

Intern-S1-Pro: A Trillion-Parameter Model That Beats GPT-5.2 on Scientific Reasoning

Kabui, Charles

2026-03-31

[Read at ToKnow.ai](#)



Shanghai AI Laboratory released Intern-S1-Pro, the first [one-trillion-parameter](#) scientific multimodal foundation model. It uses a Mixture-of-Experts (MoE) design with 512 experts, activating 8 per token for 22 billion active parameters. The model was pre-trained on 6 trillion tokens of text and image data, including 270 billion tokens of scientific image-text captions generated through a custom pipeline. On [SciReasoner](#), a benchmark spanning nine scientific

disciplines, Intern-S1-Pro scores 55.5, roughly 4x GPT-5.2's 13.6 and Gemini-3-Pro's 14.7. It also leads on materials prediction (MatBench: 72.8), chemistry (SmolInstruct: 74.8), and biology tasks (Mol-Instructions: 48.8). On general reasoning, it hits 93.1 on AIME-2025 and 86.6 on MMLU-Pro, matching top open-source models like [Qwen3-VL-235B](#).

A single massive model doing science better than separate specialists is the real result here. When trained on the same biology data, Intern-S1-Pro averaged 52.5 across biological tasks compared to 39.2 for a dedicated biology model, because the general reasoning capacity helps it use specialized data more effectively. Two architecture tricks keep the trillion-scale training stable: grouped routing that guarantees perfect load balance across GPUs, and a Straight-Through Estimator that lets all 512 expert routers learn on every token instead of only the 8 selected ones.

The trend in scientific AI has been building separate specialist models per domain. Intern-S1-Pro suggests the opposite: scale one model large enough and train it jointly, and the general intelligence helps rather than hurts. For fields like drug design and materials discovery, one API call to a generalist model may soon replace an ensemble of narrow tools.

Sources:

- [Intern-S1-Pro: Scientific Multimodal Foundation Model at Trillion Scale \(arXiv\)](#)
- [Intern-S1-Pro on HuggingFace](#)
- [Intern-S1 GitHub Repository](#)
- [HuggingFace Daily Papers, March 27, 2026](#)

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