

GrandCode: An AI That Beats Every Human in Live Competitive Programming

Kabui, Charles

2026-04-08

[Read at ToKnow.ai](#)

GrandCode: An AI That Beats Every Human in Competitive Programming

Multi-agent RL system wins 3 consecutive Codeforces live contests

- 3 / 3** live Codeforces contests won with first place
- 175 → 1** AI rank progression in one year (o3 to GrandCode)
- 4 Agents** jointly trained with Agentic GRPO algorithm

April 8, 2026

ToKnow.ai

DeepReinforce released [GrandCode](#), a multi-agent reinforcement learning system that placed first in three consecutive live [Codeforces](#) competitions: Round 1087 (March 21), Round 1088 (March 28), and Round 1089 (March 29, 2026). It beat every human participant, including legendary grandmasters. For context, the previous best AI result was Google’s Gemini 3 Deep Think at 8th place, and it wasn’t even tested under live conditions. GrandCode coordinates

specialized agent modules (hypothesis proposer, solver, test generator, summarizer) and trains them jointly through post-training and online test-time RL. A new algorithm called Agentic GRPO handles two hard problems in multi-step agent training: delayed rewards (you only learn if the solution was accepted after the full pipeline runs) and off-policy drift (the policy shifts significantly during long rollouts). The system is built on [Qwen](#) models.

The progression is striking: OpenAI's o3 reached 175th place in April 2025, Gemini 3.1 Pro climbed to 8th in February 2026, and now GrandCode holds first. Competitive programming requires deep algorithmic reasoning, creative problem-solving under time pressure, and the ability to debug on the fly. An engineer or student preparing for coding interviews can now study GrandCode's [published solutions](#) to see how an AI approaches problems that challenge the world's best programmers.

This is the competitive programming equivalent of AlphaGo beating Lee Sedol in Go. The multi-agent + RL architecture also signals something broader: jointly training specialized agents with reinforcement learning can produce results that no single model achieves alone. For a counterpoint on when multi-agent setups fall short, see [AI Agent Teams Look Amazing but Rarely Work](#).

Sources:

- [GrandCode Paper \(arXiv\)](#)
- [DeepReinforce Project Page](#)
- [GrandCode Solutions on GitHub](#)
- [Hugging Face Daily Papers](#)

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