the community.

Code <https://github.com/OpenSeeker/OpenSeeker>
Model <https://huggingface.co/PolarSeeker/OpenSeeker-v2-30B-SFT>




Figure 1: OpenSeeker-v2 achieves state-of-the-art performance within its model scale and paradigm on four representative benchmarks, remarkably accomplishing this via simple SFT and outperforming Tongyi DeepResearch that is trained via extensive continual pre-training, SFT, and RL.

Preprint. Technical Report.

潜在 Future Version

Self-Evolving Score Curve



01 高质量深度搜索数据合成

The image displays the OpenSeeker GitHub repository and its associated dataset page. The top section shows the repository interface, including the repository name 'OpenSeeker', the owner 'duyuwen-duen', and the commit history. A blue box highlights the 'Starred 756' button. The bottom section shows the dataset viewer for 'PolarSeeker/OpenSeeker-v1-Data', which includes a table of dataset statistics and a blue box highlighting the '1,393' download count.

Repository Information:

- Repository: OpenSeeker
- Owner: duyuwen-duen
- Public
- Starred: 756
- Watch: 4
- Fork: 58
- Branch: main
- Tags: 0

Dataset Information:

- Dataset: PolarSeeker/OpenSeeker-v1-Data
- Like: 48
- Follow: 10
- Tasks: Question Answering
- Modalities: Text
- Formats: json
- Languages: English
- Size: 10K - 100K
- ArXiv: arxiv:2603.15594
- Tags: agent

Dataset Viewer Statistics:

question	answer	number of tool calls	trajectory	trajectory
string · lengths	string · lengths	int64	list · lengths	string · c
13	1	1	3	2 values
3.15k	187	201	403	

Dataset Card:

- Dataset card
- Data Studio
- Files and versions
- xet
- Community 1

Dataset Viewer:

- Split (1)
- train · 11.7k rows
- Search this dataset
- Number of rows: 11,677
- Total file size: 3.22 GB
- Downloads last month: 1,393

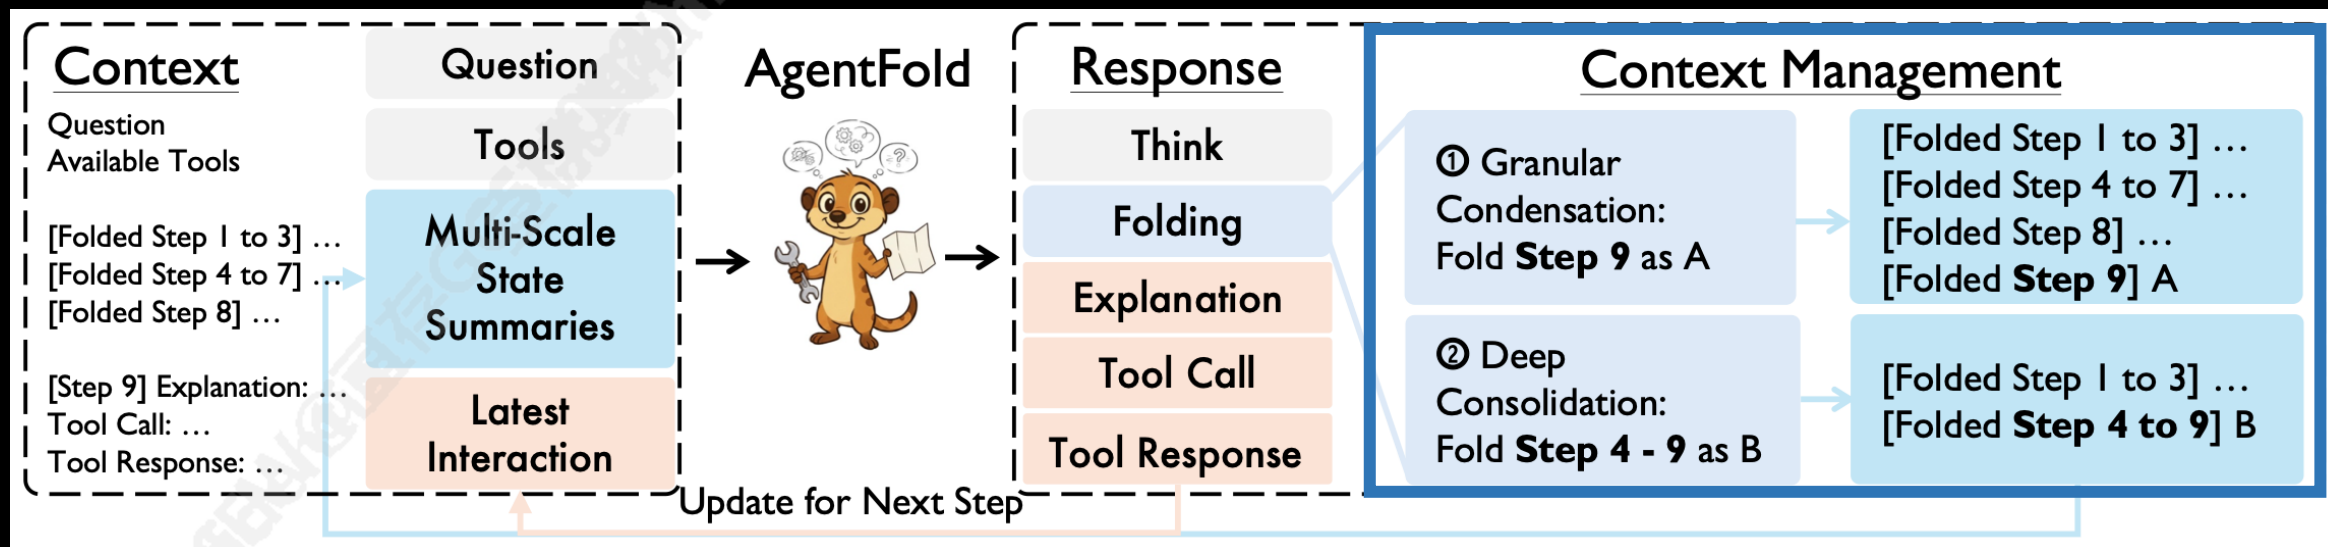
02 主动上下文管理新范式



人会**主动压缩与管理上下文**；若智能体只累加上下文，最终会面临**信噪比下降与上下文溢出**。

02 主动上下文管理新范式

AgentFold: 提出智能体上下文折叠概念, 通过折叠, 主动管理上下文



在普遍使用的 ReAct 基础上:

每次思考, 不仅要想下一步做什么, 还要想上下文要保留什么 (两个能力并不互斥)

02 主动上下文管理新范式

AgentFold: 提出智能体上下文折叠概念, 通过折叠, 主动管理上下文

Question

I am looking for a specialty food shop that has at least one location in both Mexico and California as of December 1, 2023. One of the locations in Mexico is a 7.8-kilometer walk from a McDonald's located at Anillo Perif. 137, Polanco, Polanco I Secc, Miguel Hidalgo, 11510 Ciudad de México, CDMX, Mexico as of December 1, 2023. The location in California is located within a market that held a Thanksgiving Floral Arrangement Class event in November 2023 as of December 1, 2023. What is the name of the specialty food shop?

Previous Steps

[Compressed Step 0] <Compressed summary omitted for visualization>

[Compressed Step 1 to 2] <Compressed summary omitted for visualization>

[Compressed Step 3] <Compressed summary omitted for visualization>

[Compressed Step 4] <Compressed summary omitted for visualization>

[Compressed Step 5] <Compressed summary omitted for visualization>

[Compressed Step 6 to 8] Steps 6-8 focused on obtaining GPS coordinates for the McDonald's at Blvd. Manuel Ávila Camacho 137 in Polanco, Mexico City, to define the 7.8 km radius around it. <Compressed summary omitted for visualization> **The coordinates remain unknown, and the cross-border shop investigation is preliminary.**

[Compressed Step 9] <Compressed summary omitted for visualization>

[Compressed Step 10] <Compressed summary omitted for visualization>

[Compressed Step 11] <Compressed summary omitted for visualization>

[Compressed Step 12] <Compressed summary omitted for visualization>

[Compressed Step 13] <Compressed summary omitted for visualization>

[Compressed Step 14] <Compressed summary omitted for visualization>

[Compressed Step 15] <Compressed summary omitted for visualization>

[Step 16] <lengthy web page content omitted for visualization>

<think> <omitted for visualization> </think>

<folding>

{"range": [6, 16], "summary": "**Multiple attempts (Steps 6-16)** to obtain GPS coordinates for the McDonald's at Blvd. Manuel Ávila Camacho 137, Polanco, Mexico City, via searches and visiting Infobelpro, Yelp, and TripAdvisor pages were unsuccessful. **No coordinates were found, hindering identification of the 7.8 km radius around the McDonald's.**"}

</folding>

Intermediate steps are

folded into one conclusion

<motivation>

Since direct methods to retrieve the GPS coordinates have failed, I am **shifting focus to identifying the specialty food shop by searching for potential candidates that operate in both Mexico and California**. This includes verifying if known candidates like Bazaar del Mundo have a Mexican location, and checking specific Polanco shops (e.g., Amado Pasteleria, Le Fromager) for California presence, as well as exploring concepts linking Northgate Market to cross-border operations.

</motivation>

Dynamically switch to
other search directions

<tool_call>


{"name": "search", "arguments": {"query": ["Bazaar del Mundo Mexico location", "specialty food shop Polanco Mexico California locations", "Northgate Market specialty food shop California Mexico", "Amado Pasteleria California", "Le Fromager California"]}}

</tool_call>

智能体意识到过去若干步骤均为无效操作, 选择将该部分上下文折叠; 并转变方向

02 主动上下文管理新范式

AgentFold: 提出智能体上下文折叠概念, 通过折叠, 主动管理上下文




➤ 30B 超越 DeepSeek-V3.1-671B 模型

➤ 可用来解决超长程任务, 并保持精炼上下文

02 主动上下文管理新范式


问题：如何折叠上下文中的任意部分？



💡 现有连续折叠范式难以支持对上下文任意位置的灵活管理。

02 主动上下文管理新范式

LongSeeker: 通用灵活精准的上下文管理范式



Standard Tools 之外
新增一套用于上下文管理的
Meta Tools:

- Keep
- Compress
- Rollback
- Snippet
- Delete

02 主动上下文管理新范式

LongSeeker：通用灵活精准的上下文管理范式

Managed Context

Question

Find a brand that:

- Uses a clinically supported element to solve customer issues
- Was acquired by a company whose founder graduated in the 1900s
- Has a co-founder from a university that recently announced a merger
- Sold products in 30+ countries by 2023

Previous Steps

Compressed Step 1 to 4:

Steps 1–3 focused on the co-founder education clue and L'Oréal as a likely acquirer, but found no co-founders from City or St George's universities. Step 4 revealed that L'Oréal acquired Medik8—a clinically supported skincare brand available in over 30 countries.

Step 5:

Motivation: We have confirmed L'Oréal as the acquirer ...

Standard tool call: {"name": "search_web", "arguments": ...}

Tool response: [Web search results, omitted due to length]

Deleted Step 6:

[Step deleted]

Rollbacked to Step 7:

Rollbacked to step 7 because the current approach is not yielding results and produces useless information. The searches for co-founders ...

Output

<think>

We've confirmed CeraVe is in 30+ countries. Still need to verify: 1) Saint Joseph's University's recent merger with Pennsylvania College of Health Sciences, 2) Tom Allison's graduation from there. Let's search for reliable news articles. Also double-check: CeraVe was acquired by L'Oréal in 2017; founder Eugène Schueller graduated in 1904 (matches the 1900s clue). The clinically supported element is ceramides.

</think>

<meta_tool_call>

{ "name": "skip", "arguments": {} }

</meta_tool_call>

<motivation>

We've verified CeraVe is in 30+ countries. Need to confirm the remaining two clues: Saint Joseph's University's recent merger and Tom Allison's graduation from there. We'll perform targeted web searches to gather evidence.

</motivation>

<standard_tool_call>

{ "name": "search_web", "arguments": { "query": ["Saint Joseph's University merger with Pennsylvania College of Health Sciences", "Tom Allison CeraVe co-founder Saint Joseph's University", "CeraVe co-founder Tom Allison education"] } }

</standard_tool_call>

Meta Tool Call

Standard Tool Call

02 主动上下文管理新范式

LongSeeker: 通用灵活精准的上下文管理范式


Theorem 3.1 (Expressive Completeness). *The meta-action set $\mathcal{O} = \{\text{SKIP}, \text{COMPRESS}, \text{ROLLBACK}, \text{SNIPPET}, \text{DELETE}\}$ is expressively complete: for any $H_{\text{in}}, H_{\text{target}} \in \mathcal{H}$, there exists a finite sequence of operations from $\mathcal{O}$ that transforms H_{in} into H_{target} .*

Proof. It suffices to show that COMPRESS alone is a universal string rewriting operator over $\mathcal{H}$. By definition, $\text{COMPRESS}(H, 1, |H|, \Sigma)$ replaces the *entire* history H with an arbitrary string $\Sigma \in \mathcal{V}^*$. Setting $\Sigma = H_{\text{target}}$ yields $H' = H_{\text{target}}$ in a single operation. Since a single element of $\mathcal{O}$ can reach any target from any source, the full set $\mathcal{O}$ is trivially complete. $\square$

理论上, 仅靠 Compress 一个 meta tool 可将任意上下文管理成任意上下文
其他 meta tools 则在实际实现上带来作用: 如降幻觉、省token等

02 主动上下文管理新范式

LongSeeker：通用灵活精准的上下文管理范式



显著超越前人方案。甚至能让 30B 模型超越闭源模型

03 总结



03 总结

感谢聆听

yr991129@sjtu.edu.cn 欢迎联系~