

# AI Learns Scientific Taste: Training Models to Judge and Propose High-Impact Research

Kabui, Charles

2026-03-18

---

[Read at ToKnow.ai](#)

---



The banner features a dark blue background with white and light blue text. On the right side, there is a faint network diagram with nodes and connecting lines. The main title 'AI Learns Scientific Taste' is prominently displayed, with 'Taste' in a larger, light blue font. Below the title, a subtitle reads 'Training models to judge and propose high-impact research'. Three key statistics are presented in separate boxes: '700K Matched paper pairs', '228 HF upvotes, #1 Paper', and 'Beats GPT-5.2 At predicting paper impact'. The date 'March 18, 2026' and the 'ToKnow.ai' logo are located at the bottom left and right respectively.

## AI Learns Scientific Taste

Training models to judge and propose high-impact research

- 700K** Matched paper pairs
- 228** HF upvotes, #1 Paper
- Beats GPT-5.2** At predicting paper impact

March 18, 2026 ToKnow.ai

OpenMOSS at Fudan University trained AI to develop scientific taste: the ability to judge which research ideas have high potential impact and to propose new ones. Their method, [Reinforcement Learning from Community Feedback \(RLCF\)](#), uses citation patterns as training signal. They built [SciJudgeBench](#), 696,758 field- and time-matched paper pairs from 2.1 million arXiv publications. A “Scientific Judge” model predicts which paper in each pair will receive

more citations, controlling for field, year, and venue so the signal reflects quality rather than popularity. A second model, “Scientific Thinker,” uses the Judge as a reward signal to generate research ideas rated higher in impact than baselines. [Scientific Judge outperforms GPT-5.2 and Gemini 3 Pro](#) at predicting paper impact, generalizing to future years, unseen fields, and peer-review preferences. All models and data are [openly released](#).

Most AI-for-science work focuses on running experiments or writing papers. This targets the upstream question: what should researchers work on? A lab could use Scientific Judge to screen proposals or prioritize grants. The approach is data-efficient, needing just 11,700 training samples and one round of fine-tuning. The paper earned [228 upvotes and #1 Paper of the Month](#) on Hugging Face, with [263 GitHub stars](#) in days.

Citations encode visibility and collaboration networks, not just quality. If AI optimizes for what the community rewards, it risks reinforcing mainstream consensus over the unconventional ideas that drive breakthroughs. As [one commenter noted](#), most scientific breakthroughs come from stubborn people going against the crowd.

Read more: [Without Wonder, LLMs Are Analysts, Not Thinkers](#)

Sources:

- [AI Can Learn Scientific Taste \(arXiv paper\)](#)
- [Hugging Face Paper Page](#)
- [GitHub Repository](#)
- [OpenMOSS Model Collection](#)
- [SciJudgeBench Dataset](#)

---

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