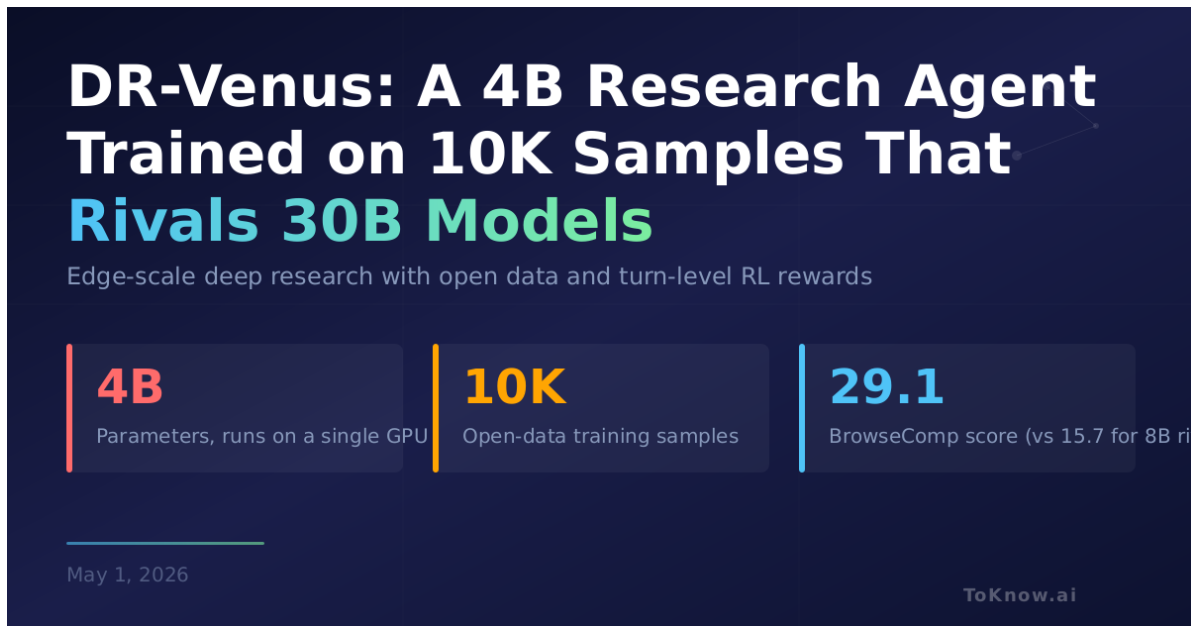


DR-Venus: A 4B Research Agent Trained on 10K Samples That Rivals Models 8x Its Size

Kabui, Charles

2026-05-01

[Read at ToKnow.ai](#)



DR-Venus: A 4B Research Agent Trained on 10K Samples That Rivals 30B Models

Edge-scale deep research with open data and turn-level RL rewards

4B Parameters, runs on a single GPU	10K Open-data training samples	29.1 BrowseComp score (vs 15.7 for 8B riv)
---	--	--

May 1, 2026

ToKnow.ai

Inclusion AI released [DR-Venus](#), a 4B-parameter deep research agent built entirely on open data that outperforms every prior agent under 9B parameters on multiple research benchmarks. The model is built on a Qwen3-4B backbone and trained in two stages. First, agentic supervised fine-tuning on cleaned trajectories from REDSearcher: 10,001 raw research traces are pruned to 9,365 correct ones, then upweighted via turn-aware resampling into 18,745 training instances.

Second, reinforcement learning using IGPO with turn-level rewards. Instead of a single score at the end of a multi-step research session, each turn gets two signals: an information gain reward (did the agent actually learn something new from that search or page visit?) and a format-aware penalty (is it maintaining structured notes and proper citations?). On BrowseComp, DR-Venus-4B-RL scores 29.1, nearly doubling the 15.7 from WebExplorer-8B-RL, a model twice its size. It handles up to 200 tool-call steps with 256K context at inference.

The practical upside: a capable research agent that fits on a single consumer GPU. At 4B parameters with GGUF quantized versions available, anyone with a laptop GPU can run an agent that autonomously searches the web, reads pages, and synthesizes answers. The 10K open-data requirement means any lab can reproduce and extend this without proprietary access.

The same [Inclusion AI team](#) behind LLaDA2.0-Uni is building a pattern: high-impact open-source releases that compress frontier capabilities into reproducible, small-scale packages.

Sources:

- [DR-Venus: Towards Frontier Edge-Scale Deep Research Agents with Only 10K Open Data \(arXiv\)](#)
- [DR-Venus GitHub Repository](#)
- [DR-Venus-4B-RL Model \(Hugging Face\)](#)
- [REDSearcher Open Trajectory Data](#)

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