BraintankAI ASIA

DISPATCH · OLYMPIAD REASONING · 3 MIN

HONG KONG · Chinese University of Hong Kong

A compact recipe and heavy self-checking take a mid-sized model to olympiad gold level

Given a compact training recipe, an existing reasoning model scored 35 of 42 on IMO 2025 and USAMO 2026 problems, but only with a long, costly draft-and-check loop.

14 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Shubham Dhage

Yafu Li and colleagues report that a mid-sized language model, given a compact three-stage recipe, can reach gold-medal-level scores in two mathematics olympiads, provided it spends heavily on drafting, checking and rewriting its own proofs. Their system, SU-01, scored 35 of 42 points on the problems of both the 2025 International Mathematical Olympiad (IMO) and the 2026 USA Mathematical Olympiad (USAMO), as marked by human experts. The 2026 paper is an arXiv preprint, not yet peer reviewed. The authors note that earlier systems reaching silver or gold IMO standard, such as AlphaProof, AlphaGeometry 2 and Gemini Deep Think, relied on larger search and verification budgets. SU-01 applies one compact recipe to an existing model.

Olympiad marking rewards a complete, rigorous argument, not just a correct final number. The starting model, P1-30B-A3B, is a mixture of experts: a router sends each piece of text to a selection of specialised sub-networks rather than through the whole network. Its label conventionally means about 30 billion parameters (the numbers adjusted during training), roughly 3 billion of them active at a time. It was already good at final answers, scoring 69.3% on the AnswerBench test. It was weak at proofs, managing 20.0% overall on IMO-ProofBench and 6.2% on its advanced split.

The recipe begins with supervised fine-tuning, in which the model imitates about 338,000 worked examples written by a larger model, DeepSeek-V3.2-Speciale. The examples were ordered by perplexity, a measure of how surprising a text is to the model, with the least familiar first. The authors tested this ordering on a different data mix. It scored 55.8 on AnswerBench, against 39.5 for random order and 24.3 for familiar-first. Compared with random order, it also cut the share of answers truncated at the length limit, a sign of rambling, from 7.3% to 0.3%. Next comes reinforcement learning, which rewards successful attempts. The first reward is for a correct, automatically checked final answer; the later reward is for a whole proof judged valid by another model, DeepSeekMath-V2. This later phase also hands the model its own failed proofs to fix and replays rare successes on hard problems. The whole stage used about 25,000 prompts over 200 training steps.

The third stage happens at answer time. In a loop adapted from earlier work by Huang and Yang, the same model drafts a solution, refines it and writes a structured bug report. It then accepts, rejects or revises the proof, repeating until the proof consistently passes its own check or the budget runs out. This is test-time scaling, meaning extra computation spent when answering, and it is costly. Models read and write text in word fragments called tokens. On USAMO 2026 the median first draft ran to about 106,000 tokens and the median refinement to 83,000.

Without the loop, SU-01 scored 57.6% on IMO-ProofBench. That put it ahead of every similar-sized model tested, the best of which managed 52.9%, and slightly ahead of its teacher at 56.0%. With the loop, SU-01 reached 70.2%, near Gemini 3.1 Pro Thinking at 72.6% but behind GPT-5.5-High at 80.7%, although no other model was given the loop. On benchmarks with checkable final answers it essentially tied the strongest similar-sized rival, Qwen3.6-35B-A3B, at 77.3% against 77.4%. The contests show how much the loop carries. Answering directly, SU-01 scored 21 on IMO 2025, a bronze-level result, and 15 on USAMO 2026. With the loop it reached 35 on both. That is exactly the IMO gold line, and it equals the top score reported among 340 human USAMO competitors. On the International Physics Olympiad of 2024 and 2025 it cleared the gold lines without the loop. Every similar-sized model cleared the 2024 line too, including the starting model.

Much remains open. Two of the 12 contest problems went unsolved. The authors say the model is weaker when the crux is preserving combinatorial structure or proving that a property holds throughout a process. No other model was tested on the two maths contests. Loop answers were marked by human experts, while direct answers were marked automatically. The recipe leans on outside models as teacher, proof judge and backup answer checker. The authors' case for a small training footprint counts data and training steps, not computing hours. They also acknowledge that a language model judging proofs is vulnerable to quirks of the judge and does not eliminate noise in the reward.

PAPER	*Achieving Gold-Medal-Level Olympiad Reasoning via Simple and Unified Scaling*
AUTHORS	Yafu Li, Runzhe Zhan, Haoran Zhang, Shunkai Zhang, Yizhuo Li, Zhilin Wang, Jiacheng Chen, Futing Wang, Xuyang Hu, Yuchen Fan, Bangjie Xu, Yucheng Su et al.
VENUE	arXiv preprint — not yet peer reviewed · 13 May 2026
LINK	arxiv.org/abs/2605.13301
INDEX	stacks.braintank.ai
CODE	github.com/Simplified-Reasoning/SU-01