Open LM Agents:

Foundations to Learn and Discover

Akari Asai

Carnegie Mellon University / Allen Institute for AI

aasai@andrew.cmu.edu / akariasai.github.io

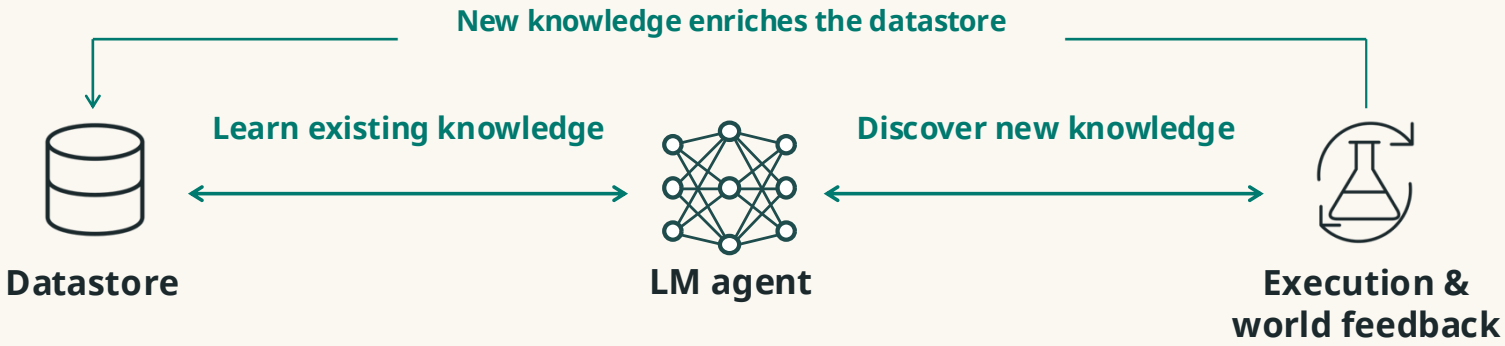




LTI Lightning Talk

LM agents for learning and discovery

Open-ended goals and long-horizon tasks



Tell me about recent research on agents.

Recent work focuses on planning [1], tool use, and learning from feedback [2].

Execute new experiments.

```
{ "experiment": "A",
  "results": [0.72, 0.81] }
```

Three foundations for open LM agents



DR Tulu

Scaling RL for long-horizon, open-ended tasks

Open agent training

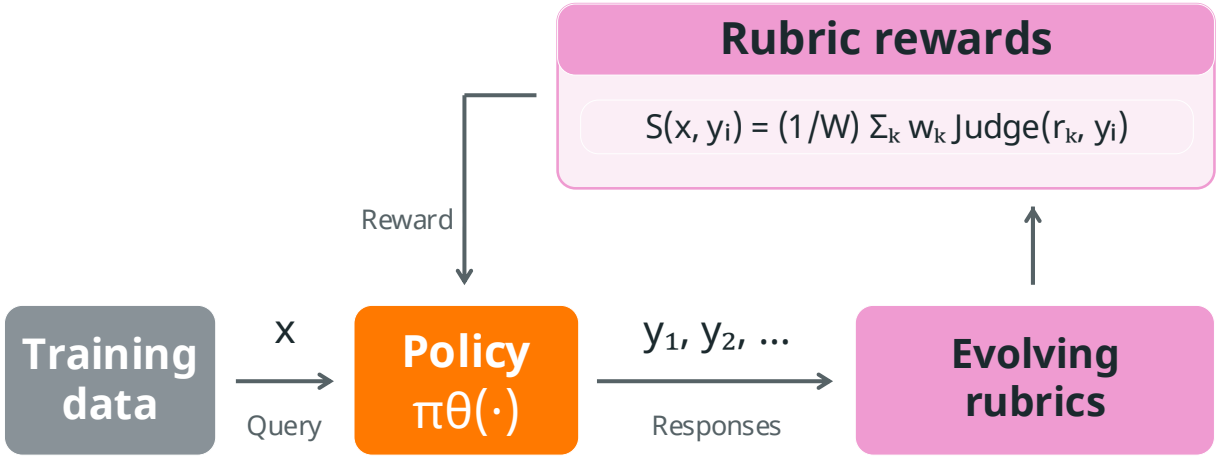
Long-horizon tasks

Safety & utility

Real-world deployment

AI for science

High-stakes applications



Feedback evolves with the agent.

Shao*, Asai* et al. DR Tulu: Reinforcement Learning with Evolving Rubrics for Deep Research. ICML (Oral) 2026.

DR Tulu

Open agent training

Long-horizon tasks

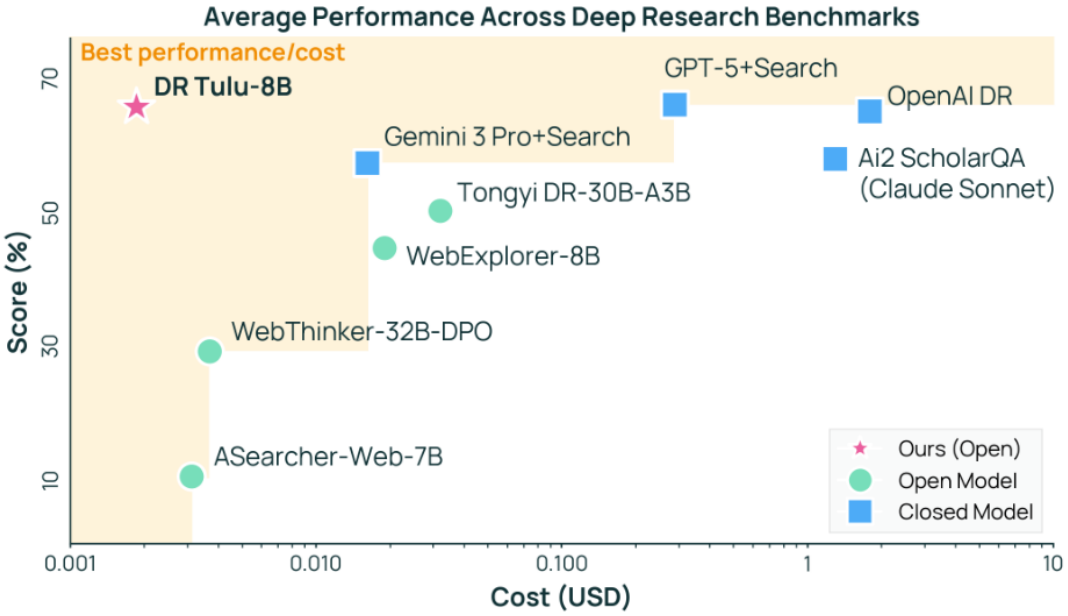
Safety & utility

Real-world deployment

AI for science

High-stakes applications

Parato-frontier by fully open 8B deep research agents



Shao*, Asai* et al. DR Tulu: Reinforcement Learning with Evolving Rubrics for Deep Research. ICML (Oral) 2026.

Open agent training

Long-horizon tasks

Safety & utility

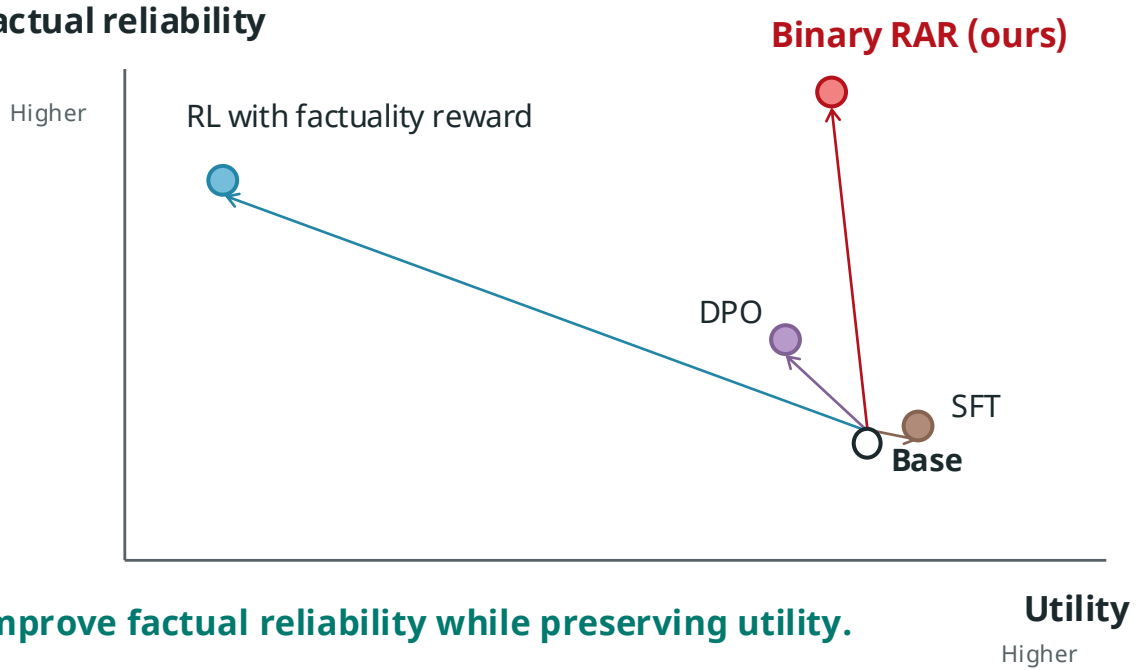
Real-world deployment

AI for science

High-stakes applications

Binary RAR

Balancing reliability and general capabilities of LM agents via RL



Improve factual reliability while preserving utility.

Chen et al. Train for Truth, Keep the Skills: Binary Retrieval-Augmented Reward Mitigates Hallucinations. ICML (Spotlight) 2026.

Open agent training

Long-horizon tasks

Safety & utility

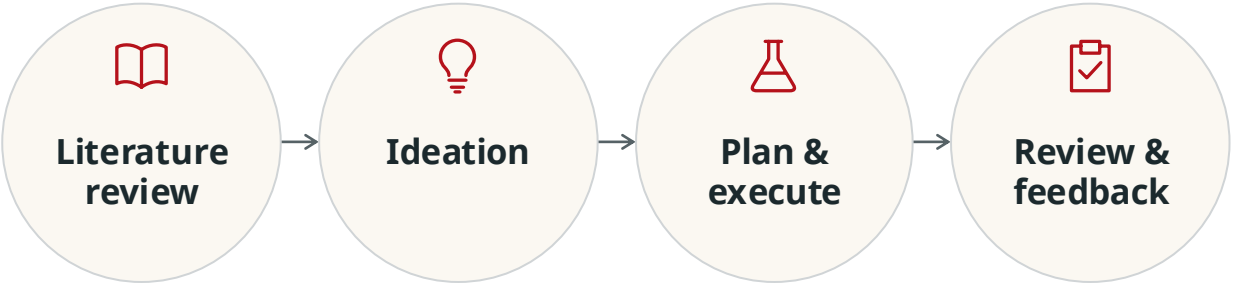
Real-world deployment

AI for science

High-stakes applications

AI for science

Benchmarking and building agents throughout the scientific workflow



OpenScholar

ScholarCatalyst

Into the Unknown

Protocol Design Agent

Science Terminal Agent

AI Reviewer

Open agent training

Long-horizon tasks

Safety & utility

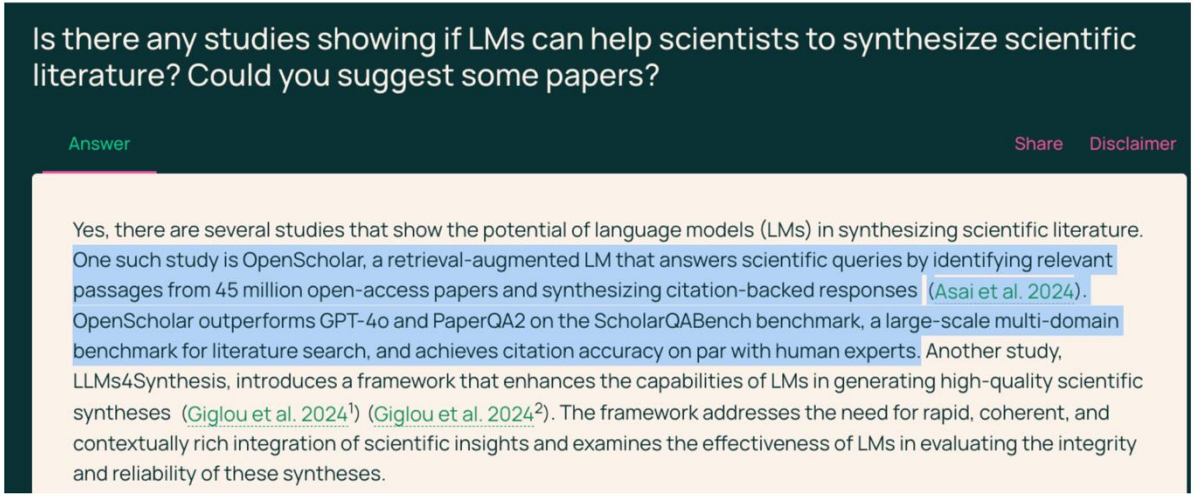
Real-world deployment

AI for science

High-stakes applications

OpenScholar

Can open models synthesize scientific literature?



Open Search Agents match expert performance

Asai et al. Synthesizing scientific literature with retrieval-augmented language models. Nature 2026.

PlateBench

Learning to design and execute wet-lab experiments

Open agent training

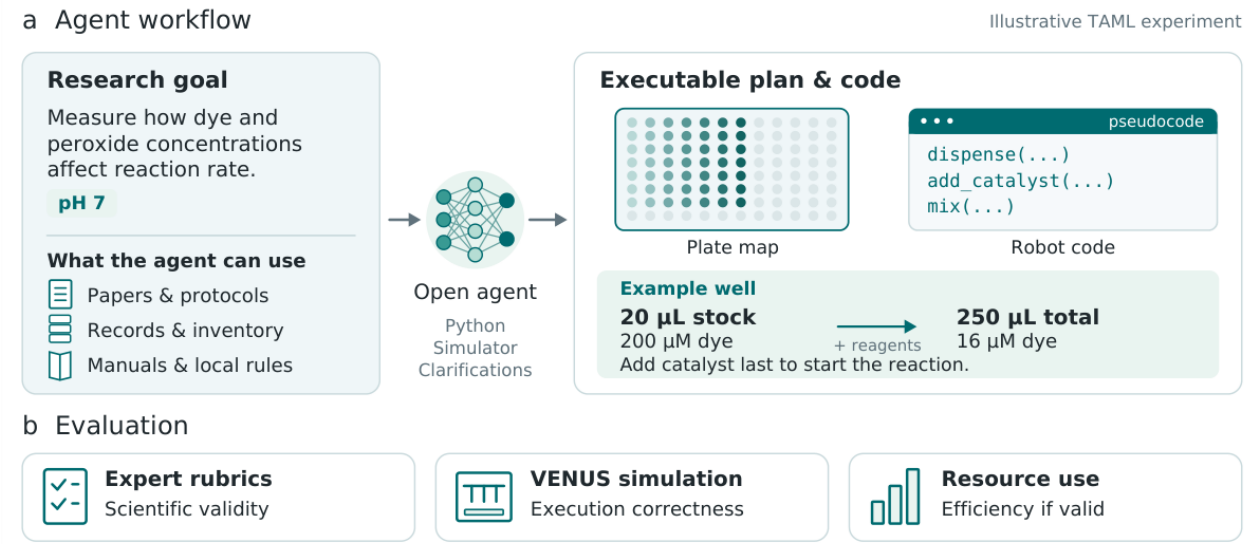
Long-horizon tasks

Safety & utility

Real-world deployment

AI for science

High-stakes applications



Ongoing work.

Moving forward

Open agent training

Long-horizon tasks

Safety & utility

Real-world deployment

AI for science

High-stakes applications

Long-horizon, multi-agent training

Synthesize tasks and environments that teach coordination while accounting for cost and resource constraints.

Safety across agent swarms

Monitor cumulative risks and design interventions for long-horizon agent swarms in high-stakes tasks.

AI for science: creativity and exploration

Explore new scientific ideas with world models and simulations, and test them through physical experimentation.