

MiroThinker-H1: A Research Agent That Checks Its Own Work at Every Step

Kabui, Charles

2026-03-19

[Read at ToKnow.ai](#)

A Research Agent That Checks Its Own Work at Every Step

MiroThinker-H1: dual verification for deep research from MiroMind AI

88.2% BrowseComp (SOTA, all models)	7.3K GitHub stars (open-source framework)	30B Open-source mini model (Apache 2.0)
---	---	---

March 19, 2026

ToKnow.ai

MiroMind AI released [MiroThinker-1.7](#) and its proprietary extension MiroThinker-H1, a family of research agents built for long, multi-step reasoning tasks like literature review, financial analysis, and scientific problem-solving. MiroThinker-1.7 comes in two open-source sizes (30B and 235B parameters, Apache 2.0) trained with an agentic mid-training stage that teaches structured planning, contextual reasoning, and tool use across a 256K context window with

up to 300 tool calls per task. The 235B model scores [74.0% on BrowseComp](#), [75.3% on BrowseComp-ZH](#), and [82.7% on GAIA-Val-165](#). MiroThinker-H1 adds a verification layer at two levels: local verification evaluates and refines intermediate reasoning steps during inference, while global verification audits the full chain of evidence before committing to a final answer. The result is an [88.2% score on BrowseComp](#), surpassing every major open and closed model on deep research benchmarks.

Verification actually makes the agent faster, not slower. By filtering out steps that produce no useful information, H1 takes fewer total interaction steps while reaching more accurate answers. For anyone doing serious research work with AI, whether financial due diligence, competitive analysis, or literature surveys, this means fewer hallucinated reasoning chains and more trustworthy final outputs. The open-source 30B mini model alone hits 72.3% on BrowseComp-ZH, making strong research-agent capabilities accessible without massive compute.

The “deep research” category is rapidly becoming the most contested space in AI. [OpenSeeker](#) recently showed that open training data can match proprietary search agents. MiroThinker-H1 takes a different approach: rather than better data, it bets on better self-correction. Both share the conviction that open-source research agents can compete with closed ones.

Sources:

- [MiroThinker-1.7 & H1 Blog Post \(MiroMind AI\)](#)
- [arXiv Paper: MiroThinker-1.7 & H1](#)
- [MiroThinker GitHub Repository \(7.3K stars\)](#)
- [MiroThinker-1.7 Models on HuggingFace](#)

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. Read more: [/terms-of-service](#)