

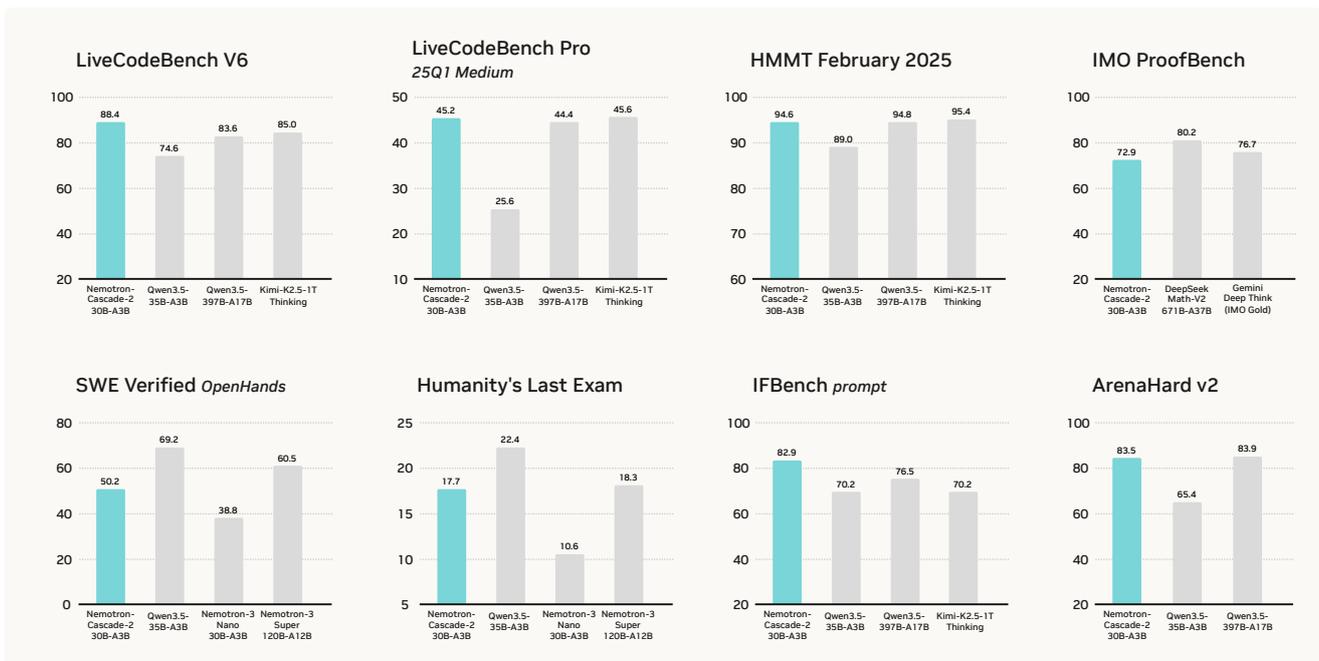
# Nemotron-Cascade 2: Post-Training LLMs with Cascade RL and Multi-Domain On-Policy Distillation

Zhuolin Yang\*, Zihan Liu\*, Yang Chen\*, Wenliang Dai\*, Boxin Wang\*, Sheng-Chieh Lin, Chankyu Lee, Yangyi Chen, Dongfu Jiang, Jiafan He<sup>‡</sup>, Renjie Pi, Grace Lam, Nayeon Lee, Alexander Bukharin, Mohammad Shoeybi, Bryan Catanzaro, Wei Ping\*<sup>†</sup>

## Abstract

We introduce Nemotron-Cascade 2, an open 30B MoE model with 3B activated parameters that delivers best-in-class reasoning and strong agentic capabilities. Despite its compact size, its mathematical and coding reasoning performance approaches that of frontier open models. It is the second open-weight LLM, after DeepSeek-V3.2-Special-671B-A37B, to achieve **Gold Medal**-level performance in the 2025 International Mathematical Olympiad (IMO), the International Olympiad in Informatics (IOI), and the ICPC World Finals, demonstrating remarkably high intelligence density with 20× fewer parameters. In contrast to Nemotron-Cascade 1, the key technical advancements are as follows. After SFT on a meticulously curated dataset, we substantially expand Cascade RL to cover a much broader spectrum of reasoning and agentic domains. Furthermore, we introduce multi-domain on-policy distillation from the strongest intermediate teacher models for each domain throughout the Cascade RL process, allowing us to efficiently recover benchmark regressions and sustain strong performance gains along the way. We release the collection of model checkpoint and training data.

- 🤗 **Nemotron-Cascade-2-30B-A3B**: the post-trained model based on Nemotron-3-Nano-30B-A3B-Base.
- 🤗 **Nemotron-Cascade-2-SFT-Data**: collection of SFT datasets for Nemotron-Cascade-2.
- 🤗 **Nemotron-Cascade-2-RL-Data**: collection of RL datasets for Nemotron-Cascade-2.



\*Equal contribution, with authors listed in reverse alphabetical order by first name.

<sup>‡</sup>Reviewed and scored our model-generated solutions for IMO 2025 as a gold medalist at the IMO 2015. Correspondence to: <jiafanhe19@ucla.edu>.

<sup>†</sup>Leads the effort. Correspondence to: <wping@nvidia.com>.

## Contents

|          |   |           |
|----------|---|-----------|
| <b>1</b> | <b>Introduction</b>                                       | <b>4</b>  |
| <b>2</b> | <b>Main Results</b>                                       | <b>4</b>  |
| <b>3</b> | <b>Supervised Fine-Tuning</b>                             | <b>6</b>  |
| 3.1      | Training Framework  | 6         |
| 3.1.1    | Overview  | 6         |
| 3.1.2    | Chat Template   | 6         |
| 3.2      | SFT Data Curation   | 7         |
| 3.2.1    | Math  | 7         |
| 3.2.2    | Code Reasoning  | 7         |
| 3.2.3    | Science   | 8         |
| 3.2.4    | Long Context  | 8         |
| 3.2.5    | General Chat  | 8         |
| 3.2.6    | Instruction Following                                     | 8         |
| 3.2.7    | Safety  | 8         |
| 3.2.8    | Conversational Agent                                      | 9         |
| 3.2.9    | Software Engineering Agent                                | 9         |
| 3.2.10   | Terminal Agent  | 9         |
| <b>4</b> | <b>Cascade RL and Multi-Domain On-Policy Distillation</b> | <b>9</b>  |
| 4.1      | Training Framework  | 9         |
| 4.1.1    | What determines the ordering of Cascade RL                | 10        |
| 4.1.2    | RL Training Configuration                                 | 10        |
| 4.2      | Instruction-Following Reinforcement Learning (IF-RL)      | 11        |
| 4.2.1    | Dataset   | 11        |
| 4.2.2    | Training recipe   | 11        |
| 4.3      | Multi-domain RL   | 12        |
| 4.4      | Multi-domain On-Policy Distillation (MOPD)                | 12        |
| 4.5      | Reinforcement Learning from Human Feedback (RLHF)         | 14        |
| 4.5.1    | Dataset   | 14        |
| 4.5.2    | Training recipe   | 14        |
| 4.5.3    | Hyper-parameters  | 14        |
| 4.6      | Long-context RL   | 15        |
| 4.7      | Code RL   | 15        |
| 4.7.1    | Data Curation   | 15        |
| 4.7.2    | Training Details  | 15        |
| 4.8      | Software Engineering Reinforcement Learning (SWE RL)      | 15        |
| 4.8.1    | Agentless RL  | 15        |
| 4.8.2    | Execution-based RL for Agentic SWE Scaffold               | 16        |
| <b>5</b> | <b>International Mathematical Olympiad (IMO)</b>          | <b>16</b> |
| 5.1      | IMO 2025  | 16        |
| 5.2      | IMO-ProofBench  | 17        |
| <b>6</b> | <b>Competitive Coding</b>                                 | <b>17</b> |
| 6.1      | IOI 2025 and ICPC World Finals 2025                       | 17        |
| 6.2      | Competitive Coding Benchmark Results                      | 18        |

|  |           |
|--|-----------|
| <b>7 Acknowledgments</b>   | <b>19</b> |
| <b>A Benchmarks and Evaluation Setups</b>                        | <b>20</b> |
| A.1 Math . . . . .   | 20        |
| A.1.1 Non-proof Math . . . . .                                   | 20        |
| A.1.2 Math Proof . . . . .                                       | 20        |
| A.2 Code Reasoning . . . . .                                     | 21        |
| A.3 Knowledge and STEM . . . . .                                 | 21        |
| A.4 Alignment and Instruction-Following . . . . .                | 22        |
| A.5 Long Context and Context Learning . . . . .                  | 22        |
| A.6 Agentic Tasks . . . . .                                      | 23        |
| A.7 Multilingual . . . . .                                       | 24        |
| <b>B Training Hyperparameters</b>                                | <b>24</b> |
| <b>C Prompt Templates</b>  | <b>26</b> |
| C.1 Prompt Templates for Test-Time Scaling on IOI 2025 . . . . . | 26        |
| C.2 HLE Judge Prompt . . . . .                                   | 27        |
| <b>D ELO Rating Analysis</b>                                     | <b>27</b> |
| <b>E IMO 2025 Model Solutions</b>                                | <b>30</b> |

## 1. Introduction

Reinforcement Learning (RL) (Guo et al., 2025; Ouyang et al., 2022) has emerged as the cornerstone of LLM post-training, driving advances in reasoning, agentic capabilities, and real-world problem-solving. As models are tasked with increasingly sophisticated requirements, the primary challenge lies in successfully incorporating a broader array of RL environments and very diverse reasoning and agentic tasks. Scaling RL to encompass multifaceted, real-world applications necessitates robust frameworks capable of handling varied reward signals and complex environmental feedback without destabilizing the training process.

Our previous work, Nemotron-Cascade 1 (Wang et al., 2025), introduced Cascade RL, a framework that orchestrates sequential, domain-wise RL training across specialized task domains. Cascade RL significantly simplifies the engineering complexity associated with multi-domain RL while achieving state-of-the-art performance across a wide range of benchmarks. The advantages of Cascade RL are threefold. First, domain-specific RL stages are remarkably resistant to *catastrophic forgetting*. They rarely degrade benchmark performance attained in earlier domains and may even improve it. Second, it allows RL hyperparameters and the training curriculum to be carefully tailored to each specific domain, enabling optimized learning dynamics and improved final performance. Third, task homogeneity within each RL stage also yields substantial compute savings, as response lengths and verification wall-clock times are more uniform within a domain than across multiple domains trained jointly.

In this work, we introduce Nemotron-Cascade 2, an open 30B Mixture-of-Experts (MoE) model with 3B activated parameters. Similar to its predecessor, Nemotron-Cascade 2 further scales Cascade RL on high-priority domains to preserve the benefits of domain-wise training, enabling us to push the limits of reasoning performance in key domains to state-of-the-art levels. Furthermore, we incorporate on-policy distillation (Xiao et al., 2026; Zeng et al., 2026) into Cascade RL training stages. By distilling knowledge from the best-performing intermediate teacher models within each specific domain during Cascade RL, this mechanism effectively recovers any benchmark regressions that can occur when training in increasingly complex RL environments. In addition, we integrate multi-domain RL into Cascade RL for groups of tasks with similar response formats and comparable verification costs, allowing them to be trained jointly to scale up for more RL environments and improve training efficiency when cross-task interference is minimal.

Our Nemotron-Cascade-2-30B-A3B achieves breakthrough performance in mathematical and coding reasoning, securing gold-medal results in both the 2025 International Mathematical Olympiad (IMO) and the International Olympiad in Informatics (IOI) despite being only a 30B MoE model,<sup>1</sup> while also delivering best-in-class performance across a broad range of benchmarks, including alignment, instruction-following, long context (e.g., 1M context window), and agentic tasks. See Table 1 for the full results. We fully open source the model weights, training data, and methodological details, enabling the research community to reproduce, analyze, and extend the proposed Cascade RL training paradigm.

We organize the remainder of this report as follows. Section §2 summarizes the main results. Section §3 describes the supervised fine-tuning (SFT) with details on data curation. Section §4 presents Cascade RL framework intergrated with the multi-domain on-policy distillation. Section §5 details the evaluation setup and results on IMO, while Section §6 presents the evaluation setup and results on IOI and the ICPC World Finals.

## 2. Main Results

We evaluate Nemotron-Cascade 2 on a comprehensive suite of benchmarks covering mathematical and coding reasoning, knowledge and STEM, alignment and instruction following, long-context understanding and in-context learning, multilingual capabilities, and agentic tasks. The main results are shown in Table 1, and the benchmarks and detailed evaluation setups are described in Appendix A.

---

<sup>1</sup>Our model is the second open-weight LLM, after DeepSeek-V3.2-Speciale-671B-A37B (Liu et al., 2025), to achieve gold-medal performance in both the IMO and IOI.

Table 1: **Main results.** Nemotron-Cascade-2-30B-A3B achieves gold-medal performance in both the IMO 2025 and IOI 2025, which demonstrate remarkably high intelligence density. † Numbers in brackets refers to Tool-Integrated Reasoning (TIR) results. ‡ For the baseline models, we use official numbers when available, otherwise evaluate them using the recommended settings.

| Benchmark<br>Metric: pass@1                  | Nemotron-3-Nano<br>30B-A3B | Nemotron-3-Super<br>120B-A12B | Qwen3.5<br>35B-A3B | Nemotron-Cascade-2<br>30B-A3B   |
|--|----------------------------|-------------------------------|--------------------|---|
| <b>Math</b>                                  |                            |                               |                    |   |
| IMO 2025                                     | –                          | –                             | –                  |  <b>35 pts</b> |
| IMO AnswerBench                              | 70.4 <sup>‡</sup>          | 77.2 <sup>‡</sup>             | 74.8 <sup>‡</sup>  | 79.3  |
| IMO ProofBench                               | –                          | –                             | –                  | 72.9  |
| AIME 2025                                    | 89.1                       | 90.2                          | 91.9 <sup>‡</sup>  | 92.4 (98.6) <sup>†</sup>  |
| AIME 2026                                    | 89.9 <sup>‡</sup>          | 89.8 <sup>‡</sup>             | 91.1 <sup>‡</sup>  | 90.9 (95.0) <sup>†</sup>  |
| HMMT Feb25                                   | 84.6 <sup>‡</sup>          | 93.7                          | 89.0               | 94.6  |
| <b>Code Reasoning</b>                        |                            |                               |                    |   |
| IOI 2025                                     | –                          | –                             | 348.6 <sup>‡</sup> |  <b>439.28</b> |
| ICPC World Finals 2025                       | –                          | –                             | –                  |  <b>10/12</b>  |
| LiveCodeBench v6 (2408-2505)                 | 68.3                       | 78.7                          | 74.6               | 87.2 (88.4) <sup>†</sup>  |
| LiveCodeBenchPro 25Q2 (Easy)                 | 54.5 <sup>‡</sup>          | 81.7 <sup>‡</sup>             | 81.1 <sup>‡</sup>  | 87.0 (89.3) <sup>†</sup>  |
| LiveCodeBenchPro 25Q2 (Med)                  | 3.50 <sup>‡</sup>          | 23.2 <sup>‡</sup>             | 17.8 <sup>‡</sup>  | 27.6 (36.8) <sup>†</sup>  |
| SciCode                                      | 33.3                       | 42.1                          | 38.0               | 36.4  |
| <b>Knowledge &amp; STEM</b>                  |                            |                               |                    |   |
| MMLU-Redux                                   | –                          | –                             | 93.3               | 86.3  |
| MMLU-Pro                                     | 78.3                       | 83.7                          | 85.3               | 79.8  |
| GPQA-Diamond                                 | 73.0                       | 79.2                          | 84.2               | 76.1  |
| HLE (no tool)                                | 10.6                       | 18.3                          | 22.4               | 17.7  |
| <b>Alignment &amp; Instruction Following</b> |                            |                               |                    |   |
| ArenaHard v2 (Avg.)                          | 67.7                       | –                             | 65.4 <sup>‡</sup>  | 83.5  |
| – Hard Prompt                                | 72.1                       | 73.9                          | 64.5 <sup>‡</sup>  | 88.2  |
| – Creative Writing                           | 63.2                       | –                             | 66.3 <sup>‡</sup>  | 78.7  |
| IFBench (prompt)                             | 71.5                       | 72.6                          | 70.2               | 82.9  |
| Scale AI Multi-Challenge                     | 38.5                       | 55.2                          | 60.0               | 45.3  |
| <b>Long Context &amp; Context Learning</b>   |                            |                               |                    |   |
| AA-LCR                                       | 35.9                       | 58.3                          | 58.5               | 39.1  |
| LongBench v2                                 | 39.6                       | –                             | 59.0               | 40.3  |
| NIAH@1M (RULER Subset)                       | 94.8                       | 98.3                          | 94.3 <sup>‡</sup>  | 99.0  |
| CL-Bench                                     | 12.0 <sup>‡</sup>          | –                             | 15.5 <sup>‡</sup>  | 12.2  |
| <b>Agentic</b>                               |                            |                               |                    |   |
| BFCL v4                                      | 53.8                       | –                             | 67.3               | 52.9  |
| $\tau^2$ -Bench                              | 49.0                       | 61.2                          | 81.2               | 58.9  |
| Terminal Bench 2.0                           | 8.5                        | 31.0                          | 40.5               | 21.1  |
| SWE Verified (OpenHands)                     | 38.8                       | 60.5                          | 69.2               | 50.2  |
| <b>Multilingual</b>                          |                            |                               |                    |   |
| MMLU-ProX                                    | 59.5                       | 79.4                          | 81.0               | 72.5  |
| WMT24++ (en -> xx)                           | 86.2                       | 86.7                          | 87.6 <sup>‡</sup>  | 84.1  |

From Table 1, Nemotron-Cascade-2-30B-A3B outperforms both the latest released Qwen3.5-35B-A3B (2026-02-24) (Qwen Team, 2026) and the larger Nemotron-3-Super-120B-A12B (2026-03-11) (Blakeman et al., 2025), and achieves best-in-class performance across benchmarks in mathematics, code reasoning, alignment, and instruction following. Notably, despite being only a 30B MoE model, Nemotron-Cascade 2 achieves gold-medal

Table 2: Performance of Nemotron-Cascade-2-30B-A3B model on IMO 2025, IOI 2025, and ICPC World Finals 2025 competitions. Nemotron-Cascade-2 model achieved solid gold medal on all these top-tier competitions. Our IMO 2025 solutions are evaluated by human expert (IMO 2015 Gold medalist) while IOI 2025 and ICPCWF 2025 solutions are verified through OnlineJudge with official testcases.

| Competition | P1 | P2             | P3  | P4  | P5    | P6 | Overall    | Medal |
|-------------|----|----------------|-----|-----|-------|----|------------|-------|
| IMO 2025    | 7  | 7 <sup>†</sup> | 7   | 7   | 7     | 0  | 35/42      | Gold  |
| IOI 2025    | 39 | 88.53          | 100 | 100 | 28.75 | 83 | 439.28/600 | Gold  |

| Competition            | A | B | C | D | E | F | G | H | I | J | K | L | Overall | Medal |
|------------------------|---|---|---|---|---|---|---|---|---|---|---|---|---------|-------|
| ICPC World Finals 2025 | + | - | + | + | + | + | - | + | + | + | + | + | 10/12   | Gold  |

<sup>†</sup> For IMO 2025 P2, we use LLM grader with reference solution and marking schema from ProofBench (Ma et al., 2025) due to the extensive analytic geometry approach of the model, which human expert could be hard to verify all the intermediate derivation steps.

performance on IMO 2025, IOI 2025 and ICPC World Finals 2025, results previously thought to be attainable only by frontier proprietary models (Gemini Team, 2025) (i.e., Gemini Deep Think) and frontier-sized open models (Liu et al., 2025) (i.e., DeepSeek-V3.2-Speciale-671B-A37B). The detailed performance of our model is reported in Table 2. It underperforms Qwen3.5-35B-A3B primarily on knowledge-intensive and agentic tasks, highlighting the importance of stronger knowledge-intensive pretraining and agentic RL in future work.

Nemotron-Cascade-2-30B-A3B also outperforms Nemotron-3-Nano-30B-A3B on nearly all benchmarks, even though both models are post-trained from the same pretrained model, Nemotron-3-Nano-30B-A3B-Base (NVIDIA, 2025). This result further demonstrates the effectiveness of our Cascade RL plus MOPD training pipeline.

### 3. Supervised Fine-Tuning

In this section, we describe the training framework and data curation process for supervised fine-tuning (SFT), the first stage of our post-training pipeline. This stage equips the model with foundational capabilities, including reasoning, conversational ability, instruction following, and agentic and software engineering skills.

#### 3.1. Training Framework

##### 3.1.1. Overview

Our SFT data spans a broad range of domains, including mathematics, coding, science, tool use, agentic tasks, and software engineering, as well as more general domains such as multi-turn dialogue, knowledge-intensive question answering, creative writing, role-playing, safety, and instruction following.

We pack all SFT samples into sequences of up to 256K tokens and train the model in a single stage. Empirically, we find that the SFT model reaches optimal performance after approximately 1.5 epochs. The SFT training hyperparameters can be found in Appendix B.

##### 3.1.2. Chat Template

Our chat template is depicted in Figure 1. There are two changes to the chat template compared with Nemotron-Cascade (Wang et al., 2025). First, we remove the `/think` and `/no_think` tags for simplicity. Second, we prepend an empty `<think></think>` block to activate the non-thinking mode.

For tool calling task, we specify all available tools in the system prompt within the `<tools>` and `</tools>` tags, and instruct the model to perform tool calls wrapped within the `<tool_call>` and `</tool_call>` tags.

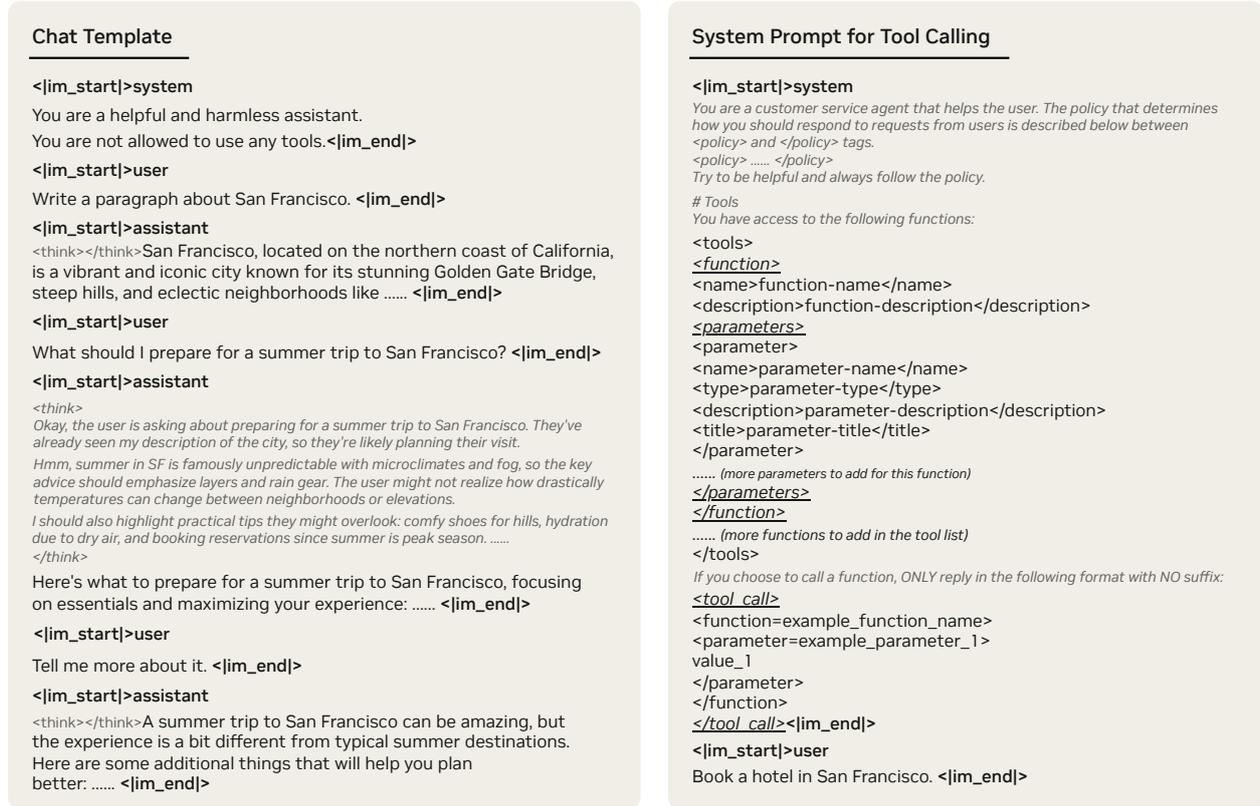


Figure 1: (Left) The chat template uses adjacent `<think></think>` tokens to indicate non-thinking mode, and a single `<think>` followed by `\n` to indicate thinking mode. (Right) For tool calling, the available tools are listed in the system prompt. The model is instructed to call tools within the `<tool_call>` and `</tool_call>` tags.

## 3.2. SFT Data Curation

### 3.2.1. Math

Our non-proof math prompts are primarily sourced from Nemotron-Cascade (Wang et al., 2025) and Nemotron-Math-v2 (Du et al., 2025), from which we collect 1.8M tool calling (i.e., python) samples and 1.9M non-tool samples, with responses generated by DeepSeek-V3.2 and DeepSeek-V3.2-Special (Liu et al., 2025), respectively. In addition, we collect 676K samples from the generation-selection category (without tool calling) of Nemotron-3-Nano (Blakeman et al., 2025), with responses generated by GPT-OSS-120B (Agarwal et al., 2025). In total, the competition math SFT comprises 1.8M tool-calling samples and 2.6M samples without tool use.

For mathematical natural language proof, we collect 98K mathematical proof problems from the AOPS split of Nemotron-Math-Proofs-v1 (Du et al., 2025). We generate multiple samples per problem to cover two capabilities including proof generation (410K) and proof verification (400K) using DeepSeek-V3.2-Special (Liu et al., 2025), resulting in a total of 816K samples.

### 3.2.2. Code Reasoning

Built on Nemotron-Cascade 1 (Wang et al., 2025), we curate approximately 165K unique coding prompts from several open-source datasets, including OpenCode-Stage2 (Huang et al., 2024), OpenCodeReasoning (Ahmad et al., 2025), and HardTests (He et al., 2025). These prompts are originally sourced from competitive programming platforms such as Codeforces, AtCoder, AIZU, and CodeChef. To encourage prompt diversity and reduce redundancy in our SFT training set, we apply strict deduplication using two methods: (1) sample I/O fingerprinting and (2) n-gram-based text analysis. This process removes approximately 24.2% of self-duplicated coding prompts.

We choose GPT-OSS-120B (Agarwal et al., 2025) as our SFT teacher model due to its strong code reasoning capabilities. For each coding prompt with verifiable test cases, we apply correctness filtering to the teacher’s reasoning traces, retaining only those that generate correct code. For prompts without verifiable test cases, we generally select longer reasoning traces under the assumption that they reflect more thorough problem analysis. This pipeline yields a final dataset comprising 1.9M Python reasoning traces, 1.0M C++ reasoning traces, and 1.3M Python tool-calling reasoning traces for competitive coding.

**Scientific Coding:** We further collect scientific research coding prompts spanning the domains of biology, material science, physics, chemistry, and mathematics. The responses to these prompts are generated by GPT-OSS-120B (Agarwal et al., 2025), resulting in a total of 1.1M SFT samples.

### 3.2.3. Science

The science prompts we collect span physics, chemistry, and biology. We use 1.4M science SFT samples from Nemotron-Cascade (Wang et al., 2025) and an additional 1.3M samples from Nemotron-3-Nano (Blakeman et al., 2025). Responses in both datasets are generated by GPT-OSS-120B (Agarwal et al., 2025).

### 3.2.4. Long Context

We adopt the 160K long context SFT data from Nemotron-3-Nano (Blakeman et al., 2025), which has an average sequence length of 128K tokens. In addition, we collect another 74K long context SFT from ChatQA-2 (Xu et al., 2024), which has an average length of 29K tokens.

### 3.2.5. General Chat

We source prompts from Nemotron-Cascade 1 (Wang et al., 2025) and construct 4.9M reasoning-on and 372K reasoning-off samples. Responses for reasoning-on samples are generated by GPT-OSS-120B (Agarwal et al., 2025). For reasoning-off samples, 300K responses are drawn from high-quality annotated short answers within the dataset itself, while an additional 330K are generated by DeepSeek-V3-0324 (Liu et al., 2024) to improve response quality.

To enhance multi-turn dialogue capabilities, we synthesize approximately 700K multi-turn conversation samples using two GPT-OSS-120B (Agarwal et al., 2025) instances in a role-playing setup, where one instance plays the user and the other the assistant. The user-side model may terminate the conversation at any point to prevent repetitive exchanges.

We additionally incorporate 4.6M reasoning-on chat samples from Nemotron-3-Nano (Blakeman et al., 2025), with prompts drawn from LMSYS (Zheng et al., 2023) and WildChat (Zhao et al., 2024). Responses are generated by GPT-OSS-120B (Agarwal et al., 2025), Qwen3-235B-A22B-Thinking-2507, and Qwen3-235B-A22B-Instruct-2507 (Yang et al., 2025).

### 3.2.6. Instruction Following

We source prompts from Nemotron-Cascade 1 (Wang et al., 2025) and generate approximately 230K reasoning-on responses using GPT-OSS-120B (Agarwal et al., 2025) and 64K reasoning-off responses using DeepSeek-V3-0324 (Liu et al., 2024). In addition, we incorporate 497K instruction-following samples from Nemotron-3-Nano (Blakeman et al., 2025), including 457K reasoning-on and 40K reasoning-off responses. These responses are generated by GPT-OSS-120B (Agarwal et al., 2025), Qwen3-235B-A22B-Thinking-2507, and Qwen3-235B-A22B-Instruct-2507 (Yang et al., 2025).

### 3.2.7. Safety

We collect 4K safety SFT samples from Nemotron-3-Nano (Blakeman et al., 2025) to enable models to exhibit appropriate refusal behavior when encountering unsafe inputs. The SFT prompts are originally sourced from Nemotron Content Safety v2 (Ghosh et al., 2025), Gretel Safety Alignment v1 (gre, 2024), Harmful Tasks (Hasan et al., 2024), and Red-Team-2K (Luo et al., 2024).

### 3.2.8. Conversational Agent

Aside from the Python tool-use data for math and code reasoning, we further gather tool-use samples in multi-turn conversational settings, where multiple tools are available and the assistant must determine which tools to invoke and how to use them effectively. We collect 822K conversational tool-use samples from Nemotron-3-Nano (Blakeman et al., 2025), with responses generated by Qwen3-235B-A22B-Thinking-2507, Qwen3-32B, Qwen3-235B-A22B-Instruct-2507 (Yang et al., 2025), and GPT-OSS-120B (Agarwal et al., 2025).

### 3.2.9. Software Engineering Agent

We curate the software engineering (SWE) data using various agentic scaffolds, including OpenHands (Wang et al., 2025), SWE-Agent (Yang et al., 2024), Mini-SWE-Agent, and the agentless scaffold proposed by Wei et al. (2025), to enhance the models’ agentic software engineering capabilities. First, we utilize the data from Nemotron 3 Nano (Blakeman et al., 2025) and Super (Blakeman et al., 2025), which includes SWE agentic trajectories generated using Qwen3-Coder-480B-A35B-Instruct (Yang et al., 2025). The problem instances are drawn from SWE-Gym (Pan\* et al., 2025), SWE-rebench (Badertdinov et al., 2025), and R2E-Subset (Jain et al., 2025). Second, we employ SWE agentless data from Nemotron-Cascade 1 (Wang et al., 2025), which includes three main tasks: (1) buggy code localization, (2) code repair, and (3) test case generation. Following the established procedure in Wang et al. (2025), we reconstruct the code repair data using DeepSeek-V3.2 (Liu et al., 2025).

Our preliminary study shows that incorporating SWE agentless data improves models’ effectiveness on SWE agentic tasks. For example, fine-tuning solely on agentic data achieves Pass@1 of 48.9 and Pass@4 of 62.8, whereas fine-tuning on a combination of agentic and agentless data improves performance to Pass@1 of 49.9 and Pass@4 of 65.2 on SWE-bench Verified using OpenHands. Based on this observation, we combine 125K agentic samples and 389K agentless samples as the supervised fine-tuning (SFT) data for SWE tasks. Our models are trained in non-thinking mode on SWE agentic data and in thinking mode on SWE agentless data.

### 3.2.10. Terminal Agent

To enhance agentic capabilities for terminal use, we adopt the *Terminal-Task-Gen* methodology (Pi et al., 2026) to curate our training tasks. This framework consists of (1) dataset adapters that transform static data into interactive terminal formats, and (2) synthetic tasks generated from both diverse seed prompts and a structured terminal skill taxonomy. Using this framework, we curate 490K samples in total. Specifically, we first adapt 162K math, 32K code, and 32K SWE-specific samples from existing high-quality sources (Wang et al., 2025), which establishes broad foundational coverage. To further improve targeted skill refinement, we synthesize 120K seed-based and 140K skill-based tasks. For trajectory construction, we leverage the tasks curated from above, and employ DeepSeek-V3.2 (Liu et al., 2025) as the core engine to generate step-by-step solution traces via an execution-feedback loop within isolated Docker environments. The Terminus 2 agent framework (Merrill et al., 2026) serves as the underlying scaffolding and tool-use protocol, enabling the model to interact with the terminal and complete complex tasks.

## 4. Cascade RL and Multi-Domain On-Policy Distillation

Following a similar approach to Nemotron-Cascade 1 (Wang et al., 2025), we apply Cascaded Reinforcement Learning (Cascade RL) as our post training pipeline. In particular, we integrated the Multi-Domain On-Policy Distillation (MOPD) along the Cascade RL process.

### 4.1. Training Framework

We illustrate our training process in Figure 2. In this work, we start the Cascade RL process with IF-RL (§4.2) to establish foundational instruction adherence, followed by multi-domain RL (§4.3) to enhance the model’s tool-calling capabilities, STEM reasoning, and response format adherence. We then transition to Multi-domain On-policy Distillation (§4.4) to unify specialized expertise into a single, cohesive policy to mitigate performance

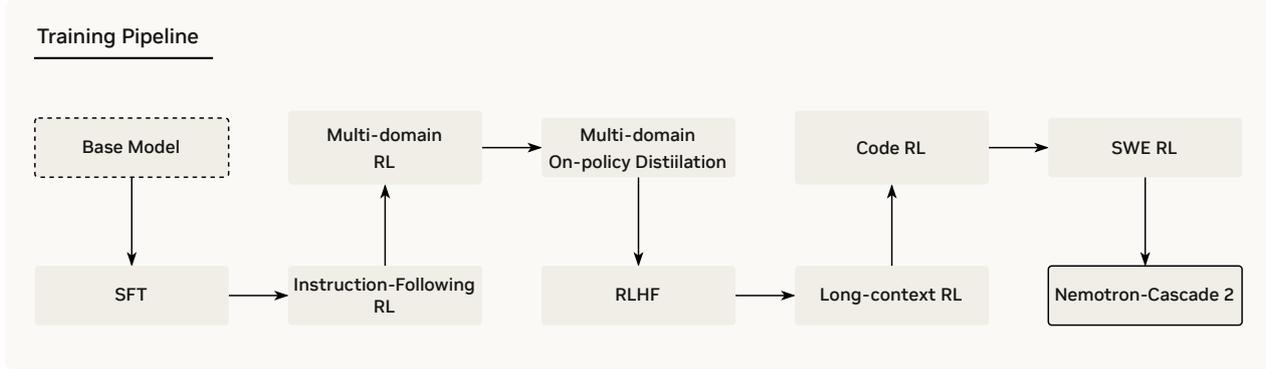


Figure 2: Nemo-RL Cascade 2 applies Cascade RL with the sequential, domain-wise ordering after SFT, leading to substantial improvements across the corresponding domains.

degradation. We continue with RLHF (§4.5) for human alignment, Long-context RL (§4.6) to enhance reasoning over massive input sequences, Code RL (§4.7) for competitive coding problems, and finally SWE RL (§4.8) for mastering agentic software interactions.

#### 4.1.1. What determines the ordering of Cascade RL

The optimal ordering of stages within a Cascade RL pipeline is not a universal constant; rather, it is a dynamic function of the model’s underlying behaviors and learning trajectories. In contrast to the original Nemo-RL Cascade (Wang et al., 2025), our current work Nemo-RL Cascade 2 introduces significant improvements in SFT data quality and substantially scales the complexity of the RL environments and tasks. These advancements have fundamentally altered the model’s behavioral dynamics, which require us to adopt a different order to better accommodate the evolving capabilities of LLMs.

**Rule of thumb: Mitigating Inter-Domain Interference.** Specifically, the rationale for this ordering is primarily driven by the need to mitigate catastrophic forgetting as the model interacts with increasingly diverse environments. Cascade RL provides a granular lens through which we can observe how specific domains compete or conflict, such as strict instruction adherence in IF-RL versus human preference alignment in RLHF. Our core design principle is to identify an ordering that minimizes negative interference across domains while thoroughly optimizing the highest-priority domains. By identifying which tasks serve as foundational priors and which act as specialized refinements, we can mitigate inter-domain interference.

**Scaling via Multi-Domain Integration.** Following this principle, the Cascade RL pipeline can incorporate multi-domain RL stages when specific domains are found to be non-conflicting or beneficial to the overall performance. This integrated approach is particularly effective as RL environments and datasets grow in complexity, while ensuring that the model maintains a broad performance profile across various benchmarks, as detailed in §4.3.

**Stabilization through On-policy Distillation.** Furthermore, We find that Multi-domain On-policy Distillation (§4.4) serves as a critical stabilization point in this ordering. It is effective at recovering benchmark performance that may have regressed during earlier, more specialized stages of the cascade RL, leading to a more balanced and robust final policy model.

#### 4.1.2. RL Training Configuration

Throughout the entire Cascade RL process, we use Group Relative Policy Optimization (GRPO) algorithm (Shao et al., 2024) with strict **on-policy** training following Nemo-RL Cascade (Wang et al., 2025). We adopt on-policy training for improved stability and higher accuracy. We conduct our training using the Nemo-RL repository (NVIDIA, 2025).

At each iteration, we generate a group of  $G$  rollouts from the current policy  $\pi_\theta$  and then perform a *single* gradient update. This ensures that the policy used for data collection always matches the one being updated, making the importance sampling ratio exactly 1. This on-policy setup contributes to stable RL training and mitigates entropy collapse. In addition, we remove KL divergence term entirely, which simplifies the GRPO objective to the standard *REINFORCE* objective (Williams, 1992) with group-normalized rewards and token-level loss (Yu et al., 2025):

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_\theta(\cdot|q)} \left[ \frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \hat{A}_{i,t} \right], \text{ where } \hat{A}_{i,t} = \frac{r_i - \text{mean}(\{r_i\}_{i=1}^G)}{\text{std}(\{r_i\}_{i=1}^G)} \text{ for all } t, \quad (1)$$

and  $\{r_i\}_{i=1}^G$  denotes the group of  $G$  rewards assigned to the sampled responses  $\{o_i\}_{i=1}^G$  for a given question  $q$  drawn from the dataset  $\mathcal{D}$ , verified against the ground-truth answer  $a$  in RLVR. For RLHF,  $r_i$  is the aggregated reward score from the generative reward model for response  $o_i$  and question  $q$ . Details of the reward functions for different domains will be provided in the corresponding subsections.

## 4.2. Instruction-Following Reinforcement Learning (IF-RL)

In this subsection, we describe our instruction-following RL recipe, which serves as the first stage of our Cascade RL. We demonstrate that applying verifiable IF-RL significantly improves instruction adherence, achieving a state-of-the-art accuracy of 83.13% on IFBench (Pyatkin et al., 2025).

### 4.2.1. Dataset

We use the same instruction-following training data used for NVIDIA Nano-v3 post-training (Blakeman et al., 2025). The instructions in this dataset are designed for objective verifiability, for instance, requiring a response to be under 200 words. This making the dataset well-suited for training and evaluating models on strict adherence. Given the high baseline quality of the data, our curation process mainly resolves formatting inconsistencies within the keyword arguments for certain instruction types (e.g., `count_increment_word`).

### 4.2.2. Training recipe

Following (Wang et al., 2025), we also apply dynamic filtering (Yu et al., 2025). This technique filters out samples where all rollouts are either entirely correct or entirely incorrect. By ensuring that every prompt in a batch provides effective gradients, dynamic filtering stabilizes IF-RL training and pushes the upper bound of model performance. Furthermore, we observed that extended IF-RL training can lead to excessive token usage, which is often unnecessary for fulfilling specific constraints in general chat domains. To mitigate this, we apply overlong penalty, which penalizes samples that fail to complete generation within the maximum sequence length with a zero reward.

Unlike Nemotron Cascade (Wang et al., 2025), we position IF-RL as the first stage of our Cascade RL training for two primary reasons: (i) IF-RL can negatively impact human alignment capabilities (e.g., ArenaHard), while our subsequent generative-reward-model-based RLHF has a negligible impact on instruction following scores. By prioritizing instruction adherence first, we can focus on maximizing instruction following performance and then utilize the later stages to recover and refine human preference alignment. (ii) An early IF-RL stage produces a model with superior instruction-following capabilities, which serves as a strong teacher for subsequent multi-domain on-policy distillation. Another difference from Nemotron Cascade (Wang et al., 2025) is that our IF-RL is trained exclusively in “thinking mode” without incorporating a reward model. We found that the “thinking mode” yields higher accuracy on instruction-following benchmarks (e.g., IFBench (Pyatkin et al., 2025)). Because subsequent RL stages recover any regressions in human preference alignment introduced during IF-RL, we can focus entirely on maximizing instruction adherence without incurring the computational overhead of an auxiliary reward model.

We use a batch size of 128, sampling 16 responses per prompt with temperature 1.0 and top-p 1.0. We adopt

a learning rate of  $3e-6$  with AdamW (Kingma, 2014), and set both the entropy loss coefficient and KL loss coefficient to 0. Our IF-RL with dynamic filtering takes around 180 steps. The full set of hyperparameters is provided in Appendix B.

### 4.3. Multi-domain RL

Following IF-RL, we conduct an additional stage of multi-domain RL that covers three capabilities: multi-choice question answering (MCQA) in the STEM domain, agentic tool calling, and structured output for instruction following. The datasets are drawn from the NVIDIA Nano-v3 RL training blend (Blakeman et al., 2025). The data mixture consists of approximately 55% MCQA, 30% agentic tool calling using the Workplace Assistant setup (Blakeman et al., 2025), and 15% structured output.

We group these domains into a single multi-domain RL stage for two main reasons. First, we do not observe performance degradation across evaluation benchmarks when training on the blended domains. Instead, the model exhibits consistent improvements on benchmarks including MMLU-Pro,  $\tau^2$ -Bench, and IF-Bench. Second, the response lengths and verification times of these datasets are similar, which minimizes training inefficiencies caused by waiting for longer generations or slower environment verification.

During training, we use a batch size of 128 and sample 16 responses per prompt with temperature 1.0 and top-p 1.0 (see Appendix B). We adopt a learning rate of  $3 \times 10^{-6}$  with AdamW (Kingma, 2014), and set both the entropy loss coefficient and KL loss coefficient to zero. This multi-domain RL stage runs for approximately 70 training steps.

### 4.4. Multi-domain On-Policy Distillation (MOPD)

While well-designed Cascade RL substantially reduces catastrophic forgetting compared with vanilla sequential RL in an arbitrary order, it does not fully eliminate capability drift as the number of training environments increases. In practice, we observe noticeable fluctuations across different benchmark categories tracked throughout training, and the dominant trade-offs differ by stage. For example, certain RLVR training often reduces model entropy and shortens reasoning traces, thus can negatively impact mathematical reasoning performance, while RLHF-oriented optimization can partially trade off against instruction-following behavior. These observations motivate an additional training stage for re-balancing capabilities within the Cascade RL process.

We therefore adopt **multi-domain on-policy distillation (MOPD)** (Agarwal et al., 2024; Lu and Lab, 2025; Xiao et al., 2026; Yang et al., 2025; Zeng et al., 2026) as a complementary post-training stage. In our setting, MOPD is particularly attractive for three reasons. First, teacher checkpoints can be selected directly from the Cascade RL pipeline by choosing the strongest validation checkpoint for each benchmark category, which makes it easy to assemble a capability-diverse teacher pool without introducing external model families. Second, because these teachers are derived from the same SFT initialization, they share the same tokenizer and vocabulary as the student, reducing distribution shift and avoiding additional alignment issues. Third, MOPD provides a dense token-level training advantage, which is especially useful compared with sparse outcome rewards, and in Figure 3(c) we show its training-efficiency benefits compared with GRPO.

#### MOPD objective.

Let  $\pi^{inf}$  denote the student policy used for response generation in the inference engine, and let  $\pi^{train}$  denote the student policy optimized by the training engine. For each prompt  $x$ , we sample a response  $y = (y_1, \dots, y_T) \sim \pi^{inf}(\cdot | x)$ . We then select a domain teacher  $\pi^{domain_i}$  for that training example, where  $domain_i$  indicates the capability domain associated with the chosen teacher. Writing  $s_t = (x, y_{<t})$  for the decoding state at step  $t$ , we define the token-level distillation advantage using reverse-KL as

$$a_t^{MOPD} = \log \pi^{domain_i}(y_t | s_t) - \log \pi^{train}(y_t | s_t). \quad (2)$$

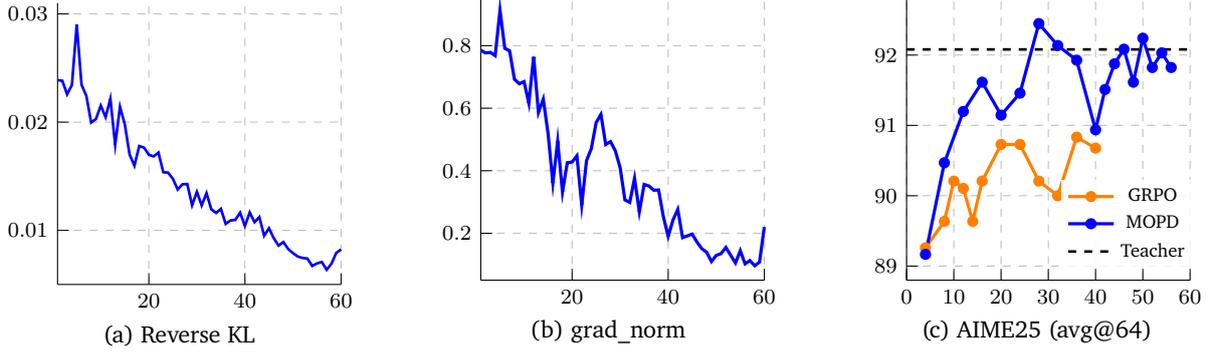


Figure 3: Training dynamics and downstream evaluation.

Intuitively, this term is positive when the domain teacher assigns a higher probability to the sampled token than the current training policy, and therefore serves as a dense token-level distillation advantage that converges toward 0 during training. The log-probability difference is computed only on the student-sampled token rather than over the full vocabulary.

Because responses are sampled under  $\pi^{inf}$  but optimized under  $\pi^{train}$ , we apply truncated importance weighting to account for train–inference mismatch:

$$r_t = \frac{\pi^{train}(y_t | s_t)}{\pi^{inf}(y_t | s_t)}, \quad w_t = \text{sg}[r_t] \mathbf{1}[\epsilon_{low} \leq r_t \leq \epsilon_{high}], \quad (3)$$

where  $\text{sg}[\cdot]$  denotes stop-gradient. We then optimize the surrogate objective

$$\mathcal{L}_{MOPD} = -\mathbb{E}_{x \sim \mathcal{D}, y \sim \pi^{inf}(\cdot | x)} \left[ \frac{1}{|\mathcal{V}(y)|} \sum_{t \in \mathcal{V}(y)} w_t \text{sg}[a_t^{MOPD}] \log \pi^{train}(y_t | s_t) \right], \quad (4)$$

where  $\mathcal{V}(y)$  is the set of valid response tokens retained by the token mask.

### Hyperparameters.

Unless otherwise specified, we use a rollout size of 4 and 128 prompts per update, giving an effective batch size of 512 responses. In later experiments, we find that using 512 prompts with rollout size 1 yields slightly more stable optimization while producing similar final results. We use a learning rate of  $2 \times 10^{-6}$  with linear warm-up over the first 30 optimization steps, starting from  $2 \times 10^{-7}$ . Training typically converges within 40-50 optimization steps (Fig. 3(a)). We find the warm-up stage important for stability: gradient norms are substantially larger at the beginning of training and decrease rapidly after the warm-up phase (Fig. 3(b)). For truncated importance weighting, we set  $\epsilon_{low} = 0.5$  and  $\epsilon_{high} = 2.0$ . In the main experiments, we use three domain teachers corresponding to math, RLHF, multi-domain. The math teacher is the initial SFT checkpoint, which already exhibits strong mathematical reasoning capabilities thanks to the meticulously curated SFT dataset. The RLHF teacher is a checkpoint optimized through RLHF from the initial SFT checkpoint. The multi-domain teacher is selected from the checkpoints after previous IF-RL + Multi-domain RL stages. We sample prompts accordingly from the RL training data pools (RLHF, IF-RL, and Multi-domain), as well as from AceReason-Math for math (Chen et al., 2025).

### Training efficiency advantage.

MOPD provides a dense token-level distillation advantage, whereas GRPO relies on a sparse sequence-level outcome reward that is shared across all generated tokens. This makes MOPD substantially more sample- and step-efficient in practice. Starting from the same initial checkpoint, MOPD consistently reaches stronger performance in fewer optimization steps. On AIME25 (Figure 3(c)), under math-only training, GRPO improves from

Table 3: Comparison of MOPD and RLHF at matched evaluation checkpoints on ArenaHard V2.0.

| Method  | Steps | ArenaHard v2 |                  |
|---------|-------|--------------|------------------|
|         |       | Hard Prompt  | Creative Writing |
| Initial | 0     | 71.5         | 40.6             |
| RLHF    | 100   | 81.7         | 68.6             |
|         | 160   | 80.7         | 71.2             |
| MOPD    | 52    | 85.5         | 71.0             |

89.9 to 91.0 after 25 steps, while MOPD reaches 92.0 within 30 steps and recovers teacher-level performance. A similar trend appears on ArenaHard v2 (Table 3). After 52 steps, MOPD improves Hard Prompt from 71.5 to 85.5 and Creative Writing from 40.6 to 71.0. In contrast, RLHF training requires 160 steps to reach 80.7 on Hard Prompt and 71.2 on Creative Writing. These results show that the dense token-level advantage in on-policy distillation lead to much faster training convergence.

## 4.5. Reinforcement Learning from Human Feedback (RLHF)

Building on multi-domain on-policy distillation, our RLHF recipe focuses on human preference learning. This process further enhances creative writing and non-verifiable problem-solving in coding and mathematics, as measured by ArenaHard v2 (Li et al., 2024), while maintaining performance across other domains without degradation.

### 4.5.1. Dataset

We adopt the RLHF training dataset from NVIDIA Nano-v3 (Blakeman et al., 2025), which comprises HelpSteer3 (Wang et al., 2025), a commercially-friendly subset of the arena-human-preference-140k dataset (Chiang et al., 2024), and a synthetic safety blend (Blakeman et al., 2025). Following the NVIDIA Nano-v3 (Blakeman et al., 2025), we utilize Qwen3-235B-A22B-Thinking-2507 (Yang et al., 2025) as our generative reward model (GenRM), trained via the HelpSteer3 framework (Wang et al., 2025). Given a conversation history, a user request, and two candidate responses, the GenRM first reasons through the strengths and weaknesses of each response before producing individual helpfulness scores and a final comparative ranking.

### 4.5.2. Training recipe

Following a training recipe similar to NVIDIA Nano-v3 (Blakeman et al., 2025), we conduct RLHF using the GenRM. To ensure the training signals are of high quality, we adopt pair-wise comparisons for all pairs of rollouts per prompt. We aggregate the reward scores in the same way as NVIDIA Nano-v3 RLHF training, and apply the same length-normalized reward adjustment and quality-gated conciseness bonus (Blakeman et al., 2025). These mechanisms encourage shorter responses without sacrificing quality, effectively mitigating the rapid growth of inference token usage.

Different from Nemotron Cascade (Wang et al., 2025), we train RLHF exclusively in the thinking mode. While incorporating both thinking and non-thinking modes can improve training convergence and yield slight gains on evaluation benchmarks, we observe a significant degradation in instruction-following performance. The resulting drop is substantial enough that the gains obtained in the earlier RLVR stage cannot be fully recovered.

### 4.5.3. Hyper-parameters

We use a batch size of 128, generating 16 rollout per prompt with a temperature of 1.0 and a top-p value of 1.0. We use a maximum response length of 16K during RLHF without applying overlong filtering. We adopt a learning rate of  $3e-6$  with AdamW (Kingma, 2014). We set the entropy loss coefficient to 0 and the KL loss coefficient to 0.03 to keep the model capabilities on other domains. The training takes around 30 steps.

## 4.6. Long-context RL

Following RLHF, we conduct a stage of long-context RL to further enhance the model’s long-context understanding and reasoning capabilities. We use the NVIDIA Nano-v3 RL data blend (Blakeman et al., 2025), but restrict this phase to long-context datasets only. In our experiments, incorporating other domains during long-context RL negatively affects performance on unrelated benchmarks, motivating this domain-specific training setup.

We adopt the Nemo-Gym RL environment (NVIDIA, 2025) and use Qwen3-235B-A22B-Instruct-2507 as an LLM judge to evaluate model rollouts for question answering tasks. During training, input sequences are limited to 32K tokens, and the maximum sequence length is set to 49K tokens without applying overlength filtering.

We train with a batch size of 128, generating 16 rollouts per prompt with temperature 1.0 and top-p 1.0. Optimization is performed using AdamW (Kingma, 2014) with a learning rate of  $3 \times 10^{-6}$ , while both the entropy and KL loss coefficients are set to zero. Training runs for approximately 30 steps, as we observe a rapid increase in generated tokens beyond that point.

## 4.7. Code RL

### 4.7.1. Data Curation

We construct our Code RL training set from the Nemotron-Cascade coding corpus (Wang et al., 2025), which contains coding prompts sourced from modern competitive programming platforms such as AtCoder, Codeforces, and AIZU with robust test cases for reward verification. To improve training efficiency and strengthen deep reasoning, we aggressively filter out prompts that GPT-OSS-120B solves correctly in all 8 of 8 rollouts, yielding a compact final set of only 3.5K samples. We find that high-difficulty prompts paired with strong test cases are critical for further boosting model performance.

### 4.7.2. Training Details

We conduct Code RL using a batch size of 128 and a learning rate of  $3 \times 10^{-6}$  with the AdamW optimizer. Compared to Nemotron-Cascade, we increase the maximum response length during RL to 118K tokens and the number of rollouts per sample to 16, enabling the policy to better capture sparse reward signals on extremely difficult problems that require long reasoning traces. We adopt the strict binary reward function to avoid potential reward hacking and keep the whole training to be fully on-policy for stability. To support the resulting verification throughput of  $128 \times 16 = 2,048$  code executions per RL step, we deploy an asynchronous reward verification server that completes each batch in 427.2 seconds across 384 CPU cores.

## 4.8. Software Engineering Reinforcement Learning (SWE RL)

### 4.8.1. Agentless RL

#### Training Details and Hyperparameters.

To enhance the models’ code repair capability, we adopt the same data source as Wang et al. (2025) for agentless code repair reinforcement learning (RL) training. Since most instances do not provide executable Docker environments, we employ GPT-OSS-120B as a reward model to evaluate the quality of code repairs generated by our models. Following Wang et al. (2025), for each instance we construct prompts using both the golden localization and the top-5 retrieved localizations, and filter out relatively easy samples. We perform agentless SWE RL with a batch size of  $128 \times 16 = 2,048$  (128 prompts with 16 rollouts per prompt), a maximum sequence length of 98,304, and a learning rate of  $3 \times 10^{-6}$  using the AdamW optimizer. We sample responses with temperature 1.0 and top-p 1.0. During training, we mask the loss for prompts for which none of the rollouts receives a reward greater than 0.5. We observe that these difficult prompts degrade the stability and effectiveness of agentless SWE RL training. Our agentless RL training typically converges within 40–50 steps.

Table 4: Effectiveness of Agentless RL on SWE-bench Verified.

| Scaffold           | Agentless Mini |        | OpenHands |        |
|--------------------|----------------|--------|-----------|--------|
|                    | avg@4          | pass@4 | avg@4     | pass@4 |
| Init.              | 41.9%          | 55.2%  | 49.8%     | 64.2%  |
| after Agentless RL | 44.3%          | 57.4%  | 50.8%     | 65.0%  |

### Can Agentless RL Training Helps Agentic Tasks?

Table 4 shows that agentless RL training not only improves model performance within the agentless framework but also enhances the models’ ability to solve SWE tasks in agentic settings. Note that for Agentless Mini evaluation, we employ a code embedding model, NV-Embed-Code (Sohrabizadeh et al., 2025), to retrieve 5 candidate files whose code contents are semantically similar to the problem context. This result suggests that improving models’ code repair capability alone can generalize across different scaffolds, consistent with the observations from Yang et al. (2026).

#### 4.8.2. Execution-based RL for Agentic SWE Scaffold

Modern software engineering agents rely on scaffolding frameworks that coordinate repository interaction, tool calling, code editing, and test execution. Training agents to operate effectively within these environments requires optimizing not only individual model outputs but the entire problem-solving trajectory. To address this, we apply Reinforcement Learning from Verifiable Rewards (RLVR) directly within agentic SWE scaffolds, enabling end-to-end optimization of the full agent workflow. Our training environments integrate established OpenHands frameworks (Wang et al., 2025), which provide structured tool usage, repository interaction, and iterative patch generation.

We train agents using execution-based reinforcement learning in fully executable software environments, where each episode corresponds to resolving a software issue instance from benchmarks such as SWE-bench. The agent operates inside an instrumented repository that exposes tools for file inspection, search, code editing, and test execution. Candidate patches generated by the agent are executed within the environment, which returns verifiable signals from compilation results and unit test outcomes, enabling automatic reward computation without human annotation. Through the OpenHands scaffolding framework, the agent iteratively localizes defects, proposes patches, and validates them through test execution. Environment feedback—including compilation errors, failing tests, or successful test passes—provides deterministic rewards that directly reflect functional correctness.

Specifically, we conduct execution-based agentic reinforcement learning with a batch size of 1024, corresponding to 16 prompts with 64 rollouts per prompt. The maximum context length is set to 256k tokens, and the agent is allowed up to 200 interaction turns, providing a larger reasoning token budget during agentic coding problem solving. Training data is drawn from SWE-Gym (Pan\* et al., 2025) and R2E-Subset (Jain et al., 2025). We generate 16 rollouts per instance using our intermediate model and evaluate them using the verification pipeline. Instances for which all rollouts pass verification (100% accuracy), indicating overly simple problems, are removed from the dataset. For instances where none of the rollouts pass verification (0% accuracy), indicating extremely difficult problems, we randomly discard 90% of such cases to reduce their proportion in the training data.

## 5. International Mathematical Olympiad (IMO)

### 5.1. IMO 2025

In Table 2, we evaluate Nemotron-Cascade-2-30B-A3B on the IMO 2025 problem set using a self-improving test-time scaling framework (Shao et al., 2025), in which the model iteratively generates candidate solutions,

verifies them, and refines them based on its own feedback. Remarkably, despite its relatively modest 30B-A3B scale, the model successfully solves the first five problems. We provide the full model solutions in Appendix E, together with comments from the human expert. These results are particularly encouraging, as they suggest that strong olympiad-level mathematical reasoning can emerge from a comparatively compact model when paired with effective inference-time scaling. There remain several promising directions for improvement: expert review indicates that some proofs are longer than necessary, include superfluous intermediate steps or definitions, occasionally expose traces of intermediate reasoning, and sometimes contain minor typographical issues. For Problem 2, the model adopts an analytic solution strategy, similar to OpenAI’s approach, rather than a more geometric approach such as that used by Gemini Deep Think (IMO Gold).

## 5.2. IMO-ProofBench

Table 5: IMO-ProofBench (Luong et al., 2025) reports scores split into the Basic (30 problems) and Advanced (30 problems) subtasks, as well as Overall (60 problems). Expert-evaluated results are taken from the IMO-ProofBench leaderboard (accessed on 2026/3/9).

| Model   | IMO-ProofBench |               |              |
|---|----------------|---------------|--------------|
|   | Basic (30)     | Advanced (30) | Overall (60) |
| Aletheia (Feng et al., 2026)                        | -              | 91.9          | -            |
| Gemini 3 Deep Think (Gemini Team, 2026)             | -              | 76.7          | -            |
| Gemini Deep Think (IMO Gold) (Gemini Team, 2025)    | 89.0           | 65.7          | 76.7         |
| DeepSeek-Math-V2-671B-A37B (Shao et al., 2025)      | 99.0           | 61.9          | 80.2         |
| DeepSeek-Math-V2-671B-A37B (our reproduced score) † | 99.5           | 57.7          | 78.6         |
| <b>Nemotron-Cascade-2-30B-A3B</b> †                 | 92.5           | 53.4          | 72.9         |
| GPT-5.2-Thinking (high) (OpenAI, 2025)              | -              | 35.7          | -            |
| Gemini 3 Pro (Gemini Team, 2025)                    | -              | 30.0          | -            |
| GPT-5 Pro (OpenAI, 2025)                            | -              | 28.6          | -            |

† Use DeepSeek-V3.2-Spiale as the judge model with LLM ProofAutoGrader prompt (Luong et al., 2025).

As shown in Table 5, Nemotron-Cascade-2-30B-A3B achieves 72.9 on IMO-ProofBench with generate-verify-refine test-time scaling, placing it within 8 points of DeepSeek-Math-V2-671B-A37B despite using  $10\times$  fewer active parameters. It reaches 90+ on Basic split and surpass the QED-Nano-4B (54.0) (LM-Provers et al., 2026) by 18 points, though the latter is not directly comparable due to judge model. Re-evaluating the provided DeepSeek-Math-V2 proofs under our LLM-judge setup yields a score within 4 points of the reported human rating, suggesting that our protocol does not substantially overestimate performance (more details in Appendix A.1.2). In Figure 4, we show that increasing test-time compute improves Nemotron-Cascade-2-30B-A3B on IMO-ProofBench (Advanced), raising the score from 40.7 at round 1 to 53.4 at round 5 and narrowing the gap to DeepSeek-Math-V2 under the same grader.

## 6. Competitive Coding

### 6.1. IOI 2025 and ICPC World Finals 2025

For IOI 2025, we adapt the IOI Test-Time Scaling pipeline from Nemotron-Cascade (Wang et al., 2025), which can be viewed as a multi-round *generate-select-submit* framework that exploits the model’s reasoning ability under IOI’s official rules. Each subtask is allotted at most 50 rounds. Within each round, we prompt our model to generate 40 candidate solutions, aggregated with (1) submission history with official judge verdicts from previous rounds, and (2) shared insights from high scored or fully solved subtasks within the same main task. The complete chat template is provided in Appendix C.1. Using this approach, we achieved full score on Problem 3 and 4, achieving a gold-medal score of 439.28 within at most  $40 \times 50 = 2000$  model generations, while the score of 507.66 is achievable within 5000 generations. Notably, on Problem 2 which requires designing and optimizing a heuristic algorithm, our pipeline reached over 86 points in just 5 rounds (at most 200 model generations), demonstrating the effectiveness of self-refinement and cross-subtask insights.

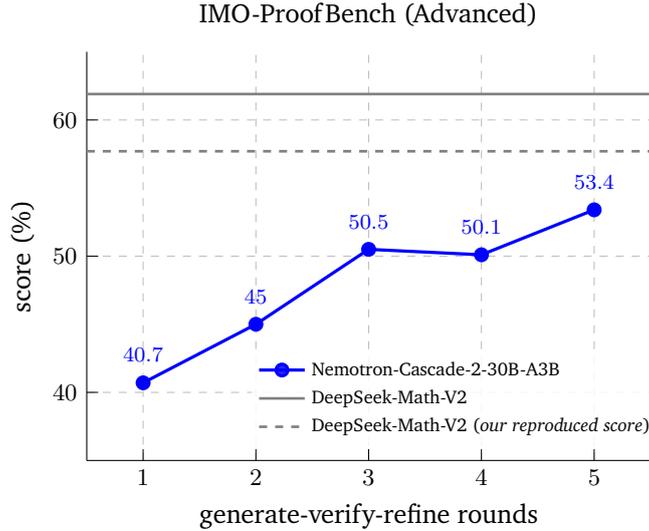


Figure 4: IMO-ProofBench (Advanced) score graded by LLM ProofAutoGrader (DeepSeek-V3.2-Speciale).

Table 6: Competitive programming results on comprehensive benchmarks, evaluated against a significantly expanded set of proprietary and open-source baseline models.

| Models                           | LiveCodeBench v6<br>2408 - 2505 | LiveCodeBench Pro |             |            |             |             |            | Codeforces<br>2501 - 2507 |             |
|----------------------------------|---------------------------------|-------------------|-------------|------------|-------------|-------------|------------|---------------------------|-------------|
|                                  |                                 | 25Q1              |             |            | 25Q2        |             |            | ELO                       | Percentile  |
|                                  |                                 | Easy              | Med         | Hard       | Easy        | Med         | Hard       |                           |             |
| GPT-5.2 (high)                   | -                               | 96.6              | 75.0        | 5.9        | 91.8        | 59.6        | 23.1       | 2590                      | 99.9        |
| Gemini-3 Pro                     | 90.7                            | 94.4              | 70.0        | 5.9        | 94.8        | 45.6        | 7.7        | 2440                      | 99.8        |
| GPT-o4-mini (high)               | 80.2                            | 85.4              | 51.7        | 0.0        | 84.5        | 29.8        | 0.0        | 2266                      | 99.5        |
| DeepSeek-v3.2-Speciale           | 88.7                            | 89.7              | 48.1        | 0.0        | 88.5        | 43.1        | 0.0        | 2353                      | 99.7        |
| GPT-OSS-120B (high)              | 87.0                            | 88.8              | 41.9        | 0.7        | 88.5        | 31.1        | 0.0        | 2320                      | 99.6        |
| Kimi-K2.5-1T-thinking            | 85.0                            | 88.5              | 45.6        | 0.0        | 90.2        | 37.9        | 0.0        | 2333                      | 99.7        |
| Qwen-3.5-397B-A17B               | 83.6                            | 89.3              | 44.4        | 0.0        | 88.1        | 31.4        | 0.0        | 2350                      | 99.7        |
| Qwen-3.5-122B-A10B               | 78.9                            | 87.6              | 35.6        | 0.0        | 84.3        | 24.2        | 0.0        | 2233                      | 99.4        |
| Qwen-3.5-35B-A3B                 | 74.6                            | 84.6              | 25.6        | 0.0        | 81.1        | 17.8        | 0.0        | 2181                      | 99.1        |
| Nemotron-3-Super-120B-A12B       | 78.7                            | 83.0              | 31.0        | 0.0        | 81.7        | 23.2        | 0.0        | 2212                      | 99.4        |
| Qwen3-235B-A22B-Thinking-2507    | 78.7                            | 75.8              | 18.8        | 0.0        | 77.6        | 17.5        | 0.0        | 2119                      | 98.6        |
| Nemotron-Cascade-14B             | 74.6                            | 71.6              | 16.3        | 0.0        | 68.9        | 10.5        | 0.0        | 2004                      | 97.9        |
| Qwen3-Next-80B-A3B-Thinking      | 73.2                            | 68.5              | 16.3        | 0.0        | 69.1        | 7.5         | 0.0        | 1894                      | 96.8        |
| Nemotron-3-Nano-30B-A3B          | 68.3                            | 60.3              | 6.0         | 0.0        | 54.5        | 3.5         | 0.0        | 1681                      | 93.1        |
| Nemotron-Cascade-2-30B-A3B       | <b>87.2</b>                     | <b>88.1</b>       | <b>39.2</b> | <b>0.7</b> | <b>87.0</b> | <b>27.6</b> | <b>0.0</b> | <b>2320</b>               | <b>99.6</b> |
| Nemotron-Cascade-2-30B-A3B (TIR) | <b>88.4</b>                     | <b>91.0</b>       | <b>45.2</b> | <b>2.2</b> | <b>89.3</b> | <b>36.8</b> | <b>0.0</b> | <b>2345</b>               | <b>99.7</b> |

For ICPC World Finals 2025, we generate up to 1000 solutions per problem and submit them for official evaluation after initial filtering. We successfully solved 10 out of 12 problems, achieving the #4 Gold medal placement, with 8 problems (except Problems A and I) solved within only 100 submissions.

## 6.2. Competitive Coding Benchmark Results

We evaluate our Nemotron-Cascade-2-30B-A3B model on various competitive coding benchmarks, including LiveCodeBench v6 (Jain et al., 2024), and LiveCodeBench Pro (Zheng et al., 2025)’s 25Q1 and 25Q2 splits. We also estimate Codeforces ELO score through simulated participation on 40 Div.1/Div.2 Codeforces Rounds held from 2501 to 2507. We report our avg@8 results under 128K-token thinking budget, the sampling temperature of 1.0 and the  $\text{top}_p$  of 0.95. For Tool-Integrated Reasoning (TIR) results, we allow our model to call a stateful Python executor for up to 100 calls. For baseline model evaluation, we follow their recommended inference configurations, ensuring a thinking budget of at least 128K tokens to at most 256K tokens. More evaluation

details can be found in Appendix A and Appendix D.

As shown in Table 6, Nemotron-Cascade-2-30B-A3B achieves magnificent Pass@1 accuracy and ELO rating, even compared with frontier open-source models with over 100B total params, such as Nemotron-3-Super-120B-A12B, GPT-OSS-120B, and Qwen-3.5-122B-A10B. With Tool-Integrated Reasoning (TIR), our model’s performance can be further boosted especially on hard problems, and match the strongest open-source models with more than 300B total parameters, such as Kimi-K2.5-1T-Thinking, Qwen-3.5-397B-A17B, and DeepSeek-v3.2-Speciale, which either lack TIR support for deep reasoning or perform poorly with Python TIR. Notably, Nemotron-Cascade-2-30B-A3B achieves above 0% on the LiveCodeBench Pro hard split within 8 attempts, demonstrating strong reasoning ability on problems that are extremely difficult even for humans.

## 7. Acknowledgments

We would like to extend our gratitude to the NVIDIA Nemo team for the valuable discussion and collaboration on building reasoning models. We especially wish to thank Boris Ginsburg, Oleksii Kuchaiev, Igor Gitman, Olivier Delalleau, Zhilin Wang, Olivier Delalleau, Tugrul Konuk, Wei Du, Somshubra Majumdar, Wasi Uddin Ahmad, Siddhartha Jain, Jiaqi Zeng, Yi Dong, Alexander Bukharin, Vahid Noroozi, Khushi Bhardwaj, Sugam Dipak Devare, Jian Zhang, and Jonathan Cohen.

We thank Ying Lin for helpful discussions and useful input in building the knowledge-intensive SFT dataset. We also thank Atefeh Sohrabizadeh, Jialin Song, and Jonathan Raiman for valuable discussions on SWE-bench.

# Appendix

## A. Benchmarks and Evaluation Setups

### A.1. Math

#### A.1.1. Non-proof Math

For non-proof math reasoning tasks, we include

- **AIME 2025** (MAA, 2025) consists of 30 problems from American Invitational Mathematics Examination at 2025.
- **AIME 2026** (MAA, 2026) consists of 30 problems from American Invitational Mathematics Examination at 2026.
- **HMMT Feb 2025** (HMMT, 2025) consists of 30 problems from Harvard-MIT Mathematics Tournament 2025 February math competition.
- **IMO-AnswerBench** (Luong et al., 2025) consists of 400 problems with verifiable answers carefully chosen from past Olympiad competitions and then altered by experts to avoid memorization.

For Nemotron-Cascade-2-30B-A3 evaluated on AIME 2025, AIME 2026 and HMMT 2025 Feb, we set the thinking budget (maximum response length) to 131K tokens, the sampling temperature to 1.0, the top-p value to 1.0. For the with-tool setting, we enable tool use by appending a system-prompt postfix, allowing the model to call a stateful Python executor for up to 100 tool calls with a maximum response length of 131K tokens. For IMO-AnswerBench, we set to 256K tokens because we found the questions are significantly more difficult. We use and report the LLM-Judge score using GPT-OSS-120B (Agarwal et al., 2025) as the judge and the AnswerAutoGrader prompt (Luong et al., 2025) for answer correctness on IMO-AnswerBench as the short answers are complicated for rule-based verifier to compute. Following Liu et al. (2024, 2026), we report avg@64 for AIME/HMMT and avg@16 for IMO-AnswerBench.

For baseline models, we use official numbers from their reports or evaluate them with the recommended settings if the official numbers are unavailable.

#### A.1.2. Math Proof

For math proof tasks, we include

- **IMO 2025** (IMO, 2025) consists of 6 problems from IMO 2025.
- **IMO-ProofBench** (Luong et al., 2025) is designed to evaluate the ability of AI models to construct comprehensive and valid mathematical arguments. This benchmark consists of 60 proof-based problems, curated to mirror the kinds of problems found in the IMO.

For Nemotron-Cascade-2-30B-A3, we apply test-time scaling following the DeepSeek-Math-V2 generate-verify-refine pipeline, using the same instructions. We implement this pipeline with NeMo-Skills (NVIDIA, 2025). We use the default hyperparameters from DeepSeek-Math-V2: 128 proof generations, 64 verifications per proof, selection of the top 32 proofs for refinement, and 8 verification analyses paired with each proof, prioritizing the lowest-rated analyses. We then generate 4 refined proofs and continue for up to 8 rounds, or until the average proof score reaches the threshold of 0.99999. We set the maximum generation length to 256K tokens, with temperature 1.0 and top-p 0.95.

For IMO-ProofBench Basic and 11 problems from the Advanced split (i.e., Problems 1, 4, 7, 13, 14, 17, 19, 22, 25, 26, and 28), we reduce the compute budget to 32 proof generations, 16 verifications, top 8 proofs, and 2 rounds to save compute. For IMO-ProofBench evaluation, we use DeepSeek-V3.2-Special to make sure the results are reproducible later and run 64 grading attempts with the ProofAutoGrader prompt (Luong et al., 2025). We found that reporting mean score yields 73.8 for DeepSeek-Math-V2 on the Advanced

split, which is substantially more generous than the human rating of 61.9. We therefore adopt a simple aggregation rule based on analysis: if any judge assigns a score of 0, the final score is set to 0; otherwise, return the mean score. Under this rule, DeepSeek-Math-V2 obtains 57.7, which is much closer to the human rating and reduces the discrepancy from 11.9 points to 4.2 points.

## A.2. Code Reasoning

For code generation tasks, we include

- **LiveCodeBench** (Jain et al., 2024) contains diverse algorithm coding problems with unit tests, collected from AtCoder, LeetCode platforms. We evaluate models competitive coding capability on LiveCodeBench v6 (2024/08-2025/05, 454 problems in total). We report pass@1 accuracy in *thinking* mode, averaged over 8 generations (avg@8).
- **LiveCodeBench Pro** (Zheng et al., 2025) contains daily-updated challenging competitive coding problems with strong unit tests, collected mainly from top-tier coding contests. We report pass@1 accuracy on Easy/Med difficulty splits in *thinking* mode, averaged over 8 generations (avg@8) on two recently released subsets: 2025Q1 (2025/01-2025/04, 166 problems in total) and 2025Q2 (2025/04-2025/07, 167 problems in total).
- **IOI and ICPC World Finals** represent the most challenging and prestigious annual algorithmic coding competitions, gathering the world’s top human contestants. The IOI awards gold medals to approximately the top 8.3% (one-twelfth) of participants, while the ICPC World Finals (ICPCWF) limits gold medals to only the top 4 teams globally.
- **SciCode** (Tian et al., 2024) serves as a challenging benchmark to evaluate model’s ability on solving realistic scientific research tasks from STEM domains. It contains 338 subproblems from 80 main tasks.

For Nemotron-Cascade-2-30B-A3B evaluated on LiveCodeBench v6 and LiveCodeBench Pro, we use a 128K-token thinking budget, a sampling temperature of 1.0, a top-p of 0.95. For the with-tool setting, we enable tool use by appending a system-prompt postfix, allowing the model to call a stateful Python executor for up to 100 tool calls with a maximum response length of 131K tokens. We evaluate baseline models with their recommended inference configurations, ensuring a thinking budget of at least 128K tokens.

## A.3. Knowledge and STEM

For knowledge reasoning tasks, we include:

- **MMLU-Redux** (Gema et al., 2024) is a benchmark consisting of a subset of 3,000 manually re-annotated questions across 30 MMLU subjects (Hendrycks et al., 2020), which eliminates the original annotation errors. We evaluate the models in *thinking* mode and, due to the large test set size, report exact match (EM) accuracy based on a single generation per question.
- **MMLU-Pro** (Wang et al., 2024) is an enhanced version of the original MMLU benchmark that mitigates model saturation by expanding to over 12,000 graduate-level questions and increasing answer choices from four to ten. We report EM accuracy in *thinking* mode using one generation per question.
- **GPQA-Diamond** (Rein et al., 2024) is a benchmark for assessing an LLM’s scientific reasoning capability. It consists of the highest quality 198 GPQA questions covering graduate-level physics, biology, and chemistry. We report pass@1 accuracy in *thinking* mode, averaged over 8 generations per question (avg@8) to reduce variance.
- **HLE** (Phan et al., 2025) is a frontier academic reasoning benchmark spanning a broad range of expert-level subjects. We evaluate on its text-only split, which contains 2,158 examples.

For Nemotron-Cascade-2-30B-A3B evaluated on MMLU-Redux, MMLU-Pro, GPQA-Diamond and HLE in *thinking* mode, we use a temperature of 1.0, a top-p value of 0.95, and a 128K-token thinking budget (maximum response length). For HLE, we use the default system prompt and append “Please place your final answer inside

`\boxed{ }`” to each question, and use GPT-OSS-120B as the LLM judge for answer extraction and correctness verification with the prompt in Appendix C.2. Compared with the official HLE response format, which requests an explanation, an answer, and a confidence score, this boxed-answer prompt improves the accuracy by 6–7 points, primarily on the math subset, by better aligning with the answer format used in our math SFT data.

#### A.4. Alignment and Instruction-Following

For alignment tasks, we include:

- **ArenaHard 2.0** (Li et al., 2024) is a human-preference alignment benchmark featuring 750 diverse and rigorous real-user prompts. The dataset is specifically structured with 500 prompts targeting open-ended software engineering problems and complex mathematical questions, while the remaining 250 focus on creative writing. It uses an automatic LLM-as-Judge approach to estimate human preferences relative to a baseline model, enabling fully automated, low-cost, and fast evaluation without human intervention. In our experiments, we report results without style control to allow for straightforward comparison with the officially reported numbers of other models. We evaluate the models in *thinking* mode, and use GPT-4.1 as the automated judge.
- **IFBench** (Pyatkin et al., 2025) extends IFEval (Zhou et al., 2023) by introducing 58 new, diverse, and challenging verifiable out-of-domain instruction constraints. It provides a separate constraint list to ensure no overlap between training and test constraints, enabling evaluation of an LLM’s generalization ability. The test set contains 294 prompts. We report pass@1 accuracy in *thinking* mode, averaged over 8 generations (avg@8).
- **Scale AI Multi-Challenge** (Deshpande et al., 2025) is a benchmark designed to evaluate LLMs in multi-turn conversations with human users. It consists of four challenge categories: Instruction Retention, Inference Memory, Reliable Versioned Editing, and Self-Coherence. These tasks require models to simultaneously perform accurate instruction following, effective context management, and in-context reasoning. The test set contains 273 conversations in total. We report pass@1 accuracy in *thinking* mode, averaged across 10 generations (avg@10).

For Nemotron-Cascade models evaluated on IFEval in *non-thinking* mode, on IFBench and ArenaHard in *thinking* mode, we use a temperature of 0.6, a top-p value of 0.95, and a maximum response length of 32K tokens. For baseline models, we use officially reported results whenever available; if such results are absent, we evaluate them using their recommended inference configuration or the same settings as ours.

#### A.5. Long Context and Context Learning

For long context and context learning tasks, we include:

- **AA-LCR** (Team, 2025) consists of 100 challenging text-based questions that require reasoning over multiple long, real-world documents, including company reports, government consultations, legal documents, and academic papers. Each sample contains a document set averaging approximately 100k tokens. The questions are designed such that answers cannot be directly retrieved from the documents and instead require reasoning across multiple sources of information. We report pass@1 accuracy in *thinking* mode, averaged over 16 generations (avg@16).
- **LongBench v2** (Bai et al., 2025) contains 503 challenging multiple-choice questions with context lengths ranging from 8k to 2M words. The benchmark spans six task categories: single-document QA, multi-document QA, long in-context learning, long dialogue history understanding, code repository understanding, and long structured data understanding. The questions are designed to be difficult; even human experts equipped with document search tools may require substantial time to answer them correctly. We evaluate models in *thinking* mode and report pass@1 accuracy averaged over four generations (avg@4).
- **NIAH@1M (Ruler Subset)** refers to the needle-in-a-haystack (NIAH) tasks from the RULER benchmark (Hsieh et al., 2024). The NIAH test (Kamradt, 2023) assesses an LLM’s long-context ability to retrieve

a specific piece of information (the “needle”) embedded within long distractor text (the “haystack”). The RULER benchmark defines four variants of this task: Single NIAH, Multi-keys NIAH, Multi-values NIAH, and Multi-queries NIAH. Following [Blakeman et al. \(2025\)](#), we evaluate 100 instances from each category using a 1M-token context setting. Models are evaluated in reasoning-off mode, and we report pass@1 accuracy from a single generation (avg@1).

- **CL-Bench** ([Dou et al., 2026](#)) evaluates an LLM’s ability to learn from provided context and apply the acquired knowledge to solve tasks, a process referred to as context learning. The benchmark contains 1,899 test samples spanning 500 complex contexts and 31,607 verification rubrics, all developed by experienced domain experts. The knowledge required to complete these tasks largely falls outside what existing models typically learn during pre-training, requiring models to learn directly from the provided context. Models are evaluated in thinking mode, and we report pass@1 accuracy from a single generation (avg@1).

## A.6. Agentic Tasks

For agentic tasks, we include:

- **BFCL v4** ([Patil et al., 2025](#)) offers a comprehensive agentic evaluation framework for LLMs, covering tasks such as web search, memory reading and writing, and function invocation across multiple programming languages. We follow the official BFCL V4 evaluation protocol and report scores across a combination of Agentic, multi-turn, live, and non-live categories. Models are evaluated in thinking mode, and we report pass@1 accuracy based on a single generation (avg@1).
- **SWE-bench Verified** ([OpenAI, 2024](#)) is a subset of the original test set from SWE-bench ([Jimenez et al., 2023](#)), consisting of 500 samples verified to be non-problematic by human annotators. We evaluate models in *non-thinking* mode and report pass@1 accuracy, averaged over 4 generations per prompt (avg@4).
- **$\tau^2$ -Bench** ([Barres et al., 2025](#)) evaluates multi-turn customer-service agents in environments with explicit policies, tool use, and shared world-state updates. We evaluate on the three official subsets: airline (50 examples), retail (114 examples), and telecom (114 examples). To keep the standard error within 1.5, we report avg@16 on airline and avg@8 on both retail and telecom.
- **Terminal Bench 2.0** ([Merrill et al., 2026](#)) is adopted for evaluating agents in terminal-based environments, which comprises of 89 human-validated tasks across specialized fields such as scientific computing, machine learning, and system administration. Moving beyond simple code generation, this benchmark focuses on end-to-end workflows, requiring agents to demonstrate proficiency in holistic operations like model training, system configuration, and software debugging rather than just producing isolated functions. We evaluate the model using the default Terminus-2 scaffolding. We report avg@5 task success rate.

For SWE-bench Verified, we use the OpenHands scaffold ([Wang et al., 2025](#)) as the agentic coding evaluation framework. We adopt a full interaction retention policy for agent trajectories, preserving the complete history of tool calls, observations, and model outputs across turns. This includes prior file views, search results, executed commands, and intermediate patches, enabling the model to maintain state and reason effectively over long-horizon debugging processes. We set the maximum context length to 256K tokens and allow up to 200 turns, consistent with our execution-based agentic SWE-RL training configuration. Notably, this evaluation setup closely mirrors our training environment, as both rely on execution-based feedback and multi-turn interaction within the same tool-augmented scaffold. This alignment reduces train–test mismatch and enables the model to more effectively transfer learned behaviors, such as iterative debugging, hypothesis refinement, and tool-driven reasoning, to the evaluation setting.

For  $\tau^2$ -Bench evaluation, we adopt a *latest-turn thought retention* policy for managing reasoning traces in multi-turn interactions: we retain the model’s reasoning content after the most recent user turn, while discarding reasoning content from earlier turns. The official  $\tau^2$ -Bench evaluation code follows a *no thought carry-over* policy, which removes all prior reasoning content; in our experiments, this evaluation setup consistently reduces scores by 3–5 points relative to latest-turn thought retention. We attribute this gap to train–test mismatch,

since our SFT data for  $\tau^2$ -style interactions is constructed with the same latest-turn thought retention policy, which is also the thought-state management strategy used in Nemotron-3-Nano-v3 and DeepSeek-V3.2. For the telecom subset, we additionally modify the system prompt to emphasize the dual-control setting by repeating the instruction “Make sure you guide the user through the steps, do not perform user-side actions yourself.” three times. We also tested a *full thought retention* policy, which preserves reasoning content from all previous turns and more closely matches RL training, but found it gives similar accuracy to latest-turn thought retention while incurring substantially longer contexts. We therefore report our final  $\tau^2$ -Bench results using latest-turn thought retention.

## A.7. Multilingual

For multilingual tasks, we include:

- **MMLU-ProX** (Xuan et al., 2025) expands the challenging MMLU-Pro benchmark to include 29 languages. Following Blakeman et al. (2025), six languages are selected for evaluation: English (en), German (de), Spanish (es), French (fr), Italian (it), and Japanese (ja). The model is evaluated in thinking mode, and we report pass@1 accuracy from a single generation (avg@1).
- **WMT24++** (Deutsch et al., 2025) extends the WMT24 machine translation benchmark to cover 55 languages. Following Blakeman et al. (2025), we evaluate on five translation pairs: English to German (en → de), English to Spanish (en → es), English to French (en → fr), English to Italian (en → it), and English to Japanese (en → ja). We use XCOMET-XXL (Guerreiro et al., 2024) as the evaluation metric to assess the translation quality. Our model is evaluated in thinking mode, and we report pass@1 accuracy based on a single generation (avg@1).

## B. Training Hyperparameters

We list the training hyperparameters for the Nemotron-Cascade-2-30B-A3B during all stages in Table 7, 9, 10.

Table 7: Training hyperparameters for Nemotron-Cascade-2-30B-A3B in SFT.

| Hyperparameters            |                                 |
|----------------------------|---------------------------------|
| Global batch size          | 64                              |
| Packed sequence length     | 256K                            |
| Max learning rate          | $5 \times 10^{-5}$              |
| Min learning rate          | $5 \times 10^{-6}$              |
| Learning rate warmup steps | 200                             |
| Scheduler                  | cosine                          |
| Max Steps                  | 40,000                          |
| Optimizer                  | AdamW                           |
| Optimizer config           | $\beta_1 = 0.9, \beta_2 = 0.98$ |
| Weight decay               | 0.1                             |
| # of training steps        | 33,000                          |

Table 8: Training hyperparameters of Nemotron-Cascade-2-30B-A3B in Cascade RL (IF-RL, Multi-domain RL, MOPD).

| Hyper-parameters    | IF-RL              | Multi-domain RL    | MOPD               |
|---------------------|--------------------|--------------------|--------------------|
| Max response length | 49K                | 49K                | 98K                |
| Batch size          | 128                | 128                | 128                |
| # Rollout size      | 16                 | 16                 | 4                  |
| Learning rate       | $3 \times 10^{-6}$ | $3 \times 10^{-6}$ | $3 \times 10^{-6}$ |
| Steps               | 180                | 70                 | 52                 |
| Optimizer           | AdamW              | Adam               | AdamW              |
|                     | $\beta_1 = 0.9$    | $\beta_1 = 0.9$    | $\beta_1 = 0.9$    |
|                     | $\beta_2 = 0.95$   | $\beta_2 = 0.95$   | $\beta_2 = 0.95$   |
| Temperature         | 1.0                | 1.0                | 1.0                |
| Top-p               | 1.0                | 1.0                | 1.0                |
| Overlong filtering  | False              | True               | False              |

Table 9: Training hyperparameters of Nemotron-Cascade-2-30B-A3B in Cascade RL (RLHF, Long-context RL, Code RL).

| Hyper-parameters    | RLHF               | Long-context RL    | Code RL            |
|---------------------|--------------------|--------------------|--------------------|
| Max response length | 16K                | 49K                | 118K               |
| Batch size          | 128                | 128                | 128                |
| # Rollout size      | 16                 | 16                 | 16                 |
| Learning rate       | $3 \times 10^{-6}$ | $3 \times 10^{-6}$ | $3 \times 10^{-6}$ |
| Steps               | 25                 | 30                 | 22                 |
| Optimizer           | AdamW              | Adam               | AdamW              |
|                     | $\beta_1 = 0.9$    | $\beta_1 = 0.9$    | $\beta_1 = 0.9$    |
|                     | $\beta_2 = 0.95$   | $\beta_2 = 0.95$   | $\beta_2 = 0.95$   |
| Temperature         | 1.0                | 1.0                | 1.0                |
| Top-p               | 1.0                | 1.0                | 0.95               |
| Overlong filtering  | True               | True               | True               |

Table 10: Training hyperparameters of Nemotron-Cascade-2-30B-A3B model in execution-based agentic SWE-RL.

| Hyperparameters            |                    |
|----------------------------|--------------------|
| # prompts per step         | 16                 |
| # rollout                  | 64                 |
| Temperature                | 0.8                |
| Max sequence length        | 256k               |
| Max turn                   | 200                |
| Max learning rate          | $3 \times 10^{-6}$ |
| Min learning rate          | 0                  |
| Learning rate warmup steps | 10                 |

## C. Prompt Templates

### C.1. Prompt Templates for Test-Time Scaling on IOI 2025

Write Python code to solve the problem. Please place the solution code in the following format:

```
``python
```

```
# Your solution code here
```

```
``
```

```
{problem_statement}
```

Below you are provided the accepted correct solutions but with different input constraints. You may use them as a reference for your insights.

```
=====
```

```
## Different Constraints (for reference only):
```

```
{subtask_constraints}
```

```
### Accepted Code:
```

```
[CODE]
```

```
=====
```

```
## Different Constraints (for reference only):
```

```
...
```

```
=====
```

From here, you are also given your submission history containing **incorrect** code and their corresponding official judgement verdicts as reference – Official judgement verdicts and problem statement/conditions are 100% reliable. You should make improvements from them if they could help:

```
=====
```

```
### Incorrect Code
```

```
[CODE]
```

```
Judgement Verdict: [VERDICT], Score: [SCORE]
```

```
=====
```

```
### Incorrect Code
```

```
...
```

```
=====
```

## C.2. HLE Judge Prompt

Judge whether the following [response] to [question] is correct or not based on the precise and unambiguous [correct\_answer] below.

[question]: {question}

[response]: {response}

Your judgement must be in the format and criteria specified below:

extracted\_final\_answer: The final exact answer extracted from the [response]. Put the extracted answer as 'None' if there is no exact, final answer to extract from the response.

[correct\_answer]: {correct\_answer}

reasoning: Explain why the extracted\_final\_answer is correct or incorrect based on [correct\_answer], focusing only on if there are meaningful differences between [correct\_answer] and the extracted\_final\_answer. Do not comment on any background to the problem, do not attempt to solve the problem, do not argue for any answer different than [correct\_answer], focus only on whether the answers match.

correct: Answer 'yes' if extracted\_final\_answer matches the [correct\_answer] given above, or is within a small margin of error for numerical problems. Answer 'no' otherwise, i.e. if there is any inconsistency, ambiguity, non-equivalency, or if the extracted answer is incorrect.

confidence: The extracted confidence score between 0|%| and 100|%| from [response]. Put 100 if there is no confidence score available.

## D. ELO Rating Analysis

We perform ELO rating analysis on our Nemotron-Cascade-2-30B-A3B model based on **40** recent Div.1 and Div.2 Codeforces contests held between 2501–2507. Problems and evaluations are provided by LiveCodeBench Pro (Zheng et al., 2025). We adopt similar rating estimation approach as in Wang et al. (2025), by allowing model with up to  $N = 8$  submissions to each contest problems, estimating model performance and relative ranking to human contestants with expected penalty consideration. We generate the model’s responses using a temperature of 1.0,  $\text{top-p}$  of 0.95, and a maximum token budget of 128K. The performance details of our Nemotron-Cascade-2-30B-A3B model (**with** and **without** python-tool use) can be found in Table 11 and Table 12, respectively.

We observed our model’s strong code reasoning ability on solving really tough problems and achieving high ranking even on some Div. 1 rounds (Round 999, 1012, 1015, 1021 etc.), while maintaining stable performance on solving easy-medium level problems. However, the models still has weakness on dealing with problems that requiring constructive algorithms, interactive manner, and hypothesis-driven ideas.

Table 11: Nemotron-Cascade-2-30B-A3B performance details on 40 Div.1 and Div.2 Codeforces Rounds ranging from 2501 to 2507 **without** python-tool use. We attempt each problem with  $N = 8$  times in total. For regular codeforces rounds, we present the score after considering expected penalties for each problem. For ICPC style rounds, we mark passed/failed problems as + and - correspondingly. We compute the estimated rank to human contestants and the corresponding Elo score as shown in rightmost two columns.

| Contest Name  | Contest Problems |         |         |         |         |         |         |         |         |     | Score    | Penalty  | Est. Rank | ELO        |      |
|---|------------------|---------|---------|---------|---------|---------|---------|---------|---------|-----|----------|----------|-----------|------------|------|
| Hello 2025  | A                | B       | C       | D       | E1      | E2      | F       | G       | H       |     | 10779.46 | -        | 13/16703  | 3449       |      |
|   | 500.00           | 1000.00 | 1493.75 | 2235.71 | 0.0     | 1900.00 | 0.0     | 3650.00 | 0.0     |     |          |          |           |            |      |
| Codeforces Round 996 (Div. 2)                                     | A                | B       | C       | D       | E       | F       |         |         |         |     | 5793.75  | -        | 2/21232   | 2198       |      |
|   | 500.00           | 993.75  | 1475.00 | 0.0     | 2825.00 | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 997 (Div. 2)                                     | A                | B       | C       | D       | E       | F1      | F2      |         |         |     | 9378.75  | -        | 1/18823   | 2198       |      |
|   | 493.75           | 1250.00 | 1475.00 | 0.0     | 2225.00 | 2710.00 | 1225.00 |         |         |     |          |          |           |            |      |
| IAEPC Preliminary Contest (Codeforces Round 999, Div. 1 + Div. 2) | A                | B       | C       | D       | E       | F1      | F2      | G       | H1      | H2  | I        | 9278.75  | -         | 43/12647   | 3076 |
|   | 500.00           | 1000.00 | 1500.00 | 1493.75 | 1960.00 | 0.0     | 0.0     | 0.0     | 2825.00 | 0.0 | 0.0      |          |           |            |      |
| Codeforces Round 1000 (Div. 2)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          | 10935.71 | -         | 1/17169    | 2200 |
|   | 500.00           | 985.71  | 1500.00 | 2250.00 | 2687.50 | 1687.50 | 1325.00 |         |         |     |          |          |           |            |      |
| Ethflow Round 1 (Codeforces Round 1001, Div. 1 + Div. 2)          | A                | B       | C       | D       | E1      | E2      | F       | G       | H       |     |          | 2493.75  | -         | 1727/16234 | 1898 |
|   | 500.00           | 993.75  | 1000.00 | 0.0     | 0.0     | 0.0     | 0.0     | 0.0     | 0.0     |     |          |          |           |            |      |
| Codeforces Round 1002 (Div. 2)                                    | A                | B       | C       | D       | E1      | E2      |         |         |         |     |          | 3300.00  | -         | 1102/19443 | 1882 |
|   | 500.00           | 975.00  | 0.0     | 1825.00 | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1004 (Div. 1)                                    | A                | B       | C       | D1      | D2      | E       | F       |         |         |     |          | 2681.25  | -         | 145/1030   | 2666 |
|   | 0.0              | 687.50  | 1250.00 | 743.75  | 0.0     | 0.0     | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1004 (Div. 2)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          | 5397.50  | -         | 8/16749    | 2098 |
|   | 500.00           | 960.00  | 0.0     | 0.0     | 1637.50 | 2250.00 | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1005 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 9198.21  | -         | 1/17621    | 2260 |
|   | 493.75           | 1000.00 | 1243.75 | 1735.71 | 2075.00 | 2650.00 |         |         |         |     |          |          |           |            |      |
| Educational Codeforces Round 174 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 4        | 2.86      | 156/16701  | 2242 |
|   | +                | +       | +       | +       | -       | -       |         |         |         |     |          |          |           |            |      |
| Educational Codeforces Round 175 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 4        | 0.00      | 234/16060  | 2195 |
|   | +                | +       | +       | +       | -       | -       |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1007 (Div. 2)                                    | A                | B       | C       | D1      | D2      | E       | F       |         |         |     |          | 8429.46  | -         | 1/16254    | 2198 |
|   | 500.00           | 1000.00 | 1485.71 | 1743.75 | 1225.00 | 2475.00 | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1008 (Div. 1)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          | 2000.00  | -         | 355/909    | 2312 |
|   | 500.00           | 0.0     | 1500.00 | 0.0     | 0.0     | 0.0     | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1008 (Div. 2)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          | 6825.00  | -         | 9/14641    | 2008 |
|   | 500.00           | 750.00  | 1250.00 | 1575.00 | 0.0     | 2750.00 | 0.0     |         |         |     |          |          |           |            |      |
| Educational Codeforces Round 176 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 5        | 10.86     | 2/18159    | 2198 |
|   | +                | +       | +       | +       | +       | -       |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1011 (Div. 2)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          | 10137.50 | -         | 1/15906    | 2200 |
|   | 500.00           | 1250.00 | 1250.00 | 1743.75 | 2500.00 | 1993.75 | 900.00  |         |         |     |          |          |           |            |      |
| Codeforces Round 1012 (Div. 1)                                    | A                | B1      | B2      | C1      | C2      | D       | E       |         |         |     |          | 2985.00  | -         | 24/653     | 3057 |
|   | 710.00           | 975.00  | 325.00  | 975.00  | 0.0     | 0.0     | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1012 (Div. 2)                                    | A                | B       | C       | D       | E1      | E2      | F1      | F2      |         |     |          | 9945.00  | -         | 1/8536     | 2007 |
|   | 500.00           | 960.00  | 1750.00 | 1960.00 | 1975.00 | 825.00  | 1975.00 | 0.0     |         |     |          |          |           |            |      |
| Codeforces Round 1014 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 6500.00  | -         | 2/15842    | 2213 |
|   | 500.00           | 750.00  | 1250.00 | 1750.00 | 2250.00 | 0.0     |         |         |         |     |          |          |           |            |      |
| Teza Round 1 (Codeforces Round 1015, Div. 1 + Div. 2)             | A                | B       | C       | D       | E       | F       | G1      | G2      | H       |     |          | 12521.43 | -         | 4/11206    | 3830 |
|   | 750.00           | 1000.00 | 1500.00 | 1735.71 | 2235.71 | 2825.00 | 2475.00 | 0.0     | 0.0     |     |          |          |           |            |      |
| Neowise Labs Contest 1 (Codeforces Round 1018, Div. 1 + Div. 2)   | A                | B       | C       | D       | E       | F       | G       | H       |         |     |          | 4400.00  | -         | 493/12771  | 2312 |
|   | 500.00           | 750.00  | 1500.00 | 1650.00 | 0.0     | 0.0     | 0.0     | 0.0     |         |     |          |          |           |            |      |
| Codeforces Round 1019 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 4825.00  | -         | 47/14465   | 2202 |
|   | 500.00           | 1000.00 | 1500.00 | 1825.00 | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1021 (Div. 1)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 3218.75  | -         | 75/651     | 2760 |
|   | 493.75           | 900.00  | 0.0     | 1825.00 | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1021 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 8468.75  | -         | 1/5824     | 2019 |
|   | 500.00           | 1250.00 | 1493.75 | 2150.00 | 0.0     | 3075.00 |         |         |         |     |          |          |           |            |      |
| Educational Codeforces Round 178 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       | G       |         |         |     |          | 6        | 12.50     | 4/11706    | 2215 |
|   | +                | +       | +       | +       | +       | +       | -       |         |         |     |          |          |           |            |      |
| Codeforces Round 1022 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 3087.50  | -         | 308/11127  | 2132 |
|   | 500.00           | 1187.50 | 1400.00 | 0.0     | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1023 (Div. 2)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          | 6506.25  | -         | 6/11636    | 2209 |
|   | 250.00           | 750.00  | 1493.75 | 1937.50 | 0.0     | 2075.00 | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1024 (Div. 1)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 1729.46  | -         | 477/857    | 2149 |
|   | 485.71           | 1243.75 | 0.0     | 0.0     | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1024 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 3479.46  | -         | 34/11201   | 1998 |
|   | 250.00           | 500.00  | 985.71  | 1743.75 | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1025 (Div. 2)                                    | A                | B       | C1      | C2      | C3      | D       | E       | F       |         |     |          | 7985.71  | -         | 1/15945    | 2197 |
|   | 500.00           | 985.71  | 1243.75 | 575.00  | 500.00  | 1687.50 | 2493.75 | 0.0     |         |     |          |          |           |            |      |
| Codeforces Round 1026 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 9897.50  | -         | 1/17668    | 2198 |
|   | 500.00           | 750.00  | 1500.00 | 1960.00 | 2250.00 | 2937.50 |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1028 (Div. 1)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          | 2710.00  | -         | 75/956     | 2865 |
|   | 500.00           | 0.0     | 0.0     | 2210.00 | 0.0     | 0.0     | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1028 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 5453.75  | -         | 4/18314    | 2018 |
|   | 493.75           | 750.00  | 1250.00 | 0.0     | 0.0     | 2960.00 |         |         |         |     |          |          |           |            |      |
| Educational Codeforces Round 179 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       | G       |         |         |     |          | 5        | 60.00     | 94/12301   | 2231 |
|   | +                | -       | +       | +       | +       | +       | -       |         |         |     |          |          |           |            |      |
| Codeforces Round 1030 (Div. 2)                                    | A                | B       | C       | D1      | D2      | E       | F       |         |         |     |          | 7003.75  | -         | 2/18335    | 2205 |
|   | 500.00           | 975.00  | 1000.00 | 1243.75 | 960.00  | 2325.00 | 0.0     |         |         |     |          |          |           |            |      |
| Codeforces Round 1031 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 4060.71  | -         | 20/11032   | 2216 |
|   | 500.00           | 735.71  | 0.0     | 0.0     | 0.0     | 2825.00 |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1033 (Div. 2) and CodeNite 2025                  | A                | B       | C       | D       | E       | F       | G       |         |         |     |          | 9623.21  | -         | 1/12948    | 2216 |
|   | 493.75           | 750.00  | 1250.00 | 1735.71 | 2493.75 | 2900.00 | 0.0     |         |         |     |          |          |           |            |      |
| Educational Codeforces Round 180 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 5        | 33.75     | 8/17128    | 2253 |
|   | +                | +       | +       | +       | +       | +       |         |         |         |     |          |          |           |            |      |
| Codeforces Round 1035 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          | 2985.71  | -         | 587/15624  | 2008 |
|   | 500.00           | 1000.00 | 1485.71 | 0.0     | 0.0     | 0.0     |         |         |         |     |          |          |           |            |      |

Table 12: Nemotron-Cascade-2-30B-A3B performance details on 40 Div.1 and Div.2 Codeforces Rounds ranging from 2501 to 2507 **with** python-tool use. We attempt each problem with  $N = 8$  times in total. For regular codeforces rounds, we present the score after considering expected penalties for each problem. For ICPC style rounds, we mark passed/failed problems as + and - correspondingly. We compute the estimated rank to human contestants and the corresponding Elo score as shown in rightmost two columns.

| Contest Name  | Contest Problems |         |         |         |         |         |         |         |         |     | Score    | Penalty | Est. Rank  | ELO  |
|---|------------------|---------|---------|---------|---------|---------|---------|---------|---------|-----|----------|---------|------------|------|
|   | A                | B       | C       | D       | E1      | E2      | F       | G       | H       |     |          |         |            |      |
| Hello 2025  | 500.00           | 1000.00 | 1500.00 | 2225.00 | 937.50  | 1900.00 | 0.0     | 3650.00 | 0.0     |     | 11712.50 | -       | 11/16703   | 3497 |
| Codeforces Round 996 (Div. 2)                                     | A                | B       | C       | D       | E       | F       |         |         |         |     | 5025.00  | -       | 2/21232    | 2198 |
| Codeforces Round 997 (Div. 2)                                     | 500.00           | 975.00  | 1475.00 | 2075.00 | 0.0     | 0.0     |         |         |         |     | 11187.50 | -       | 1/18823    | 2198 |
| Codeforces Round 997 (Div. 2)                                     | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          |         |            |      |
| IAEPC Preliminary Contest (Codeforces Round 999, Div. 1 + Div. 2) | 493.75           | 1250.00 | 1500.00 | 1900.00 | 2075.00 | 2743.75 | 1225.00 |         |         |     | 9416.96  | -       | 40/12647   | 3097 |
| Codeforces Round 1000 (Div. 2)                                    | 500.00           | 1000.00 | 1493.75 | 1485.71 | 2000.00 | 0.0     | 0.0     | 0.0     | 2937.50 | 0.0 | 10981.25 | -       | 1/17169    | 2200 |
| Codeforces Round 1000 (Div. 2)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          |         |            |      |
| Ethflow Round 1 (Codeforces Round 1001, Div. 1 + Div. 2)          | 500.00           | 1000.00 | 1500.00 | 2243.75 | 2725.00 | 1687.50 | 1325.00 |         |         |     | 2493.75  | -       | 1727/16234 | 1898 |
| Ethflow Round 1 (Codeforces Round 1001, Div. 1 + Div. 2)          | A                | B       | C       | D       | E1      | E2      | F       | G       | H       |     |          |         |            |      |
| Codeforces Round 1002 (Div. 2)                                    | 500.00           | 993.75  | 1000.00 | 0.0     | 0.0     | 0.0     | 0.0     | 0.0     | 0.0     | 0.0 | 3300.00  | -       | 1102/19443 | 1882 |
| Codeforces Round 1002 (Div. 2)                                    | A                | B       | C       | D       | E1      | E2      |         |         |         |     |          |         |            |      |
| Codeforces Round 1004 (Div. 1)                                    | 500.00           | 975.00  | 0.0     | 1825.00 | 0.0     | 0.0     |         |         |         |     | 2743.75  | -       | 122/1030   | 2721 |
| Codeforces Round 1004 (Div. 1)                                    | A                | B       | C       | D       | D1      | D2      | E       | F       |         |     |          |         |            |      |
| Codeforces Round 1004 (Div. 2)                                    | 0.0              | 743.75  | 1250.00 | 750.00  | 0.0     | 0.0     | 0.0     | 0.0     |         |     | 2743.75  | -       | 122/1030   | 2721 |
| Codeforces Round 1004 (Div. 2)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Codeforces Round 1004 (Div. 2)                                    | 500.00           | 993.75  | 0.0     | 0.0     | 1743.75 | 2250.00 | 0.0     |         |         |     | 5487.50  | -       | 6/16749    | 2098 |
| Codeforces Round 1004 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1005 (Div. 2)                                    | 500.00           | 985.71  | 1250.00 | 1710.00 | 2210.00 | 2575.00 |         |         |         |     | 9230.71  | -       | 1/17621    | 2260 |
| Codeforces Round 1005 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Educational Codeforces Round 174 (Rated for Div. 2)               | +                | +       | +       | +       | -       | -       |         |         |         |     | 4        | 2.50    | 156/16701  | 2242 |
| Educational Codeforces Round 174 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Educational Codeforces Round 175 (Rated for Div. 2)               | +                | +       | +       | +       | +       | -       |         |         |         |     | 5        | 5.00    | 3/16060    | 2198 |
| Educational Codeforces Round 175 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1007 (Div. 2)                                    | 500.00           | 1000.00 | 1500.00 | 1725.00 | 1225.00 | 2493.75 | 0.0     |         |         |     | 8443.75  | -       | 1/16254    | 2198 |
| Codeforces Round 1007 (Div. 2)                                    | A                | B       | C       | D1      | D2      | E       | F       |         |         |     |          |         |            |      |
| Codeforces Round 1008 (Div. 1)                                    | 500.00           | 0.0     | 1500.00 | 0.0     | 0.0     | 0.0     | 0.0     |         |         |     | 2000.00  | -       | 355/909    | 2312 |
| Codeforces Round 1008 (Div. 1)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Codeforces Round 1008 (Div. 2)                                    | 500.00           | 750.00  | 1250.00 | 1725.00 | 0.0     | 2750.00 | 0.0     |         |         |     | 6975.00  | -       | 5/14641    | 2008 |
| Codeforces Round 1008 (Div. 2)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Educational Codeforces Round 176 (Rated for Div. 2)               | +                | +       | +       | +       | +       | -       |         |         |         |     | 5        | 2.50    | 2/18159    | 2198 |
| Educational Codeforces Round 176 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1011 (Div. 2)                                    | 500.00           | 1250.00 | 1250.00 | 1743.75 | 2493.75 | 2000.00 | 900.00  |         |         |     | 10137.50 | -       | 1/15906    | 2200 |
| Codeforces Round 1011 (Div. 2)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          |         |            |      |
| Codeforces Round 1012 (Div. 1)                                    | 725.00           | 975.00  | 0.0     | 993.75  | 0.0     | 0.0     | 0.0     |         |         |     | 2693.75  | -       | 66/653     | 2745 |
| Codeforces Round 1012 (Div. 1)                                    | A                | B1      | B2      | C1      | C2      | D       | E       |         |         |     |          |         |            |      |
| Codeforces Round 1012 (Div. 2)                                    | 500.00           | 1000.00 | 1750.00 | 1975.00 | 1975.00 | 0.0     | 1993.75 | 0.0     |         |     | 9193.75  | -       | 1/8536     | 2007 |
| Codeforces Round 1012 (Div. 2)                                    | A                | B       | C       | D       | E       | E2      | F1      | F2      |         |     |          |         |            |      |
| Codeforces Round 1014 (Div. 2)                                    | 500.00           | 750.00  | 1250.00 | 1750.00 | 2250.00 | 0.0     |         |         |         |     | 6500.00  | -       | 2/15842    | 2213 |
| Codeforces Round 1014 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Teza Round 1 (Codeforces Round 1015, Div. 1 + Div. 2)             | 750.00           | 1000.00 | 1500.00 | 1743.75 | 2243.75 | 0.0     | 2485.71 | 0.0     | 0.0     |     | 9723.21  | -       | 55/11206   | 3008 |
| Teza Round 1 (Codeforces Round 1015, Div. 1 + Div. 2)             | A                | B       | C       | D       | E       | F       | G1      | G2      | H       |     |          |         |            |      |
| Newwise Labs Contest 1 (Codeforces Round 1018, Div. 1 + Div. 2)   | 500.00           | 750.00  | 1500.00 | 1687.50 | 1960.00 | 0.0     | 0.0     | 0.0     |         |     | 6397.50  | -       | 70/12771   | 2933 |
| Newwise Labs Contest 1 (Codeforces Round 1018, Div. 1 + Div. 2)   | A                | B       | C       | D       | E       | F       | G       | H       |         |     |          |         |            |      |
| Codeforces Round 1019 (Div. 2)                                    | 500.00           | 1000.00 | 1500.00 | 1825.00 | 0.0     | 2900.00 |         |         |         |     | 7725.00  | -       | 2/14465    | 2202 |
| Codeforces Round 1019 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1021 (Div. 1)                                    | 493.75           | 985.71  | 1460.00 | 1960.00 | 0.0     | 0.0     |         |         |         |     | 4899.46  | -       | 21/651     | 3143 |
| Codeforces Round 1021 (Div. 1)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1021 (Div. 2)                                    | 500.00           | 1250.00 | 1493.75 | 2235.71 | 2710.00 | 3210.00 |         |         |         |     | 11399.46 | -       | 1/5824     | 2019 |
| Codeforces Round 1021 (Div. 2)                                    | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Educational Codeforces Round 178 (Rated for Div. 2)               | +                | +       | +       | +       | +       | +       | -       |         |         |     | 6        | 12.11   | 4/11706    | 2215 |
| Educational Codeforces Round 178 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Codeforces Round 1022 (Div. 2)                                    | 500.00           | 1250.00 | 1485.71 | 0.0     | 0.0     | 0.0     |         |         |         |     | 3235.71  | -       | 300/11127  | 2137 |
| Codeforces Round 1022 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1023 (Div. 2)                                    | 250.00           | 743.75  | 1493.75 | 1900.00 | 0.0     | 2150.00 | 0.0     |         |         |     | 6537.50  | -       | 6/11636    | 2209 |
| Codeforces Round 1023 (Div. 2)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          |         |            |      |
| Codeforces Round 1024 (Div. 1)                                    | 500.00           | 1250.00 | 0.0     | 2325.00 | 0.0     | 0.0     |         |         |         |     | 4075.00  | -       | 156/857    | 2590 |
| Codeforces Round 1024 (Div. 1)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1024 (Div. 2)                                    | 250.00           | 500.00  | 1000.00 | 1750.00 | 0.0     | 2825.00 |         |         |         |     | 6325.00  | -       | 2/11201    | 1998 |
| Codeforces Round 1024 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1025 (Div. 2)                                    | 500.00           | 985.71  | 1235.71 | 710.00  | 460.00  | 1710.00 | 2485.71 | 0.0     |         |     | 8087.14  | -       | 1/15945    | 2197 |
| Codeforces Round 1025 (Div. 2)                                    | A                | B       | C1      | C2      | C3      | D       | E       | F       |         |     |          |         |            |      |
| Codeforces Round 1026 (Div. 2)                                    | 500.00           | 750.00  | 1500.00 | 0.0     | 2243.75 | 2937.50 |         |         |         |     | 7931.25  | -       | 8/17668    | 2198 |
| Codeforces Round 1026 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1028 (Div. 1)                                    | 500.00           | 1225.00 | 0.0     | 2235.71 | 0.0     | 0.0     | 0.0     |         |         |     | 3960.71  | -       | 36/956     | 3093 |
| Codeforces Round 1028 (Div. 1)                                    | A                | B       | C       | D       | E       | F1      | F2      |         |         |     |          |         |            |      |
| Codeforces Round 1028 (Div. 2)                                    | 493.75           | 750.00  | 1250.00 | 1975.00 | 0.0     | 2985.71 |         |         |         |     | 7454.46  | -       | 1/18314    | 2018 |
| Codeforces Round 1028 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Educational Codeforces Round 179 (Rated for Div. 2)               | +                | +       | +       | +       | +       | +       | -       |         |         |     | 6        | 40.11   | 3/12301    | 2231 |
| Educational Codeforces Round 179 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Codeforces Round 1030 (Div. 2)                                    | 500.00           | 993.75  | 1000.00 | 1243.75 | 937.50  | 2485.71 | 0.0     |         |         |     | 7160.71  | -       | 2/18335    | 2205 |
| Codeforces Round 1030 (Div. 2)                                    | A                | B       | C       | D1      | D2      | E       | F       |         |         |     |          |         |            |      |
| Codeforces Round 1031 (Div. 2)                                    | 500.00           | 743.75  | 0.0     | 0.0     | 0.0     | 2900.00 |         |         |         |     | 4143.75  | -       | 20/11032   | 2216 |
| Codeforces Round 1031 (Div. 2)                                    | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |
| Codeforces Round 1033 (Div. 2) and CodeNite 2025                  | 500.00           | 750.00  | 1250.00 | 1750.00 | 2500.00 | 2900.00 | 0.0     |         |         |     | 9650.00  | -       | 1/12948    | 2216 |
| Codeforces Round 1033 (Div. 2) and CodeNite 2025                  | A                | B       | C       | D       | E       | F       | G       |         |         |     |          |         |            |      |
| Educational Codeforces Round 180 (Rated for Div. 2)               | +                | +       | +       | +       | +       | -       |         |         |         |     | 5        | 10.86   | 8/17128    | 2253 |
| Educational Codeforces Round 180 (Rated for Div. 2)               | A                | B       | C       | D       | E       | F       |         |         |         |     |          |         |            |      |

## E. IMO 2025 Model Solutions

### Problem 1

A line in the plane is called *sunny* if it is not parallel to any of the  $x$ -axis, the  $y$ -axis, or the line  $x + y = 0$ .

Let  $n \geq 3$  be a given integer. Determine all nonnegative integers  $k$  such that there exist  $n$  distinct lines in the plane satisfying both of the following:

- for all positive integers  $a$  and  $b$  with  $a + b \leq n + 1$ , the point  $(a, b)$  lies on at least one of the lines; and
- exactly  $k$  of the  $n$  lines are sunny.

### Short Answer

$k = 0, 1, 3$  for all  $n$

### Model Proof

**Notation.** For an integer  $n \geq 3$  define

$$S_n = \{(a, b) \in \mathbb{Z}^2 \mid a \geq 1, b \geq 1, a + b \leq n + 1\}.$$

A line in the plane is called *sunny* if its slope is neither 0 (horizontal),  $\infty$  (vertical), nor  $-1$  (the slope of the line  $x + y = 0$ ).

We have to determine all non-negative integers  $k$  for which there exist  $n$  distinct lines covering all points of  $S_n$  with exactly  $k$  of them sunny.

#### 1. Constructions showing that $k = 0, 1, 3$ are attainable

$k = 0$ .

Take the  $n$  vertical lines  $x = 1, x = 2, \dots, x = n$ . Every point  $(a, b) \in S_n$  has  $a \leq n$ , hence lies on the line  $x = a$ . These lines are not sunny. Clearly there are  $n$  distinct lines, so  $k = 0$  is possible for every  $n \geq 3$ .

$k = 1$ .

For  $n \geq 3$  take the  $n - 1$  vertical lines  $x = 1, \dots, x = n - 1$  and the line

$$L : \text{through } (n, 1) \text{ and } (n - 2, 2).$$

The slope of  $L$  is  $\frac{2-1}{(n-2)-n} = -\frac{1}{2}$ ; it is not  $0, \infty, -1$ , so  $L$  is sunny.

- Points with  $a \leq n - 1$  are covered by the vertical line  $x = a$ .

- The only point with  $a = n$  is  $(n, 1)$ , which lies on  $L$ .

Thus all points of  $S_n$  are covered. Hence  $k = 1$  is attainable for all  $n \geq 3$ .

$k = 3$ .

Case  $n = 3$ .

$S_3$  consists of the six points

$$(1, 1), (1, 2), (1, 3), (2, 1), (2, 2), (3, 1).$$

Take the three sunny lines

$$\ell_1 : y = x,$$

$$\ell_2 : y = -2x + 5,$$

$$\ell_3 : y = -\frac{1}{2}x + \frac{5}{2}.$$

Their slopes are  $1, -2, -\frac{1}{2}$  - none is  $0, \infty, -1$ , so they are sunny. One checks directly that each of the six points belongs to at least one of them. Therefore  $k = 3$  is possible for  $n = 3$ .

Case  $n \geq 4$ .

Take the  $n - 3$  vertical lines  $x = 1, x = 2, \dots, x = n - 3$ .  
 The remaining points of  $S_n$  are those with  $a \geq n - 2$ :

$$W = \{(n - 2, 1), (n - 2, 2), (n - 2, 3), (n - 1, 1), (n - 1, 2), (n, 1)\}.$$

Cover  $W$  by the following three sunny lines:

- $L_a$  : through  $(n - 2, 1)$  and  $(n - 1, 2)$  (slope 1),
- $L_b$  : through  $(n - 2, 2)$  and  $(n, 1)$  (slope  $-\frac{1}{2}$ ),
- $L_c$  : through  $(n - 2, 3)$  and  $(n - 1, 1)$  (slope  $-2$ ).

All slopes are different from  $0, \infty, -1$ , hence these lines are sunny.  
 Together with the  $n - 3$  vertical lines we have exactly  $n$  distinct lines. Every point with  $a \leq n - 3$  lies on a vertical line; the six points of  $W$  lie on  $L_a, L_b, L_c$ . Thus the whole  $S_n$  is covered. Consequently  $k = 3$  is attainable for all  $n \geq 3$ .

**2. Impossibility: no other value of  $k$  can occur**

We prove by induction on  $n$  that any covering of  $S_n$  by  $n$  lines must have  $k \in \{0, 1, 3\}$ .

**2.1. The boundary set  $X$  and a lemma**

Define the three *special* lines

$$S_1 : x = 1, \quad S_2 : y = 1, \quad S_3 : x + y = n + 1.$$

These are the sides of the right triangle that contains  $S_n$ . Let  $X$  be the set of points of  $S_n$  lying on these three sides; i.e.

$$X = \{(1, b) \mid 1 \leq b \leq n\} \cup \{(a, 1) \mid 1 \leq a \leq n\} \cup \{(a, b) \mid a + b = n + 1\}.$$

A1

There are  $|X| = 3n - 3$  distinct points (the three corners are counted twice)<sup>A2</sup>

**Lemma.** A line that is **not** one of  $S_1, S_2, S_3$  contains at most two points of  $X$ .

*Proof.* If such a line contained two distinct points from the same side, it would coincide with that side, contradicting the assumption. Hence it can contain at most one point from each side.<sup>A3</sup> Moreover, a straight line can intersect the boundary of a triangle in at most two points (unless it coincides with a side). Therefore it contains at most two points of  $X$ .  $\square$

**2.2. A counting argument for  $n \geq 4$**

Consider any covering of  $S_n$  by  $n$  lines, and let  $k$  be the number of sunny lines<sup>A4</sup>. For each line  $\ell$  denote by  $f(\ell)$  the number of points of  $X$  that lie on  $\ell$ . Because every point of  $X$  is covered at least once,

$$\sum_{\ell} f(\ell) \geq |X| = 3n - 3. \tag{1}$$

If a line is one of the special sides  $S_1, S_2, S_3$ , then it contains all  $n$  points of that side, so  $f(\ell) = n$ . If it is not a special side (whether sunny or not), the lemma gives  $f(\ell) \leq 2$ .

Let  $s$  be the number of the three special sides that actually appear among the  $n$  lines. Then

$$\sum_{\ell} f(\ell) \leq s \cdot n + (n - s) \cdot 2 = 2n + s(n - 2). \tag{2}$$

Combining (1) and (2) we obtain

$$3n - 3 \leq 2n + s(n - 2) \implies n - 3 \leq s(n - 2). \tag{3}$$

For  $n \geq 4$  we have  $n - 2 > 0$ <sup>A5</sup>, so (3) forces  $s \geq 1$ . Hence every covering of  $S_n$  with  $n \geq 4$  lines must contain at least one of the sides  $x = 1, y = 1$  or  $x + y = n + 1$ .

A1 |  
 Similar to definition of  $S_n$ , we need to require  $a, b$  are integers.

A2 |  
 What means of counted twice? Two options:  
 1. Each edge have  $n$  point, so totally  $3n$  in 3 edges, but three corners counted twice so  $3n - 3$  distinct.  
 2. Directly say  $3n - 3$  distinct points without mentioning counted twice is also okay.

A3 |  
 This sentence is not related to the proof.

A4 |  
 The definition is unnecessary since  $k$  is not used in 2.2.

A5 |  
 $n - 3 > 0$

**2.3. Reducing the problem to  $n - 1$**

Assume  $n \geq 4$  and that a covering of  $S_n$  with  $n$  lines and  $k$  sunny lines exists. By the previous paragraph, one of the special sides is present; we show how to produce a covering of  $S_{n-1}$  with  $n - 1$  lines and the same  $k$ .

*Case A:  $x = 1$  is present.*

Remove the line  $x = 1$ . The remaining  $n - 1$  lines cover all points of  $S_n$  with  $a \geq 2$ . Let

$$S' = \{(a, b) \in S_n \mid a \geq 2\},$$

and define the translation  $\varphi : S' \rightarrow S_{n-1}$  by  $\varphi(a, b) = (a - 1, b)$ <sup>A6</sup>. This is a bijection. Translate the remaining  $n - 1$  lines by  $(-1, 0)$ : for each line  $\ell$  (with  $\ell \neq (x = 1)$ ), set  $\ell' = \{(x - 1, y) \mid (x, y) \in \ell\}$ . Translation preserves slopes, so  $\ell'$  is sunny exactly when  $\ell$  is sunny, and distinct lines remain distinct. Because the original lines covered  $S'$ , their translates cover  $\varphi(S') = S_{n-1}$ . Thus we obtain a covering of  $S_{n-1}$  by  $n - 1$  lines with exactly  $k$  sunny lines.

*Case B:  $y = 1$  is present.*

Analogously, remove  $y = 1$ , translate the remaining lines by  $(0, -1)$ , and obtain a covering of  $S_{n-1}$  with  $n - 1$  lines and  $k$  sunny lines.

*Case C:  $x + y = n + 1$  is present.*

Remove this line. The set  $\{(a, b) \in S_n \mid a + b \leq n\}$  is exactly  $S_{n-1}$ . The remaining  $n - 1$  lines cover  $S_{n-1}$  directly (they cover all of  $S_n$ , and the removed line contains no point of  $S_{n-1}$ ). No translation is needed, and the number of sunny lines stays  $k$ .

In all three cases we have produced a covering of  $S_{n-1}$  by  $n - 1$  lines with the same  $k$ .

**2.4. Base case  $n = 3$  and induction**

**Base case  $n = 3$ .**

We have already exhibited coverings with  $k = 0, 1, 3$ . It remains to show that  $k = 2$  is impossible.

The set  $S_3$  is exactly the boundary  $X$  for  $n = 3$ . Observe that the only lines containing two points of  $S_3$  and being sunny are

- $\ell_1$  : through  $(1, 1), (2, 2)$  (slope 1),
- $\ell_2$  : through  $(1, 2), (3, 1)$  (slope  $-1/2$ ),
- $\ell_3$  : through  $(1, 3), (2, 1)$  (slope  $-2$ ).

Any other line containing two points has slope 0,  $\infty$ , or  $-1$ , hence is non-sunny.

Suppose a covering of  $S_3$  by three lines has exactly  $k = 2$  sunny lines. Then there is exactly one non-sunny line  $N$ .

If  $N$  is one of the special sides (say  $x = 1$ ), then  $N$  covers three points. The remaining three points must be covered by the two sunny lines. But any line through two of those three points is either vertical, horizontal, or of slope  $-1$ , all non-sunny. Hence each sunny line can contain at most one of these points, so they cover at most two points - contradiction.

If  $N$  is not a special side, then  $N$  contains at most two points. Let  $f(\ell)$  be the number of points of  $S_3$  on line  $\ell$ . To cover all six points we need  $\sum f(\ell) \geq 6$ . For a sunny line that is not special,  $f(\ell) \leq 2$  (by the lemma). If any line had  $f(\ell) \leq 1$ , then the total sum would be at most  $1 + 2 + 2 = 5 < 6$ . Therefore each line must satisfy  $f(\ell) = 2$ . In particular, the two sunny lines each contain exactly two points, and  $N$  also contains exactly two points. Thus the three lines partition  $S_3$  into three disjoint collinear pairs.

The only collinear pairs that can be covered by a sunny line are  $\{\ell_1\}, \{\ell_2\}, \{\ell_3\}$ . Hence the two sunny lines must be two of  $\{\ell_1, \ell_2, \ell_3\}$ . The uncovered pair is exactly the pair of the remaining  $\ell_m$ . For  $N$  to cover that pair,  $N$  would have to be the line through those two points, which is  $\ell_m$ , a sunny line - contradicting that  $N$  is non-sunny. Therefore  $k = 2$  is impossible. Consequently, for  $n = 3$  the only possible  $k$  are 0, 1, 3.

**Induction step for  $n \geq 4$ .**

Assume the statement holds for  $n - 1$ . Given any covering of  $S_n$  by  $n$  lines with  $k$  sunny lines, we have shown how to produce a covering of  $S_{n-1}$  by  $n - 1$  lines with the same  $k$ . By the induction hypothesis,  $k \in \{0, 1, 3\}$ . Hence the statement holds for  $n$  as well.

By induction, for every  $n \geq 3$  any covering of  $S_n$  by  $n$  lines must have  $k \in \{0, 1, 3\}$ . Together with the constructions in §1, we conclude that the set of all non-negative integers  $k$  for which such a configuration exists is exactly

A6 |

$\phi((a, b))$  since input is a pair  $(a, b)$  not two parameter  $a$  and  $b$ .

$\{0, 1, 3\}$ .

HUMAN EXPERT COMMENT

 $7/7$

**Problem 2**

Let  $\Omega$  and  $\Gamma$  be circles with centres  $M$  and  $N$ , respectively, such that the radius of  $\Omega$  is less than the radius of  $\Gamma$ . Suppose  $\Omega$  and  $\Gamma$  intersect at two distinct points  $A$  and  $B$ . Line  $MN$  intersects  $\Omega$  at  $C$  and  $\Gamma$  at  $D$ , so that  $C, M, N, D$  lie on  $MN$  in that order. Let  $P$  be the circumcentre of triangle  $ACD$ . Line  $AP$  meets  $\Omega$  again at  $E \neq A$  and meets  $\Gamma$  again at  $F \neq A$ . Let  $H$  be the orthocentre of triangle  $PMN$ . Prove that the line through  $H$  parallel to  $AP$  is tangent to the circumcircle of triangle  $BEF$ .

**Short Answer**

No short answer provided.

**Model Proof**

We set up a coordinate system with  $M = (0, 0)$  and  $N = (d, 0)$ ,  $d > 0$ .

Circles:  $\Omega : x^2 + y^2 = r^2$  (radius  $r$ ),  $\Gamma : (x - d)^2 + y^2 = R^2$  (radius  $R$ ),  $0 < r < R$ .

They intersect in two points; let  $A = (x_A, y_A)$  with  $y_A > 0$ ,  $B = (x_A, -y_A)$ . Solving the circle equations gives

$$x_A = \frac{d^2 + r^2 - R^2}{2d}, \quad y_A^2 = \frac{((R + r)^2 - d^2)(d^2 - (R - r)^2)}{4d^2}.$$

Define

$$U = R + r, \quad V = R - r, \quad K = d^2 - V^2, \quad L = U^2 - d^2.$$

Then  $4y_A^2 = \frac{LK}{d^2}$ . (Note that  $d$  satisfies  $V < d < U$  because the circles intersect in two distinct points.)

The line  $MN$  (the  $x$ -axis) meets  $\Omega$  at  $C = (-r, 0)$  and  $\Gamma$  at  $D = (d + R, 0)$ . The order on  $MN$  is  $C, M, N, D$ .

**1. Circumcenter  $P$  of  $\triangle ACD$**

Since  $C$  and  $D$  lie on the  $x$ -axis, the perpendicular bisector of  $CD$  is the vertical line  $x = h$  where

$$h = \frac{-r + (d + R)}{2} = \frac{d + R - r}{2} = \frac{d + V}{2}.$$

Thus  $P = (h, p)$  for some  $p$ . From  $PA = PC$  we obtain

$$(h - x_A)^2 + (p - y_A)^2 = (h + r)^2 + p^2.$$

Solving for  $p$  yields

$$p = -\frac{KT}{4dy_A}, \quad \text{where } T = d + R + r = d + U.$$

The vector  $\mathbf{v} = P - A$  is

$$\mathbf{v} = (v_x, v_y) = \left( \frac{VT}{2d}, -\frac{KUT}{4d^2y_A} \right).$$

**2. Second intersections  $E, F$  and midpoint  $S$  of  $EF$**

The line  $AP$  consists of points  $A + t\mathbf{v}$ . Substituting into  $\Omega$  (using  $A \cdot A = r^2$ ) gives

$$2t A \cdot \mathbf{v} + t^2 |\mathbf{v}|^2 = 0,$$

so the non-zero root is  $t_E = -2A \cdot \mathbf{v} / |\mathbf{v}|^2$ . Compute

$$A \cdot \mathbf{v} = -\frac{rT}{2} \implies t_E = \frac{rT}{|\mathbf{v}|^2}.$$

Similarly, for  $\Gamma$  we use  $(A - N) \cdot \mathbf{v} = -RT/2$  and obtain

$$t_F = \frac{RT}{|\mathbf{v}|^2}.$$

Hence  $E = A + t_E \mathbf{v}$ ,  $F = A + t_F \mathbf{v}$ . The midpoint  $S$  of  $EF$  is  $S = A + t_S \mathbf{v}$  with

$$t_S = \frac{t_E + t_F}{2} = \frac{(r + R)T}{2|\mathbf{v}|^2} = \frac{UT}{2|\mathbf{v}|^2}.$$

### 3. Unit vectors along and perpendicular to $AP$

Let  $\mathbf{e} = \mathbf{v}/|\mathbf{v}|$  and  $\mathbf{n} = (-v_y, v_x)/|\mathbf{v}|$ . Then  $\mathbf{e}$  is the direction of  $AP$ ,  $\mathbf{n}$  is perpendicular to  $AP$ , and  $EF$  is parallel to  $\mathbf{e}$ . The half-length of  $EF$  is

$$\ell = \frac{|t_F - t_E||\mathbf{v}|}{2} = \frac{VT}{2|\mathbf{v}|}, \quad \ell^2 = \frac{V^2 T^2}{4|\mathbf{v}|^2}.$$

Thus  $E = S - \ell \mathbf{e}$ ,  $F = S + \ell \mathbf{e}$ .

### 4. Orthocenter $H$ of $\triangle PMN$

Points:  $M = (0, 0)$ ,  $N = (d, 0)$ ,  $P = (h, p)$ .

The altitude from  $P$  is the vertical line  $x = h$ .

The altitude from  $M$  is perpendicular to  $PN$ ; slope of  $PN$  is  $-p/(d - h)$ , so the altitude from  $M$  has equation  $y = \frac{d-h}{p}x$ . Intersection gives

$$H = \left( h, \frac{h(d-h)}{p} \right).$$

Now  $h(d-h) = \frac{(d+v)(d-v)}{4} = \frac{d^2 - v^2}{4} = \frac{K}{4}$ . Using  $p = -\frac{KT}{4dy_A}$ , we obtain

$$H = \left( h, \frac{K/4}{-KT/(4dy_A)} \right) = \left( h, -\frac{dy_A}{T} \right).$$

### 5. Computation of $\beta = (S - B) \cdot \mathbf{n}$ and $\gamma = (S - H) \cdot \mathbf{n}$

Because  $S = A + t_S \mathbf{v}$  and  $\mathbf{n} \perp \mathbf{v}$ , we have

$$(S - B) \cdot \mathbf{n} = (A - B) \cdot \mathbf{n}, \quad (S - H) \cdot \mathbf{n} = (A - H) \cdot \mathbf{n}.$$

Thus

$$\beta = (A - B) \cdot \mathbf{n}, \quad \gamma = (A - H) \cdot \mathbf{n}.$$

Now  $A - B = (0, 2y_A)$  and  $\mathbf{n} = (-v_y, v_x)/|\mathbf{v}|$ , so

$$\beta = \frac{2y_A v_x}{|\mathbf{v}|} = \frac{2y_A}{|\mathbf{v}|} \cdot \frac{VT}{2d} = \frac{y_A VT}{d|\mathbf{v}|}.$$

To compute  $\gamma$ , we use  $\mathbf{w} \cdot \mathbf{n} = -\frac{\mathbf{w} \times \mathbf{v}}{|\mathbf{v}|}$  (for any  $\mathbf{w}$ ). Hence

$$\gamma = -\frac{(A - H) \times \mathbf{v}}{|\mathbf{v}|}.$$

We compute the cross product. Coordinates:

$$\begin{aligned} x_A - h &= -\frac{VT}{2d}, & y_A - H_y &= y_A + \frac{dy_A}{T} = y_A \frac{T+d}{T}, \\ v_x &= \frac{VT}{2d}, & v_y &= -\frac{KUT}{4d^2 y_A}. \end{aligned}$$

Then

$$(A - H) \times \mathbf{v} = (x_A - h)v_y - (y_A - H_y)v_x = \left(-\frac{VT}{2d}\right) \left(-\frac{KUT}{4d^2 y_A}\right) - \left(y_A \frac{T+d}{T}\right) \left(\frac{VT}{2d}\right) = \frac{VKUT^2}{8d^3 y_A} - \frac{Vy_A(T+d)}{2d}.$$

Factor  $\frac{V}{2d}$ :

$$(A - H) \times \mathbf{v} = \frac{V}{2d} \left( \frac{KUT^2}{4d^2y_A} - y_A(T + d) \right).$$

Let  $B = \frac{KUT^2}{4d^2y_A} - y_A(T + d)$ . Multiply by  $2y_A$ :

$$2y_AB = \frac{KUT^2}{2d^2} - 2y_A^2(T + d).$$

Using  $4y_A^2 = \frac{LK}{d^2}$  we have  $y_A^2 = \frac{LK}{4d^2}$ , so

$$2y_A^2(T + d) = \frac{LK(T + d)}{2d^2}.$$

Thus

$$2y_AB = \frac{K}{2d^2} (UT^2 - L(T + d)).$$

Now  $T = d + U$  and  $L = U^2 - d^2$ . Compute

$$UT^2 - L(T + d) = U(d + U)^2 - (U^2 - d^2)(2d + U) = 2d^2T.$$

Hence  $2y_AB = \frac{K}{2d^2} \cdot 2d^2T = KT$ , so  $B = \frac{KT}{2y_A}$ . Consequently,

$$(A - H) \times \mathbf{v} = \frac{V}{2d} \cdot \frac{KT}{2y_A} = \frac{VKT}{4dy_A}.$$

Therefore

$$\gamma = -\frac{VKT}{4dy_A|\mathbf{v}|}.$$

Comparing with  $\beta = \frac{y_AV T}{d|\mathbf{v}|}$ , we obtain

$$\gamma = -\frac{K}{4y_A^2} \beta.$$

Moreover, using  $4y_A^2 = \frac{LK}{d^2}$  we can write

$$\frac{K}{4y_A^2} = \frac{d^2}{L} \implies \gamma = -\frac{d^2}{L} \beta.$$

### 6. Expression for $|S - B|^2$

Write  $A - B = \alpha\mathbf{e} + \beta\mathbf{n}$ , where  $\alpha = (A - B) \cdot \mathbf{e} = \frac{(A - B) \cdot \mathbf{v}}{|\mathbf{v}|}$ . Compute

$$(A - B) \cdot \mathbf{v} = 2y_A v_y = 2y_A \left( -\frac{KUT}{4d^2y_A} \right) = -\frac{KUT}{2d^2},$$

so

$$\alpha = -\frac{KUT}{2d^2|\mathbf{v}|}.$$

Since  $S = A + t_S\mathbf{v}$  and  $\mathbf{v} = |\mathbf{v}|\mathbf{e}$ ,

$$S - B = (A - B) + t_S\mathbf{v} = (\alpha\mathbf{e} + \beta\mathbf{n}) + t_S|\mathbf{v}|\mathbf{e} = (\alpha + t_S|\mathbf{v}|)\mathbf{e} + \beta\mathbf{n}.$$

Now  $t_S|\mathbf{v}| = \frac{UT}{2|\mathbf{v}|}$ . Thus

$$\alpha + t_S|\mathbf{v}| = -\frac{KUT}{2d^2|\mathbf{v}|} + \frac{UT}{2|\mathbf{v}|} = \frac{UT}{2|\mathbf{v}|} \left( 1 - \frac{K}{d^2} \right) = \frac{UT}{2|\mathbf{v}|} \cdot \frac{V^2}{d^2} = \frac{UV^2T}{2d^2|\mathbf{v}|}.$$

Hence

$$|S - B|^2 = (\alpha + t_S |\mathbf{v}|)^2 + \beta^2 = \frac{U^2 V^4 T^2}{4d^4 |\mathbf{v}|^2} + \frac{y_A^2 V^2 T^2}{d^2 |\mathbf{v}|^2}.$$

Substitute  $y_A^2 = \frac{LK}{4d^2}$  to obtain

$$|S - B|^2 = \frac{V^2 T^2}{|\mathbf{v}|^2} \left( \frac{U^2 V^2}{4d^4} + \frac{LK}{4d^4} \right) = \frac{V^2 T^2}{4d^4 |\mathbf{v}|^2} (U^2 V^2 + LK).$$

### 7. Circumcenter $O$ of $\triangle BEF$ and the tangency condition

The perpendicular bisector of  $EF$  is the line through  $S$  parallel to  $\mathbf{n}$ , so we can write  $O = S + \lambda \mathbf{n}$  for some real  $\lambda$ . The circumradius  $\rho$  satisfies  $\rho^2 = |O - E|^2 = \ell^2 + \lambda^2$ . The condition  $|O - B|^2 = \rho^2$  gives

$$|(S - B) + \lambda \mathbf{n}|^2 = \ell^2 + \lambda^2.$$

Expanding, and using  $(S - B) = \alpha' \mathbf{e} + \beta \mathbf{n}$  with  $\alpha' = \alpha + t_S |\mathbf{v}|$  (so that  $|S - B|^2 = \alpha'^2 + \beta^2$ ), we get

$$\alpha'^2 + (\beta + \lambda)^2 = \ell^2 + \lambda^2 \implies |S - B|^2 + 2\lambda\beta = \ell^2. \quad (1)$$

Thus  $2\lambda\beta = \ell^2 - |S - B|^2$ .

Now consider the line through  $H$  parallel to  $AP$ . Its distance to  $O$  equals the circumradius of  $\triangle BEF$  iff

$$|(O - H) \cdot \mathbf{n}| = \rho.$$

Since  $O - H = (S - H) + \lambda \mathbf{n}$  and  $(S - H) \cdot \mathbf{n} = \gamma$ , we have  $(O - H) \cdot \mathbf{n} = \gamma + \lambda$ . Therefore tangency is equivalent to

$$(\gamma + \lambda)^2 = \lambda^2 + \ell^2 \iff \gamma^2 + 2\gamma\lambda = \ell^2. \quad (2)$$

Substituting  $\lambda = \frac{\ell^2 - |S - B|^2}{2\beta}$  from (1) into (2) yields

$$\gamma^2 + \frac{\gamma}{\beta} (\ell^2 - |S - B|^2) = \ell^2.$$

Multiplying by  $\beta$  and rearranging gives

$$\gamma(\ell^2 - |S - B|^2) = \beta(\ell^2 - \gamma^2). \quad (3)$$

### 8. Verification of (3)

We now use the explicit expressions:

$$\beta = \frac{y_A V T}{d |\mathbf{v}|}, \quad \gamma = -\frac{d^2}{L} \beta, \quad \ell^2 = \frac{V^2 T^2}{4 |\mathbf{v}|^2}, \quad |S - B|^2 = \frac{V^2 T^2}{4d^4 |\mathbf{v}|^2} (U^2 V^2 + LK).$$

Set  $c = -\frac{d^2}{L}$ , so that  $\gamma = c\beta$ . Then (3) becomes (after dividing by the non-zero factor  $\beta$ )

$$c(\ell^2 - |S - B|^2) = \ell^2 - c^2 \beta^2. \quad (3')$$

Compute each side:

$$c^2 \beta^2 = \frac{d^4}{L^2} \cdot \frac{y_A^2 V^2 T^2}{d^2 |\mathbf{v}|^2} = \frac{d^4}{L^2} \cdot \frac{LK V^2 T^2}{4d^4 |\mathbf{v}|^2} = \frac{KV^2 T^2}{4L |\mathbf{v}|^2}.$$

Hence

$$\ell^2 - c^2 \beta^2 = \frac{V^2 T^2}{4 |\mathbf{v}|^2} - \frac{KV^2 T^2}{4L |\mathbf{v}|^2} = \frac{V^2 T^2}{4 |\mathbf{v}|^2} \left( 1 - \frac{K}{L} \right) = \frac{V^2 T^2}{4 |\mathbf{v}|^2} \cdot \frac{L - K}{L}.$$

Next,

$$\ell^2 - |S - B|^2 = \frac{V^2 T^2}{4|\mathbf{v}|^2} - \frac{V^2 T^2}{4d^4 |\mathbf{v}|^2} (U^2 V^2 + LK) = \frac{V^2 T^2}{4|\mathbf{v}|^2} \left[ 1 - \frac{U^2 V^2 + LK}{d^4} \right].$$

Thus the left-hand side of (3') is

$$c(\ell^2 - |S - B|^2) = -\frac{d^2}{L} \cdot \frac{V^2 T^2}{4|\mathbf{v}|^2} \left[ 1 - \frac{U^2 V^2 + LK}{d^4} \right] = -\frac{V^2 T^2}{4L|\mathbf{v}|^2} \left( d^2 - \frac{U^2 V^2 + LK}{d^2} \right).$$

Equation (3') is therefore equivalent (after multiplying both sides by  $4L|\mathbf{v}|^2/V^2 T^2$ ) to

$$-d^2 + \frac{U^2 V^2 + LK}{d^2} = L - K.$$

Multiply by  $d^2$ :

$$-d^4 + U^2 V^2 + LK = (L - K)d^2.$$

Now substitute  $L = U^2 - d^2$ ,  $K = d^2 - V^2$ . The left side expands to

$$-d^4 + U^2 V^2 + (U^2 - d^2)(d^2 - V^2) = -d^4 + U^2 V^2 + U^2 d^2 - U^2 V^2 - d^4 + d^2 V^2 = d^2(U^2 + V^2 - 2d^2).$$

The right side is

$$(L - K)d^2 = ((U^2 - d^2) - (d^2 - V^2))d^2 = (U^2 + V^2 - 2d^2)d^2,$$

which is identical. Hence (3') holds, and consequently (3) and (2) are satisfied.

Therefore the line through  $H$  parallel to  $AP$  is tangent to the circumcircle of  $\triangle BEF$ .  $\square$

#### LLM JUDGE COMMENT

Instruction: adopt from ProofBench(Ma et al., 2025) **With Reference Solution and Marking Scheme (Strict)**

You are an **expert math proof grader**. You are judging the correctness of an LLM-generated proof for a math problem.

**### Input** Your input will consist of: **\*\*Problem Statement\*\***: A mathematical problem that the proof is attempting to solve.

**\*\*Reference Solution\*\***: A correct solution or proof provided for reference. This is **not necessarily the only valid solution**. If the problem requires a final numeric or algebraic answer, this section contains the correct answer, which should be the only accepted final answer (though alternative reasoning paths are valid).

**\*\*Marking Scheme\*\***: A problem-specific grading rubric (0-7 scale) with checkpoints, zero-credit items, and deductions. You must follow this scheme when assigning points.

**\*\*Proof Solution\*\***: The proof that you need to evaluate. This proof may contain errors, omissions, or unclear steps. The proof was generated by another language model.

**### Task** Analyze the proof carefully.

\* Compare the proof against the reference solution and the marking scheme.

\* Award points according to the marking scheme's checkpoints, zero-credit items, and deductions.

\* Identify logical errors, incorrect steps, or unclear reasoning.

\* Give a score between 0 and 7 with a brief overall assessment.

\* Show clearly how the score was derived:

\* Which checkpoints were earned (with awarded points).

Grade this as an IMO marker using ultra-strict standards. Do not fill in missing steps, do not repair the proof, and do not infer claims that are not explicitly proved. Treat any later use of an unstated or unreproved bound as a gap. First list every logical dependency line by line, then check whether each dependency has been established earlier in the written solution. Distinguish clearly between: (1) correct idea, (2) complete written proof, and (3) patchable gap. Score according to the written proof, not the intended argument

problem: {problem}

Reference Solution: {reference solution}

Marking Scheme: {marking scheme}

proof: {proof}

### GPT-5.4-Thinking(Extensive)

Thought for 2m 35s>

Here is the grading, done as an ultra-strict IMO marker on the written proof only.

Overall judgment: the submitted proof is mathematically sound and, despite being very calculation-heavy, it is a complete coordinate proof. It does not follow the reference solution's synthetic route, but it replaces the official constructions by explicit coordinate/algebraic identities and closes the tangency criterion directly. I do not see a fatal logical gap.

Dependency audit

1. The proof fixes coordinates ( $M=(0,0)$ ), ( $N=(d,0)$ ), with circles

$$\Omega : x^2 + y^2 = r^2, \quad \Gamma : (x - d)^2 + y^2 = R^2, \quad 0 < r < R.$$

This is a valid coordinate setup. No gap.

2. It writes

$$A = (x_A, y_A), \quad B = (x_A, -y_A),$$

and derives

$$x_A = \frac{d^2 + r^2 - R^2}{2d}, \quad y_A^2 = \frac{((R + r)^2 - d^2)(d^2 - (R - r)^2)}{4d^2}.$$

This follows from solving the two circle equations and using symmetry about the (x)-axis. Established.

3. It defines

$$U = R + r, \quad V = R - r, \quad K = d^2 - V^2, \quad L = U^2 - d^2,$$

and notes ( $V < d < U$ ) because the circles intersect in two distinct points. Hence ( $K > 0$ ), ( $L > 0$ ), ( $y_A > 0$ ). Established.

4. It identifies

$$C = (-r, 0), \quad D = (d + R, 0),$$

with order (C,M,N,D). Correct and established.

5. It finds the circumcenter (P) of ( $\triangle ACD$ ) as ( $P=(h,p)$ ), where

$$h = \frac{-r + (d + R)}{2} = \frac{d + V}{2}.$$

This uses that the perpendicular bisector of (CD) is vertical. Established.

6. From ( $PA=PC$ ), it states

$$p = -\frac{KT}{4dy_A}, \quad T = d + U.$$

This is a compressed algebra step, but it is a direct consequence of the displayed equation and the earlier definitions. I checked it; it is correct. No logical gap.

7. It defines ( $\mathbf{v} = P - A$ ) and gives

$$\mathbf{v} = (v_x, v_y) = \left( \frac{VT}{2d}, -\frac{KUT}{4d^2y_A} \right).$$

Again, compressed algebra, but correct and dependent only on already established formulas.

8. It parametrizes line (AP) as ( $A + t\mathbf{v}$ ). Substituting into ( $\Omega$ ), it derives the second intersection parameter

$$t_E = \frac{rT}{|\mathbf{v}|^2}.$$

This follows from

$$2t, A \cdot \mathbf{v} + t^2|\mathbf{v}|^2 = 0$$

and the computed identity  $(A \cdot \mathbf{v} = -rT/2)$ . Established.

9. Similarly, substituting into  $(\Gamma)$ , it derives

$$t_F = \frac{RT}{|\mathbf{v}|^2}.$$

This depends on  $((A - N) \cdot \mathbf{v} = -RT/2)$ . Established.

10. Hence

$$E = A + t_E \mathbf{v}, \quad F = A + t_F \mathbf{v},$$

and the midpoint (S) of (EF) is

$$S = A + t_S \mathbf{v}, \quad t_S = \frac{UT}{2|\mathbf{v}|^2}.$$

Established.

11. It defines

$$\mathbf{e} = \frac{\mathbf{v}}{|\mathbf{v}|}, \quad \mathbf{n} = \frac{(-v_y, v_x)}{|\mathbf{v}|},$$

so  $(\mathbf{e})$  is the direction of  $(AP)$ ,  $(\mathbf{n} \perp AP)$ , and  $(EF \parallel AP)$ . Established.

12. It computes the half-length of (EF):

$$\ell = \frac{|t_F - t_E| |\mathbf{v}|}{2} = \frac{VT}{2|\mathbf{v}|}.$$

Correct and established.

13. It computes the orthocenter (H) of  $(\triangle PMN)$ . From the altitude equations it gets

$$H = \left( h, \frac{h(d-h)}{p} \right),$$

and then

$$H = \left( h, -\frac{dy_A}{T} \right).$$

This is correct and fully supported by earlier formulas.

14. It defines

$$\beta = (S - B) \cdot \mathbf{n}, \quad \gamma = (S - H) \cdot \mathbf{n}.$$

Using  $(S = A + t_S \mathbf{v})$  and  $(\mathbf{n} \perp \mathbf{v})$ , it reduces these to

$$\beta = (A - B) \cdot \mathbf{n}, \quad \gamma = (A - H) \cdot \mathbf{n}.$$

Established.

15. It computes

$$\beta = \frac{y_A VT}{d|\mathbf{v}|}.$$

Correct and established.

16. It computes

$$\gamma = -\frac{(A - H) \times \mathbf{v}}{|\mathbf{v}|},$$

then works through the cross product and arrives at

$$\gamma = -\frac{VKT}{4dy_A|\mathbf{v}|} = -\frac{K}{4y_A^2} \beta = -\frac{d^2}{L} \beta.$$

This is the densest calculation in the proof, but every ingredient used has already been introduced, and the final identity checks out. No missing geometric lemma is being smuggled in.

17. It decomposes  $(A-B)$  into the  $((\mathbf{e}, \mathbf{n}))$ -basis, then obtains

$$|S - B|^2 = \frac{V^2 T^2}{4d^4 |\mathbf{v}|^2} (U^2 V^2 + LK).$$

This is a long but direct algebraic consequence of earlier formulas. Established.

18. It introduces (O), the circumcenter of  $(\triangle BEF)$ . Since  $(EF \parallel \mathbf{e})$ , the perpendicular bisector of (EF) is the line through (S) parallel to  $(\mathbf{n})$ , so

$$O = S + \lambda \mathbf{n}.$$

Correct.

19. Using  $(|OB| = |OE|)$ , it derives

$$|S - B|^2 + 2\lambda\beta = \ell^2. \quad (1)$$

Correct.

20. It then expresses the line through (H) parallel to (AP) as a line with direction  $(\mathbf{e})$ , so its distance from (O) is  $(|(O - H) \cdot \mathbf{n}|)$ . Tangency is therefore equivalent to

$$(\gamma + \lambda)^2 = \lambda^2 + \ell^2 \iff \gamma^2 + 2\gamma\lambda = \ell^2. \quad (2)$$

This is a valid tangent-to-circle criterion.

21. Eliminating  $(\lambda)$  between (1) and (2), it obtains

$$\gamma(\ell^2 - |S - B|^2) = \beta(\ell^2 - \gamma^2). \quad (3)$$

Correct algebra.

22. It substitutes the explicit expressions

$$\beta, \gamma, \ell^2, |S - B|^2$$

into (3), rewrites with  $(c = -d^2/L)$ , and reduces the whole claim to

$$-d^4 + U^2V^2 + LK = (L - K)d^2.$$

Correct reduction.

23. Finally, substituting  $(L = U^2 - d^2)$  and  $(K = d^2 - V^2)$ , both sides simplify to

$$d^2(U^2 + V^2 - 2d^2),$$

so (3) holds, hence (2) holds, hence the line through (H) parallel to (AP) is tangent to the circumcircle of  $(\triangle BEF)$ . Complete.

Correct idea vs complete proof vs patchable gaps

Correct idea: yes. The idea is to replace the synthetic construction by a coordinate model, compute (P,E,F,H), write the circumcenter (O) of  $(\triangle BEF)$  in the  $((\mathbf{e}, \mathbf{n}))$ -frame, and reduce tangency to an explicit scalar identity.

Complete written proof: yes. Every major object used later is explicitly defined, and the final tangency condition is actually proved, not merely asserted.

Patchable gaps: there are a few compressed algebra jumps (“solving for (p)”, the dot-product simplifications, the cross-product simplification, the final reduction in Section 8), but these are routine algebraic consequences of already displayed formulas, not new unproved geometric claims. Under IMO marking, I would not count these as logical gaps.

Checkpoint scoring

Because this is a full alternative coordinate proof, I map the official checkpoints by equivalent accomplished work rather than by the exact synthetic route.

1. Foundational Angle Chase — 1/1 Earned by equivalent foundational setup: the proof correctly derives the full coordinate description of (A,B,C,D,P), including the intersection formulas and all sign/positivity conditions needed later.

2. Parallelogram Construction and Alignment — 2/2 The proof does not use the official (A’)-parallelogram route, but it establishes the equivalent control of the (E,F) configuration directly: it computes the line (AP), the second intersections (E,F), the midpoint (S), and the direction/normal frame  $((\mathbf{e}, \mathbf{n}))$ . This is enough to replace the official CP2 machinery.

3. Properties of Point (T) — 2/2 The proof bypasses the auxiliary point (T) entirely and instead analyzes the actual circumcenter (O) of  $(\triangle BEF)$  directly. This is a valid stronger substitute for the official CP3 structure.

4. Properties of (H) and Final Conclusion — 2/2 Earned. The proof computes (H) explicitly, computes its signed normal offset ( $\gamma$ ), and then verifies the tangency criterion exactly.

Deductions

No cap applies, because the proof reaches the final tangency statement.

No (-1) deduction applies. I do not see a nontrivial unproved leap in the main line.

Final score

7/7

Brief assessment

This is a valid full-score coordinate solution. It is not elegant, and it is denser than the reference synthetic proof, but as written it does close all the required dependencies. Under ultra-strict marking, I would still award full marks because the proof is complete rather than merely suggestive.

**Problem 3**

Let  $\mathbb{N}$  denote the set of positive integers. A function  $f : \mathbb{N} \rightarrow \mathbb{N}$  is said to be bonza if

$$f(a) \text{ divides } b^a - f(b)^{f(a)}$$

for all positive integers  $a$  and  $b$ .

Determine the smallest real constant  $c$  such that  $f(n) \leq cn$  for all bonza functions  $f$  and all positive integers  $n$ .

**Short Answer**

4

**Model Proof**

We will prove that the smallest constant  $c$  such that  $f(n) \leq cn$  for every bonza function  $f$  is  $c = 4$ .

**1. Basic properties**

From the definition, for all  $a, b \in \mathbb{N}$ ,

$$f(a) \mid b^a - f(b)^{f(a)}. \tag{*}$$

- Taking  $b = a$  gives  $f(a) \mid a^a - f(a)^{f(a)}$ . Since  $f(a)^{f(a)}$  is a multiple of  $f(a)$ , we obtain  $f(a) \mid a^a$ . Hence every prime divisor of  $f(a)$  divides  $a$ .
- Taking  $a = b = 1$  yields  $f(1) \mid 1 - f(1)^{f(1)}$ . The right-hand side is congruent to 1 modulo  $f(1)$ , so  $f(1) \mid 1$ . Thus  $f(1) = 1$ .

**2. Behaviour on primes**

Let  $p$  be a prime. Because every prime divisor of  $f(p)$  must divide  $p$ , we have  $f(p) = p^t$  for some integer  $t \geq 0$  (where  $t = 0$  means  $f(p) = 1$ ). Define

$$S = \{ p \text{ prime} \mid f(p) \neq 1 \} = \{ p \mid t \geq 1 \}.$$

B1

If  $p \in S$ , then  $f(p) = p^t$  with  $t \geq 1$ . Using (\*) with  $a = p$  and an arbitrary  $b$  we get

$$p^t \mid b^p - f(b)^{p^t}.$$

In particular, reducing modulo  $p$ , we have

$$b^p \equiv f(b)^{p^t} \pmod{p}.$$

We now show that  $x^{p^t} \equiv x \pmod{p}$  for every integer  $x$ . This follows by induction on  $t$  using Fermat's little theorem: for  $t = 0$  it is trivial; assuming  $x^{p^t} \equiv x \pmod{p}$ , then

$$x^{p^{t+1}} = (x^{p^t})^p \equiv x^p \pmod{p} \equiv x \pmod{p},$$

where the last congruence is Fermat's little theorem. Hence  $f(b)^{p^t} \equiv f(b) \pmod{p}$ . Also  $b^p \equiv b \pmod{p}$  by Fermat. Therefore

$$b - f(b) \equiv b^p - f(b)^{p^t} \equiv 0 \pmod{p},$$

so

$$f(b) \equiv b \pmod{p} \quad \text{for all } b \in \mathbb{N}. \tag{1}$$

B1 |

$t$  should be a function w.r.t  $p$ , like  $t(p)$  or  $t_p$ , instead of a constant

**3. When  $S$  is infinite**

Assume that  $S$  is infinite. Then for any fixed  $n$ , the integer  $f(n) - n$  is divisible by every prime in  $S$ . An integer that is divisible by infinitely many distinct primes must be zero. Hence  $f(n) = n$  for all  $n$ . Consequently, the only bonza function with  $S$  infinite is the identity  $f(n) = n$ .

**4. Odd primes in  $S$  force  $S$  to be infinite**

Suppose that an odd prime  $p$  belongs to  $S$ . Let  $q$  be any prime with  $q \not\equiv 1 \pmod{p}$ . From (1) we have  $f(q) \equiv q \pmod{p}$ . Because  $f(q) \mid q^q$ , the only prime divisors of  $f(q)$  are  $q$ ; thus  $f(q)$  is a power of  $q$ , say  $f(q) = q^s$ . If  $s = 0$  then  $f(q) = 1$ , which would imply  $q \equiv 1 \pmod{p}$ , contradicting the choice of  $q$ . Hence  $s \geq 1$ , and therefore  $q \mid f(q)$ , so  $q \in S$ .

We now prove that there are infinitely many primes  $q$  with  $q \not\equiv 1 \pmod{p}$ . Assume, to the contrary, that only finitely many such primes exist; list them as  $q_1, q_2, \dots, q_k$  (this list contains  $p$  and 2 because both are not  $\equiv 1 \pmod{p}$ ). Consider

$$N = p \cdot (q_1 q_2 \cdots q_k) - 1.$$

**B2** Clearly  $N > 1$ . Any prime divisor  $r$  of  $N$  is different from every  $q_i$ , for otherwise  $r$  would divide  $p \cdot \prod q_i$  and hence divide  $N + 1$ , forcing  $r \mid 1$ , impossible. By our assumption, every prime divisor of  $N$  satisfies  $r \equiv 1 \pmod{p}$ . Consequently,  $N \equiv 1 \pmod{p}$ . But  $N = p \cdot \prod q_i - 1 \equiv -1 \pmod{p}$ , a contradiction. Hence there are infinitely many primes  $q$  with  $q \not\equiv 1 \pmod{p}$ . By the previous paragraph, each such prime lies in  $S$ , so  $S$  is infinite. Then by §3,  $f$  is the identity.

Thus, if  $f$  is not the identity, the set  $S$  cannot contain any odd prime. Therefore the only possibilities for a non-identity bonza function are

$$(i) S = \emptyset \quad \text{or} \quad (ii) S = \{2\}.$$

**5. Case  $S = \emptyset$**

This means  $f(p) = 1$  for every prime  $p$ . We show that then  $f$  is the constant function 1.

Suppose, for contradiction, that there exists  $a$  with  $f(a) > 1$ . Let  $r$  be a prime divisor of  $f(a)$ . From §1 we have  $r \mid a$ . Since  $r$  is prime and  $S = \emptyset$ , we have  $f(r) = 1$ . Apply (\*) with  $a$  (the original number) and  $b = r$ :

$$f(a) \mid r^a - f(r)^{f(a)} = r^a - 1.$$

Because  $r \mid f(a)$ , we obtain  $r \mid r^a - 1$ . But  $r^a \equiv 0 \pmod{r}$ , so  $r^a - 1 \equiv -1 \pmod{r}$ , contradiction. Hence no such  $a$  exists, and  $f(n) = 1$  for all  $n$ . The constant function 1 is clearly bonza because 1 divides any integer.

**6. Case  $S = \{2\}$**

Here  $f(2) \neq 1$  (and  $f(2) \mid 2^2 = 4$ , so  $f(2) = 2$  or 4), while  $f(p) = 1$  for every odd prime  $p$ . We will prove that  $f(n) \leq 4n$  for all  $n$ .

**6.1 Odd arguments**

Let  $n > 1$  be odd. Assume  $f(n) > 1$  and let  $r$  be a prime divisor of  $f(n)$ . Then  $r \mid n$  (since  $f(n) \mid n^n$ ), so  $r$  is an odd prime. Because  $r$  is odd,  $f(r) = 1$ . Using (\*) with  $a = n$  and  $b = r$  gives

$$f(n) \mid r^n - 1.$$

But  $r \mid f(n)$  implies  $r \mid r^n - 1$ , which is impossible because  $r^n \equiv 0 \pmod{r}$ . **Therefore  $f(n) = 1$  for all odd  $n$ . Also  $f(1) = 1$ .** <sup>B3</sup>

**6.2 Even arguments**

Now let  $n$  be even and write  $n = 2^k \cdot m$  with  $k = v_2(n) \geq 1$  and  $m$  odd.

First we show that  $f(n)$  is a power of 2. Suppose an odd prime  $r$  divides  $f(n)$ . Then, as in §6.1,  $r \mid n$  and  $f(r) = 1$  (since  $r$  is odd), and (\*) with  $b = r$  <sup>B4</sup> yields  $f(n) \mid r^n - 1$ , leading to the same contradiction  $r \mid r^n - 1$ . Hence all prime factors of  $f(n)$  are 2, so  $f(n) = 2^s$  for some  $s \geq 0$ .

Now **take  $b$  to be any odd prime  $p$**  <sup>B5</sup>. Because  $f(p) = 1$ , condition (\*) gives

**B2 |**  
correct, but no need to have a multiply as  $p$  here, since you already said  $p$  and 2 in the list.

**B3 |**  
Logical issue. It should be "Therefore,  $f(n) = 1$  for odd  $n > 1$ . Also  $f(1) = 1$ , so  $f(\text{odd}) = 1^n$ ."

**B4 |**  
missing  $a = n$

**B5 |**  
missing  $a = n$

$$2^s \mid p^n - 1. \tag{2}$$

Thus  $s \leq v_2(p^n - 1)$  for every odd prime  $p$ . Consequently,

$$s \leq \min_{\text{odd prime } p} v_2(p^n - 1).$$

We compute this minimum. Since  $n$  is even, write  $n = 2^k \cdot m$  with  $m$  odd<sup>B6</sup>. For an odd prime  $p$ , we have (using the lifting-the-exponent lemma or a direct argument)

$$v_2(p^n - 1) = v_2(p^{2^k} - 1) = v_2(p - 1) + v_2(p + 1) + k - 1. \tag{3}$$

*Proof of (3):* Because  $m$  is odd,  $v_2(p^n - 1) = v_2(p^{2^k} - 1)$ . For  $k \geq 1$ , one proves by induction that  $v_2(p^{2^k} - 1) = v_2(p - 1) + v_2(p + 1) + k - 1$ . The base  $k = 1$  is  $p^2 - 1 = (p - 1)(p + 1)$ . For the inductive step,  $p^{2^k} - 1 = (p^{2^{k-1}} - 1)(p^{2^{k-1}} + 1)$ . The first factor contributes  $v_2(p^{2^{k-1}} - 1)$ , and for  $k \geq 2$  one has  $v_2(p^{2^{k-1}} + 1) = 1$  because  $p^{2^{k-1}} \equiv 1 \pmod{8}$  (as  $p$  odd implies  $p^2 \equiv 1 \pmod{8}$ ), and by induction all higher powers are  $\equiv 1 \pmod{8}$ ). This gives the formula.

For an odd prime  $p$ , the two consecutive even numbers  $p - 1$  and  $p + 1$  have the following property: one of them is divisible by 2 exactly once, the other is divisible by 4 (at least). Hence

$$v_2(p - 1) + v_2(p + 1) \geq 1 + 2 = 3,$$

and the minimum 3 is attained for infinitely many primes, e.g.  $p = 3, 5, 11, 13, \dots$ . Therefore

$$\min_{\text{odd prime } p} v_2(p^n - 1) = 3 + k - 1 = k + 2.$$

Thus  $s \leq k + 2$ , and

$$f(n) = 2^s \leq 2^{k+2} = 4 \cdot 2^k.$$

Since  $n = 2^k \cdot m$  with  $m \geq 1$ , we have  $2^k \leq n$ . Hence

$$f(n) \leq 4 \cdot 2^k \leq 4n.$$

For odd  $n$  we already have  $f(n) = 1 \leq 4n$ . The identity function  $f(n) = n$  also satisfies this bound. Consequently, every bonza function satisfies  $f(n) \leq 4n$  for all  $n$ .<sup>B7</sup>

**B6 |**  
already defined at the beginning of 6.2

**B7 |**  
You divided into 3 cases, and here ignored the situation with  $f(n) = 1$  always.

### 7. Sharpness

We construct a bonza function  $F$  for which  $F(4) = 16$ , showing that  $c$  cannot be smaller than 4.

Define

$$F(1) = 1, \quad F(2) = 2,$$

$$F(n) = 1 \quad \text{for odd } n > 1,$$

and for even  $n \geq 4$ , let  $k = v_2(n)$  and set  $F(n) = 4 \cdot 2^k = 2^{k+2}$ .

(One could also take  $F(2) = 4$ ; both work, but we choose  $F(2) = 2$  for definiteness.)

We verify that  $F$  satisfies the bonza condition.

- **Case  $a$  odd:** Then  $F(a) = 1$  and  $1 \mid b^a - F(b)^1$  trivially.
- **Case  $a = 2$ :**  
 $b$  odd:  $F(2) = 2$  divides  $b^2 - 1$  because  $b^2$  is odd, so  $b^2 - 1$  is even.  
 $b$  even: Both  $b^2$  and  $F(b)^2$  are even, so their difference is even, hence divisible by 2.
- **Case  $a \geq 4$  even:** Write  $a = 2^k \cdot m$  with  $k = v_2(a) \geq 1$  and  $m$  odd, so  $F(a) = 2^{k+2}$ . Consider  $b$ .  
 $b$  odd:  $F(b) = 1$ . For any odd  $b$  and even  $a$ , we have  

$$v_2(b^a - 1) = v_2(b - 1) + v_2(b + 1) + k - 1 \geq 3 + k - 1 = k + 2,$$
 so  $2^{k+2} \mid b^a - 1$ .  
 $b$  even: Write  $F(b) = 2^t$  with  $t \geq 1$  (if  $b = 2$  then  $t = 1$ ; if  $b \geq 4$  then  $t = v_2(b) + 2 \geq 3$ ).

Then

$$v_2(b^a) = a \cdot v_2(b) \geq a \geq k + 2 \quad (\text{since } a \geq 4 \text{ and even}),$$

and

$$v_2(F(b)^{F(a)}) = t \cdot F(a) = t \cdot 2^{k+2} \geq 2^{k+2}.$$

**B8** Hence both  $b^a$  and  $F(b)^{F(a)}$  are multiples of  $2^{k+2}$ , and therefore their difference is also a multiple of  $2^{k+2}$ .

Thus in all situations  $F(a) \mid b^a - F(b)^{F(a)}$ . Therefore  $F$  is bonza. For  $n = 4$  we have  $v_2(4) = 2$ , so  $F(4) = 2^4 = 16$ , giving  $F(4)/4 = 4$ . Hence any constant  $c$  with  $f(n) \leq cn$  for all bonza functions must satisfy  $c \geq 4$ .

### 8. Conclusion

We have shown that every bonza function  $f$  satisfies  $f(n) \leq 4n$  for all  $n \in \mathbb{N}$ , and we have exhibited a bonza function for which equality  $f(4) = 4 \cdot 4$  is attained (i.e.,  $f(4)/4 = 4$ ). Consequently, the smallest possible constant is

$$\boxed{4}.$$

HUMAN EXPERT COMMENT

$$\boxed{7/7}$$

**B8** |

$k + 2$  should be enough

### Problem 4

A proper divisor of a positive integer  $N$  is a positive divisor of  $N$  other than  $N$  itself.

The infinite sequence  $a_1, a_2, \dots$  consists of positive integers, each of which has at least three proper divisors.

For each  $n \geq 1$ , the integer  $a_{n+1}$  is the sum of the three largest proper divisors of  $a_n$ .

Determine all possible values of  $a_1$ .

#### Short Answer

The set of all possible values for  $a_1$  are integers  $a$  of the form  $a = 2^x 3^y m$ , where: 1.  $x$  is a positive odd integer. 2.  $y$  is a positive integer satisfying  $y \geq \frac{x+1}{2}$ . 3.  $m$  is a positive integer whose prime factors are all greater than or equal to 7.

### Model Proof

#### 1. Preliminaries

For a positive integer  $N$  with at least three proper divisors, list all its divisors in increasing order:

$$1 = d_1 < d_2 < d_3 < d_4 < \dots < d_m = N.$$

Because of the pairing  $d_i \cdot d_{m+1-i} = N$ , the three largest proper divisors are

$$\frac{N}{d_2}, \frac{N}{d_3}, \frac{N}{d_4}.$$

Hence we can write

$$f(N) := \text{sum of the three largest proper divisors of } N = N \left( \frac{1}{d_2} + \frac{1}{d_3} + \frac{1}{d_4} \right). \quad (1)$$

C1

The sequence is defined by  $a_{n+1} = f(a_n)$ , and we require that every term has at least three proper divisors.

C1 |

In the correct formulation, the  $d_2, d_3, d_4$  should be function of  $N$ .

#### 2. Two key lemmas

##### Lemma 1 (odd numbers).

If  $N$  is odd and has at least three proper divisors, then  $f(N) < N$  and  $f(N)$  is odd.

*Proof.*

Because all divisors of an odd number are odd, we have  $d_2 \geq 3, d_3 \geq 5, d_4 \geq 7$ . Hence

$$\frac{1}{d_2} + \frac{1}{d_3} + \frac{1}{d_4} \leq \frac{1}{3} + \frac{1}{5} + \frac{1}{7} = \frac{71}{105} < 1,$$

so  $f(N) < N$ . Moreover, each of  $N/d_2, N/d_3, N/d_4$  is odd (odd divided by odd), so their sum is odd.  $\square$

##### Lemma 2 (even numbers not divisible by 3).

If  $N$  is even,  $3 \nmid N$ , and has at least three proper divisors, then

- (i)  $f(N) < N$ , and
- (ii)  $f(N)$  is **not** a multiple of 6.

*Proof.*

Since  $N$  is even,  $d_2 = 2$ . The smallest possible values for  $d_3, d_4$  that maximize the sum in (1) are  $d_3 = 4$  and  $d_4 = 5$  (e.g.,  $N = 20$ ). In any case  $d_3 \geq 4$  and  $d_4 \geq 5$  (because 4 is the only

even number between 2 and 5, and 5 is the smallest integer  $>4$  that does not force a factor 3). Consequently

$$\frac{1}{2} + \frac{1}{d_3} + \frac{1}{d_4} \leq \frac{1}{2} + \frac{1}{4} + \frac{1}{5} = \frac{19}{20} < 1,$$

so  $f(N) < N$ .

We now show that  $f(N)$  cannot be divisible by 6.

Write  $N = 2^a \cdot m$  with  $m$  odd and  $3 \nmid m$  ( $a \geq 1$ ). Consider two cases.

*Case  $a \geq 2$ .*

Then  $4 \mid N$ , so  $d_3 = 4$ . Hence

$$f(N) = \frac{N}{2} + \frac{N}{4} + \frac{N}{d_4} = \frac{3N}{4} + \frac{N}{d_4}.$$

Because  $N/4$  is an integer,  $3N/4$  is a multiple of 3. Since  $N$  is not divisible by 3,  $N/d_4$  is also not divisible by 3. Thus  $f(N) \equiv N/d_4 \not\equiv 0 \pmod{3}$ , and therefore  $f(N)$  is not a multiple of 3; in particular it cannot be a multiple of 6.

*Case  $a = 1$ .*

Then  $N = 2m$  with  $m$  odd,  $3 \nmid m$ . Let  $p$  be the smallest prime divisor of  $m$ ; then  $p \geq 5$  and  $d_3 = p$ . The fourth divisor  $d_4$  is either

- an odd divisor  $q$  with  $p < q < 2p$  (if such a divisor exists), or
- $2p$  (if no odd divisor lies between  $p$  and  $2p$ ).

We examine the two possibilities.

- **Subcase  $d_4$  odd.**

Then  $N/d_4 = 2(m/q)$  is even,  $N/2 = m$  is odd, and  $N/p = 2(m/p)$  is even. Hence  $f(N) = \text{odd} + \text{even} + \text{even} = \text{odd}$ , so  $f(N)$  is not even and certainly not a multiple of 6.

- **Subcase  $d_4 = 2p$  (even).**

Here

$$f(N) = \frac{N}{2} + \frac{N}{p} + \frac{N}{2p} = m + 2\frac{m}{p} + \frac{m}{p} = m + 3\frac{m}{p}.$$

Since  $m$  is not divisible by 3, we have  $f(N) \equiv m \pmod{3} \not\equiv 0$ . Thus  $f(N)$  is not a multiple of 3 and therefore not a multiple of 6.

In every subcase  $f(N)$  is not a multiple of 6, completing the proof.  $\square$

### 3. All terms must be multiples of 6

Assume, for contradiction, that some term  $a_n$  is **not** a multiple of 6.

If  $a_n$  is *odd*, Lemma 1 gives that  $a_{n+1} = f(a_n)$  is odd and  $a_{n+1} < a_n$ . Repeating, we obtain an infinite strictly decreasing sequence of positive odd integers - impossible by the well-ordering principle.

If  $a_n$  is *even but  $3 \nmid a_n$* , Lemma 2 tells us that  $a_{n+1} = f(a_n) < a_n$  and  $a_{n+1}$  is **not** a multiple of 6. There are two possibilities for  $a_{n+1}$ :

- $a_{n+1}$  **is odd**. Then by Lemma 1 all subsequent terms are odd and strictly decreasing - again impossible.
- $a_{n+1}$  **is even and  $3 \nmid a_{n+1}$** . Then we can apply the same reasoning to  $a_{n+1}$ : we obtain  $a_{n+2} < a_{n+1}$  and  $a_{n+2}$  not a multiple of 6, and so on. This would produce an infinite strictly decreasing sequence  $a_n > a_{n+1} > a_{n+2} > \dots$  of positive integers, which cannot exist.

Thus, in every case, the assumption leads to a contradiction. Consequently, **every term of the infinite sequence is a multiple of 6.**

### 4. Reduction to the form $6M$

Write  $a_n = 6M_n$  with  $M_n \in \mathbb{N}$ . Because  $6M_n$  always has at least the divisors 1, 2, 3, 6, we have  $\tau(6M_n) \geq 4$  (i.e., at least three proper divisors).

For  $N = 6M$ , the three smallest divisors greater than 1 are always  $d_2 = 2$  and  $d_3 = 3$ . The fourth divisor  $d_4$  depends on  $M$ :

- If  $M$  is even  $\rightarrow 4 \mid 6M$ , and since 4 is the smallest integer larger than 3, we have  $d_4 = 4$ .
- If  $M$  is odd and  $5 \mid M \rightarrow 5 \mid 6M$  and  $5 < 6$ , so  $d_4 = 5$ .
- If  $M$  is odd and  $5 \nmid M \rightarrow$  the next divisor is 6 (because  $6 \mid 6M$  and neither 4 nor 5 divides  $6M$ ), so  $d_4 = 6$ .

Using (1) we compute

$$f(6M) = \frac{6M}{2} + \frac{6M}{3} + \frac{6M}{d_4} = 3M + 2M + \frac{6M}{d_4} = 5M + \frac{6M}{d_4}. \tag{2}$$

Now analyse the three cases.

1.  $M$  odd,  $5 \nmid M \rightarrow d_4 = 6$

$$f(6M) = 5M + M = 6M.$$

Hence  $6M$  is a fixed point.

2.  $M$  odd,  $5 \mid M \rightarrow d_4 = 5$

$$f(6M) = 5M + \frac{6M}{5} = \frac{31}{5}M.$$

Since  $M$  is odd and divisible by 5,  $M/5$  is odd, so the result is odd. Therefore it is **not** a multiple of 6. By the result of §3, such a number cannot appear in an infinite sequence.

3.  $M$  even  $\rightarrow d_4 = 4$

$$f(6M) = 5M + \frac{6M}{4} = \frac{13}{2}M.$$

For this to be a multiple of 6 (necessary for the next term to be admissible) we need

$$\frac{13}{2}M \equiv 0 \pmod{6} \iff 13M \equiv 0 \pmod{12} \iff M \equiv 0 \pmod{12},$$

C2

because  $13 \equiv 1 \pmod{12}$ . Hence, if  $M$  is even but **not** divisible by 12, then  $f(6M)$  would not be a multiple of 6, contradicting §3. Therefore, in an infinite sequence, every even  $M$  must be divisible by 12, and then

$$a_{n+1} = f(6M) = 6 \cdot \left(13 \cdot \frac{M}{12}\right),$$

so the new parameter is  $M' = 13 \cdot (M/12)$ .

C2 |  
Notation inconsistency. Just use  $\nmid$  and  $\mid$ .

### 5. Characterising admissible $M_1 = a_1/6$

Let  $M_1 = a_1/6$ . We have shown that all terms must be multiples of 6, and the recurrence for the corresponding  $M_n$  is:

- If  $M$  is odd and  $5 \nmid M \rightarrow$  fixed point.
- If  $M$  is even and divisible by 12  $\rightarrow M \mapsto 13 \cdot (M/12)$ .
- Any other situation leads to a contradiction.

For the sequence to be infinite, we must start from  $M_1$  and, after finitely many applications of the **even step**<sup>C3</sup>, reach an odd  $M$  that is not divisible by 5. This forces  $M_1$  to have a very specific structure.

Let  $k$  be the largest integer such that  $12^k \mid M_1$  (the exponent of 12 in  $M_1$ ). Write

$$M_1 = 12^k \cdot d,$$

C3 |  
even step? 1. Second cases  $M \rightarrow 13M/12$ .  
2. even  $M$

where  $d$  is a positive integer **not** divisible by 12 (i.e.,  $d$  is the “remainder” after removing all factors of 12).

We claim that the sequence is infinite **iff**  $d$  is **odd** and  $5 \nmid d$ .

### Sufficiency

Assume  $M_1 = 12^k \cdot d$  with  $k \geq 0$ ,  $d$  odd, and  $5 \nmid d$ . We prove by induction that

$$M_{i+1} = 13^i \cdot \frac{M_1}{12^i} = 12^{k-i} \cdot 13^i \cdot d \quad (0 \leq i \leq k).$$

For  $i = 0$  this is the definition of  $M_1$ .

*Inductive step:* For  $i < k$ , we have  $M_i = 12^{k-i} \cdot 13^i \cdot d$ . Since  $k-i \geq 1$ ,  $M_i$  is even and divisible by 12 (because it contains the factor  $12^{k-i}$ ). Hence we may apply the even-step rule, obtaining

$$M_{i+1} = 13 \cdot \frac{M_i}{12} = 13 \cdot 12^{k-i-1} \cdot 13^i \cdot d = 12^{k-i-1} \cdot 13^{i+1} \cdot d,$$

which matches the formula.

Thus for  $i = k-1$  we get  $M_k = 12^1 \cdot 13^{k-1} \cdot d$ , which is even and divisible by 12. Applying the rule once more gives

$$M_{k+1} = 13 \cdot \frac{M_k}{12} = 13^k \cdot d,$$

which is **odd** (since  $13^k$  is odd and  $d$  is odd) and not divisible by 5. By case 1 of §4,  $6M_{k+1}$  is a fixed point. Consequently,

$$a_{k+1} = 6 \cdot 13^k \cdot d,$$

and the sequence becomes constant thereafter. All terms have at least three proper divisors, so the sequence is infinite.

### Necessity

Suppose, for contradiction, that  $M_1$  does **not** have the form  $12^k \cdot d$  with  $d$  odd and  $5 \nmid d$ . Write  $M_1 = 12^k \cdot d$  as above (now  $d$  may be even, or odd but divisible by 5).<sup>C4</sup> We consider two cases.

- **$d$  is even.**<sup>C5</sup>

Then  $M_1$  is even. If  $k = 0$ , then  $M_1$  is even but not divisible by 12 (because  $d$  even and not a multiple of 12). By §4, case 3, this would give  $a_2$  not a multiple of 6 - contradiction. If  $k \geq 1$ , then  $M_1$  is even and divisible by 12. Apply the even step once to obtain

$$M_2 = 13 \cdot \frac{M_1}{12} = 13 \cdot 12^{k-1} \cdot d.$$

Since  $d$  is even,  $M_2$  is even. We now check whether  $M_2$  is divisible by 12.

- If  $k-1 \geq 1$  (i.e.,  $k \geq 2$ ), then  $12^{k-1}$  already contains a factor 12, so  $M_2$  is divisible by 12. But then we can continue the process. However, after **exactly**  $k$  steps we will obtain

$$M_{k+1} = 13^k \cdot d,$$

which is even (because  $d$  is even). Since  $d$  is not divisible by 12 (by the definition of  $k$ ),  $M_{k+1}$  is **not** divisible by 12. Hence at that moment we would have an even  $M$  that is not divisible by 12, which forces  $f(6M_{k+1})$  not to be a multiple of 6 - contradiction.

- If  $k = 1$ , then  $M_1 = 12 \cdot d$  with  $d$  even. Then  $M_2 = 13d$  is even but not divisible by 12 (because  $13d$  lacks the necessary factor 3 and<sup>C6</sup> the factor 4 if  $d$  is not a multiple of 12). Again,  $M_2$  fails the requirement that every even  $M$  must be divisible by 12.

In all subcases, contradiction arises.

- **$d$  is odd but  $5 \mid d$ .**

Then after  $k$  applications of the even step we obtain

$$M_{k+1} = 13^k \cdot d,$$

C4 |

This formulation should mention  $12 \nmid d$

C5 |

No need to consider for even or odd, we can directly say  $M_{k+1} = 13^k \cdot d$  and then contradictory. We already have the dynamic of  $M_i$  at the beginning.

C6 |

or

which is odd and divisible by 5. Consequently,  $a_{k+1} = 6M_{k+1}$  falls into case 2 of §4, yielding  $f(a_{k+1})$  odd. As in §3, an odd term would lead to an infinite strictly decreasing sequence of odd numbers - contradiction.

Thus, the only way to avoid contradictions is that  $d$  is odd and  $5 \nmid d$ .

---

### 6. Final answer

Combining everything, we obtain that the initial value  $a_1$  can be exactly those numbers of the form

$$a_1 = 6 \cdot 12^k \cdot d \quad (k \in \mathbb{N}_0, d \text{ odd}, 5 \nmid d).$$

Equivalently,  $a_1 = 2^{2k+1} \cdot 3^{k+1} \cdot d$  with  $d$  odd and not divisible by 5.

---

HUMAN EXPERT COMMENT

7/7

### Problem 5

Alice and Bazza are playing the inekoalaty game, a two-player game whose rules depend on a positive real number  $\lambda$  which is known to both players. On the  $n$ th turn of the game (starting with  $n = 1$ ) the following happens:

- If  $n$  is odd, Alice chooses a nonnegative real number  $x_n$  such that

$$x_1 + x_2 + \cdots + x_n \leq \lambda n.$$

- If  $n$  is even, Bazza chooses a nonnegative real number  $x_n$  such that

$$x_1^2 + x_2^2 + \cdots + x_n^2 \leq n.$$

If a player cannot choose a suitable  $x_n$ , the game ends and the other player wins. If the game goes on forever, neither player wins. All chosen numbers are known to both players.

Determine all values of  $\lambda$  for which Alice has a winning strategy and all those for which Bazza has a winning strategy.

#### Short Answer

The outcome of the inekoalty game depends on the value of  $\lambda$  as follows:

- If  $\lambda > \frac{\sqrt{2}}{2}$ , Alice has a winning strategy.
- If  $\lambda < \frac{\sqrt{2}}{2}$ , Bazza has a winning strategy.
- If  $\lambda = \frac{\sqrt{2}}{2}$ , neither player has a winning strategy, and the game can go on forever.

#### Model Proof

##### Notation and basic facts

Define  $S_n = \sum_{i=1}^n x_i$  and  $Q_n = \sum_{i=1}^n x_i^2$ .

The rules: on odd  $n$ ,  $S_n \leq \lambda n$ ; on even  $n$ ,  $Q_n \leq n$ .

From Cauchy-Schwarz, after any even turn  $2k$  we have

$$S_{2k}^2 \leq (2k)Q_{2k} \leq (2k) \cdot 2k = (2k)^2 \implies S_{2k} \leq 2k. \quad (1)$$

For  $k \geq 1$  define the **slack** after the even turn  $2k$ :

$$d_k = \lambda(2k + 1) - S_{2k}.$$

Alice can move on turn  $2k + 1$  iff  $d_k \geq 0$  (she may choose  $x_{2k+1} = 0$ ).

If at some odd turn Alice can legally choose a number  $u > \sqrt{2}$  while  $Q_{2k} = 2k$ , then  $Q_{2k+1} = 2k + u^2 > 2k + 2$  and Bazza will have no legal move on the next turn, so Alice wins immediately.

In the general case, we will use the slack to analyse the game.

##### Key strategy for Bazza

We consider the following natural strategy for Bazza on even turns: after Alice's move  $u$  on turn  $2k + 1$ , if the game has not ended, Bazza chooses

$$v = \sqrt{2k + 2 - Q_{2k+1}},$$

i.e. the largest possible number that still satisfies  $Q_{2k+2} \leq 2k + 2$ .

This choice is always legal as long as  $Q_{2k+1} \leq 2k + 2$  and makes  $Q_{2k+2} = 2k + 2$  (the

maximal possible sum of squares).

We call this the **maximal strategy** for Bazza.

*Claim 1.* As long as Alice never picks a number  $u > \sqrt{2}$  (which would make her win immediately), the maximal strategy forces  $Q_{2k} = 2k$  for every  $k \geq 1$ .

*Proof.* By induction. For  $k = 1$ , after Alice's first move  $x_1 = a$  (with  $0 \leq a \leq \lambda$ ), Bazza can choose  $v = \sqrt{2 - a^2}$  because  $a^2 \leq 2$  (this holds for all relevant  $\lambda$ ; in the cases where we apply it, we have  $\lambda \leq \sqrt{2}/2 < \sqrt{2}$ , and for larger  $\lambda$  Alice would win earlier). Then  $Q_2 = a^2 + v^2 = 2$ . Inductive step: assume  $Q_{2k} = 2k$ . Alice's move  $u$  satisfies  $Q_{2k+1} = 2k + u^2 \leq 2k + 2$  because otherwise  $u^2 > 2$  and Alice would have won (this can only happen if  $Q_{2k} = 2k$  and  $u > \sqrt{2}$ ).<sup>D1</sup> If the game continues, we must have  $u \leq \sqrt{2}$ . Then  $Q_{2k+1} = 2k + u^2 \leq 2k + 2$ . Bazza chooses  $v = \sqrt{2k + 2 - (2k + u^2)} = \sqrt{2 - u^2}$ , yielding  $Q_{2k+2} = 2k + 2$ .  $\square$

Do not understand. Reads like a thinking draft process.

**Analysis of the slack under the maximal strategy**

Suppose the maximal strategy is followed and that the game has not ended (so far Alice has never picked  $u > \sqrt{2}$ ). Let  $u = x_{2k+1}$  and  $v = x_{2k+2}$ . Then

$$S_{2k+2} = S_{2k} + u + v, \quad Q_{2k+2} = 2k + 2,$$

and therefore

$$d_{k+1} = \lambda(2k + 3) - S_{2k+2} = d_k + 2\lambda - (u + v). \tag{2}$$

For  $0 \leq u \leq \sqrt{2}$  define  $h(u) = u + \sqrt{2 - u^2}$ . One checks that  $h(u) \geq \sqrt{2}$ , with equality exactly at  $u = 0$  and  $u = \sqrt{2}$ , and maximum 2 at  $u = 1$ . Consequently

$$d_{k+1} \leq d_k + 2\lambda - \sqrt{2}, \tag{3}$$

and equality can be achieved by choosing  $u = 0$  (or  $u = \sqrt{2}$ , if allowed).

Thus the maximal possible increase of  $d_k$  (when Alice tries to keep her slack large) is  $2\lambda - \sqrt{2}$ , attained by taking  $u = 0$ .

**Three regimes**

---

1.  $\lambda < \frac{\sqrt{2}}{2}$

Then  $2\lambda - \sqrt{2} < 0$ .

*Bazza's winning strategy.*

Bazza adopts the maximal strategy described above.

First, compute an upper bound for  $d_1$ . After the first two moves,

$S_2 = a + \sqrt{2 - a^2}$  with  $0 \leq a \leq \lambda$ .<sup>D2</sup> The function  $f(a) = a + \sqrt{2 - a^2}$  satisfies  $f(a) \geq \sqrt{2}$  (minimum at  $a = 0$ ), so

D2 | Better to use  $x_1$  not  $a$

$$d_1 = 3\lambda - S_2 \leq 3\lambda - \sqrt{2}. \tag{4}$$

Because  $\lambda < \sqrt{2}/2$ , we have  $3\lambda - \sqrt{2} < \sqrt{2}/2 < \sqrt{2}$ .

Now, from (3) we obtain for any possible play

$$d_k \leq d_1 + (k - 1)(2\lambda - \sqrt{2}) \leq (3\lambda - \sqrt{2}) + (k - 1)(2\lambda - \sqrt{2}). \tag{5}$$

(The first inequality follows by induction using  $d_{k+1} \leq d_k + 2\lambda - \sqrt{2}$ , the second uses the bound on  $d_1$ .)

Since  $2\lambda - \sqrt{2} < 0$ , the right-hand side of (5) tends to  $-\infty$  as  $k \rightarrow \infty$ . Hence there exists a finite  $K$  with  $d_K < 0$ .

When  $d_K < 0$ , at the beginning of turn  $2K + 1$  Alice cannot move (even  $x = 0$  would violate the sum condition), so Bazza wins.

Moreover, we claim that Alice never gets the chance to win earlier by picking  $u > \sqrt{2}$ . Indeed, from (5) and  $2\lambda - \sqrt{2} < 0$ , we have  $d_k \leq 3\lambda - \sqrt{2} < \sqrt{2}$  for all  $k$ . Hence  $d_k$  never exceeds  $\sqrt{2}$ , so she cannot legally pick a number larger than  $\sqrt{2}$  (which would be required to win immediately). Thus Bazza's strategy is winning.

**Conclusion for  $\lambda < \sqrt{2}/2$ :** Bazza has a winning strategy.

2.  $\lambda = \frac{\sqrt{2}}{2}$

Here  $2\lambda - \sqrt{2} = 0$ .

*Bazza does not have a winning strategy.*

We exhibit a strategy for Alice that prevents Bazza from ever winning.

Let Alice always choose  $x_1 = 0$  and thereafter  $x_{2k+1} = 0$  on every odd turn.

Consider any Bazza moves. On even turns we have only Bazza's numbers, say  $y_1, \dots, y_k$  at turn  $2k$ . The constraints are  $\sum_{i=1}^k y_i^2 = Q_{2k} \leq 2k$  (because odd terms are zero). By Cauchy-Schwarz,

$$S_{2k} = \sum_{i=1}^k y_i \leq \sqrt{k \sum_{i=1}^k y_i^2} \leq \sqrt{k \cdot 2k} = \sqrt{2}k.$$

Therefore

$$d_k = \lambda(2k + 1) - S_{2k} \geq \frac{\sqrt{2}}{2}(2k + 1) - \sqrt{2}k = \frac{\sqrt{2}}{2} > 0.$$

Thus Alice can always move (she may choose 0). Also, she never wins because to win she would need to make  $Q_{2k+1} > 2k + 2$ , which would require  $x_{2k+1}^2 > 2k + 2 - Q_{2k} \geq 2$ , i.e.  $x_{2k+1} > \sqrt{2}$ . But  $x_{2k+1} \leq d_k$  and we have not shown an upper bound on  $d_k$ ;<sup>D4</sup> however, the crucial point is that with this Alice strategy, Bazza never wins because Alice never loses. So this suffices to prove that Bazza does not have a winning strategy (he cannot force a win against every Alice strategy).

*Alice does not have a winning strategy.*

We now show that Bazza has a strategy to prevent Alice from winning.

Let Bazza adopt the maximal strategy (as in Case 1). We prove that under this strategy, Alice can never win.

First, by Claim 1, as long as the game continues,  $Q_{2k} = 2k$ . Using (3) with  $\lambda = \sqrt{2}/2$  we get

$$d_{k+1} = d_k + \sqrt{2} - (u + v) \leq d_k, \tag{6}$$

because  $u + v \geq \sqrt{2}$ . Hence  $d_k$  is non-increasing. Moreover, from (4) we have

$$d_1 \leq 3\lambda - \sqrt{2} = \sqrt{2}/2.$$

Thus  $d_k \leq \sqrt{2}/2 < \sqrt{2}$  for all  $k$ . In particular,  $d_k$  never exceeds  $\sqrt{2}$ , so Alice can never pick a number  $u > \sqrt{2}$  (which would be necessary to win immediately).

Could Alice win by picking some  $u \leq \sqrt{2}$  and later build up? To win, she must at some odd turn make  $Q_{2k+1} > 2k + 2$ . Since  $Q_{2k} = 2k$ , this would require  $u^2 > 2$ , i.e.  $u > \sqrt{2}$ . Thus any winning move would necessarily be  $u > \sqrt{2}$ . Because  $d_k \leq \sqrt{2}/2 < \sqrt{2}$ , such a move is impossible. Therefore Alice never wins.<sup>D5</sup>

Additionally, we check that  $d_k$  never becomes negative, so Alice always can move. Indeed, using (2) and the fact that  $u \leq d_k$  (so  $u \leq \sqrt{2}/2 < 1$ ), the smallest possible  $d_{k+1}$  for a given  $d_k$  is attained when  $u = d_k$ , giving  $d_{k+1} = \sqrt{2} - \sqrt{2 - d_k^2} \geq 0$ . Hence  $d_k \geq 0$  for all  $k$ .<sup>D6</sup>

Thus the maximal strategy by Bazza prevents Alice from winning, and also avoids his own loss (the game can continue indefinitely, e.g., if Alice always chooses 0, then  $d_k = \sqrt{2}/2$  constant and  $Q_{2k} = 2k$ , so the play never ends). Consequently, Alice cannot force a win; she has no winning strategy.

Since neither player has a strategy that forces a win, the game is a draw for  $\lambda = \sqrt{2}/2$ .

3.  $\lambda > \frac{\sqrt{2}}{2}$

Now  $2\lambda - \sqrt{2} > 0$ .

*Alice's winning strategy.*

She will always choose 0 on her odd turns, until she can force a win.

Specifically, on turn 1 she sets  $x_1 = 0$ . On each subsequent odd turn  $2k + 1$  ( $k \geq 1$ ), she examines the current slack  $d_k = \lambda(2k + 1) - S_{2k}$ . If  $d_k > \sqrt{2k + 2 - Q_{2k}}$ , she chooses a

**D4 |**  
This conclusion only focuses on a specific  $B$ 's policy, which already discussed at the beginning. This looks like thinking draft process.

**D5 |**  
??? Reads like a thinking draft process.

**D6 |**  
No need. The proof is finished.

number  $u$  with

$$\sqrt{2k+2-Q_{2k}} < u \leq d_k,$$

such that  $u^2 > 2k+2-Q_{2k}$  (this is possible because  $d_k$  is strictly larger). She then wins on the next turn because  $Q_{2k+1} = Q_{2k} + u^2 > 2k+2$ . If  $d_k \leq \sqrt{2k+2-Q_{2k}}$ , she simply chooses  $x_{2k+1} = 0$ .

We must verify that this strategy is legal and that it indeed leads to a win.

First, note that as long as Alice always picks 0, all her moves contribute nothing to  $Q$  and  $S_{2k}$  is the sum of Bazza's even moves. For any Bazza strategy we have

$$Q_{2k} = \sum_{i=1}^k x_{2i}^2 \leq 2k, \quad S_{2k} = \sum_{i=1}^k x_{2i}.$$

By Cauchy-Schwarz,

$$S_{2k}^2 \leq k Q_{2k} \leq 2k^2 \implies S_{2k} \leq \sqrt{2} k. \quad (7)$$

Consequently,

$$d_k = \lambda(2k+1) - S_{2k} \geq \lambda(2k+1) - \sqrt{2} k = (2\lambda - \sqrt{2})k + \lambda. \quad (8)$$

Because  $2\lambda - \sqrt{2} > 0$ , the right-hand side of (8) tends to  $+\infty$  as  $k \rightarrow \infty$ . In particular, there exists an integer  $K$  such that

$$d_K > \sqrt{2K+2-Q_{2K}}.$$

(Since the right side is at most  $\sqrt{2K+2}$ , it suffices that  $(2\lambda - \sqrt{2})K + \lambda > \sqrt{2K+2}$ ; this holds for all sufficiently large  $K$ .)

At the start of turn  $2K+1$ , we have  $d_K > \sqrt{2K+2-Q_{2K}}$ . Choose

$$u = \sqrt{2K+2-Q_{2K}} + \varepsilon,$$

with  $\varepsilon > 0$  small enough that  $u \leq d_K$  (possible because  $d_K$  is strictly larger). Then

$$S_{2K+1} = S_{2K} + u \leq S_{2K} + d_K = \lambda(2K+1),$$

so the sum condition is satisfied, and

$$Q_{2K+1} = Q_{2K} + u^2 > Q_{2K} + (2K+2-Q_{2K}) = 2K+2.$$

Therefore on turn  $2K+2$  Bazza has no legal move (even  $x = 0$  would give  $Q > 2K+2$ ). Alice wins.

It remains to check that Alice can always make the choice  $x_{2k+1} = 0$  when  $d_k \leq \sqrt{2k+2-Q_{2k}}$ . The condition for choosing 0 is that  $0 \leq d_k$ , i.e.  $d_k \geq 0$ . From (8),  $d_k \geq (2\lambda - \sqrt{2})k + \lambda > 0$  for all  $k$  (since the right side is positive for  $k = 1$  because  $3\lambda - \sqrt{2} > 0$  for  $\lambda > \sqrt{2}/2$ ). Hence Alice can always pick 0 when she does not want to win yet. Thus her strategy is well defined and guarantees a win in finite time.

**Conclusion for  $\lambda > \sqrt{2}/2$ :** Alice has a winning strategy.

**Final answer**

- Alice has a winning strategy **iff**  $\lambda > \frac{\sqrt{2}}{2}$ .
- Bazza has a winning strategy **iff**  $\lambda < \frac{\sqrt{2}}{2}$ .

- For  $\lambda = \frac{\sqrt{2}}{2}$  neither player can force a win; the game can be drawn.

**HUMAN EXPERT COMMENT**

Some parts of the proof read like a thinking process. It feels like it realizes that  $c$  has not been proved, then decides that proving  $c$  is not actually necessary. The proof seems finished, but then it starts thinking again.

|     |
|-----|
| 7/7 |
|-----|

## References

- [1] Gretel synthetic safety alignment dataset, 12 2024. URL <https://huggingface.co/datasets/gretelai/gretel-safety-alignment-en-v1>. 8
- [2] Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes. In *The twelfth international conference on learning representations*, 2024. 12
- [3] Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*, 2025. 7, 8, 9, 20
- [4] Wasi Uddin Ahmad, Sean Narenthiran, Somshubra Majumdar, Aleksander Ficek, Siddhartha Jain, Jocelyn Huang, Vahid Noroozi, and Boris Ginsburg. Opencodereasoning: Advancing data distillation for competitive coding. *arXiv preprint arXiv:2504.01943*, 2025. 7
- [5] Ibragim Badertdinov, Alexander Golubev, Maksim Nekrashevich, Anton Shevtsov, Simon Karasik, Andrei Andriushchenko, Maria Trofimova, Daria Litvintseva, and Boris Yangel. Swe-rebench: An automated pipeline for task collection and decontaminated evaluation of software engineering agents. *arXiv preprint arXiv:2505.20411*, 2025. 9
- [6] Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, et al. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3639–3664, 2025. 22
- [7] Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan.  $\tau^2$ -bench: Evaluating conversational agents in a dual-control environment, 2025. URL <https://arxiv.org/abs/2506.07982>. 23
- [8] Aaron Blakeman, Aaron Grattafiori, Aarti Basant, Abhibha Gupta, Abhinav Khattar, Adi Renduchintala, Aditya Vavre, Akanksha Shukla, Akhiad Bercovich, Aleksander Ficek, et al. Nemotron 3 nano: Open, efficient mixture-of-experts hybrid mamba-transformer model for agentic reasoning. *arXiv preprint arXiv:2512.20848*, 2025. 7, 8, 9, 11, 12, 14, 15, 23, 24
- [9] Aaron Blakeman, Aaron Grattafiori, Aarti Basant, Abhibha Gupta, Abhinav Khattar, Adi Renduchintala, Aditya Vavre, Akanksha Shukla, Akhiad Bercovich, Aleksander Ficek, et al. Nvidia nemotron 3: Efficient and open intelligence. *arXiv preprint arXiv:2512.20856*, 2025. 5, 9
- [10] Yang Chen, Zhuolin Yang, Zihan Liu, Chankyu Lee, Peng Xu, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acereason-nemotron: Advancing math and code reasoning through reinforcement learning. *Advances in neural information processing systems*, 2025. 13
- [11] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating llms by human preference, 2024. 14
- [12] Kaustubh Deshpande, Ved Sirdeshmukh, Johannes Baptist Mols, Lifeng Jin, Ed-Yeremai Hernandez-Cardona, Dean Lee, Jeremy Kritz, Willow E Primack, Summer Yue, and Chen Xing. Multichallenge: A realistic multi-turn conversation evaluation benchmark challenging to frontier llms. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 18632–18702, 2025. 22

- [13] Daniel Deutsch, Eleftheria Briakou, Isaac Rayburn Caswell, Mara Finkelstein, Rebecca Galor, Juraj Juraska, Geza Kovacs, Alison Lui, Ricardo Rei, Jason Riesa, et al. Wmt24++: Expanding the language coverage of wmt24 to 55 languages & dialects. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12257–12284, 2025. 24
- [14] Shihan Dou, Ming Zhang, Zhangyue Yin, Chenhao Huang, Yujiong Shen, Junzhe Wang, Jiayi Chen, Yuchen Ni, Junjie Ye, Cheng Zhang, et al. Cl-bench: A benchmark for context learning. *arXiv preprint arXiv:2602.03587*, 2026. 23
- [15] Wei Du, Shubham Toshniwal, Branislav Kisacanin, Sadegh Mahdavi, Ivan Moshkov, George Armstrong, Stephen Ge, Edgar Minasyan, Feng Chen, and Igor Gitman. Nemotron-math: Efficient long-context distillation of mathematical reasoning from multi-mode supervision. *arXiv preprint arXiv:2512.15489*, 2025. 7
- [16] Tony Feng, Trieu H. Trinh, Garrett Bingham, Dawsen Hwang, Yuri Chervonyi, Junehyuk Jung, Joonkyung Lee, Carlo Pagano, Sang-hyun Kim, Federico Pasqualotto, Sergei Gukov, Jonathan N. Lee, Junsu Kim, Kaiying Hou, Golnaz Ghiasi, Yi Tay, YaGuang Li, Chenkai Kuang, Yuan Liu, Hanzhao Lin, Evan Zheran Liu, Nigamaa Nayakanti, Xiaomeng Yang, Heng-Tze Cheng, Demis Hassabis, Koray Kavukcuoglu, Quoc V. Le, and Thang Luong. Towards autonomous mathematics research. *arXiv preprint arXiv:2602.10177*, 2026. doi: 10.48550/arXiv.2602.10177. 17
- [17] Aryo Pradipta Gema, Joshua Ong Jun Leang, Giwon Hong, Alessio Devoto, Alberto Carlo Maria Mancino, Rohit Saxena, Xuanli He, Yu Zhao, Xiaotang Du, Mohammad Reza Ghasemi Madani, Claire Barale, Robert McHardy, Joshua Harris, Jean Kaddour, Emile van Krieken, and Pasquale Minervini. Are we done with mmlu?, 2024. 21
- [18] Gemini Team. A new era of intelligence with gemini 3. <https://blog.google/products-and-platforms/products/gemini/gemini-3/>, 2025. Google Blog, November 18, 2025. 17
- [19] Gemini Team. Advanced version of gemini with deep think officially achieves gold-medal standard at the international mathematical olympiad. <https://deepmind.google/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad/>, 2025. Google DeepMind Blog, July 21, 2025. 6, 17
- [20] Gemini Team. Gemini 3 deep think: Advancing science, research and engineering. <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-deep-think/>, 2026. Google Blog, February 12, 2026. 17
- [21] Shaona Ghosh, Prasoon Varshney, Makes Narsimhan Sreedhar, Aishwarya Padmakumar, Traian Rebedea, Jibin Rajan Varghese, and Christopher Parisien. Aegis2. 0: A diverse ai safety dataset and risks taxonomy for alignment of llm guardrails. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5992–6026, 2025. 8
- [22] Nuno M Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André FT Martins. xcomet: Transparent machine translation evaluation through fine-grained error detection. *Transactions of the Association for Computational Linguistics*, 12:979–995, 2024. 24
- [23] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 4

- [24] Adib Hasan, Ileana Rugina, and Alex Wang. Pruning for protection: Increasing jailbreak resistance in aligned llms without fine-tuning. In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 417–430, 2024. 8
- [25] Zhongmou He, Yee Man Choi, Kexun Zhang, Jiabao Ji, Junting Zhou, Dejie Xu, Ivan Bercovich, Aidan Zhang, and Lei Li. Hardtests: Synthesizing high-quality test cases for llm coding. *arXiv preprint arXiv:2505.24098*, 2025. 7
- [26] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020. 21
- [27] HMMT. Harvard-mit mathematics tournament february 2025, 2025. 20
- [28] Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekeshe, Fei Jia, Yang Zhang, and Boris Ginsburg. Ruler: What’s the real context size of your long-context language models? *arXiv preprint arXiv:2404.06654*, 2024. 22
- [29] Siming Huang, Tianhao Cheng, Jason Klein Liu, Jiaran Hao, Liuyihan Song, Yang Xu, J Yang, Jiaheng Liu, Chenchen Zhang, Linzheng Chai, et al. Opencoder: The open cookbook for top-tier code large language models. *arXiv preprint arXiv:2411.04905*, 2024. 7
- [30] IMO. International Mathematical Olympiad, 2025. 20
- [31] Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024. 18, 21
- [32] Naman Jain, Jaskirat Singh, Manish Shetty, Liang Zheng, Koushik Sen, and Ion Stoica. R2e-gym: Procedural environments and hybrid verifiers for scaling open-weights swe agents, 2025. URL <https://arxiv.org/abs/2504.07164>. 9, 16
- [33] Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2023. 23
- [34] Gregory Kamradt. Needle in a haystack - pressure testing llms. *Github*, 2023. URL [https://github.com/gkamradt/LLMTest\\_NeedleInAHaystack/tree/main](https://github.com/gkamradt/LLMTest_NeedleInAHaystack/tree/main). 22
- [35] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 12, 14, 15
- [36] Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline. *arXiv preprint arXiv:2406.11939*, 2024. 14, 22
- [37] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-V3 technical report. *arXiv preprint arXiv:2412.19437*, 2024. 8
- [38] Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, et al. Deepseek-v3. 2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*, 2025. 4, 6, 7, 9
- [39] Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. AceMath: Advancing frontier math reasoning with post-training and reward modeling. *ACL*, 2024. 20

- [40] Zihan Liu, Zhuolin Yang, Yang Chen, Chankyu Lee, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acereason-nemotron 1.1: Advancing math and code reasoning through sft and rl synergy. *ICLR*, 2026. 20
- [41] LM-Provers, Yuxiao Qu, Amrith Setlur, Jasper Dekoninck, Edward Beeching, Jia Li, Ian Wu, Lewis Tunstall, and Aviral Kumar. Qed-nano: Teaching a tiny model to prove hard theorems. <https://huggingface.co/spaces/lm-provers/qed-nano-blogpost>, 2026. Blog post. 17
- [42] Kevin Lu and Thinking Machines Lab. On-policy distillation. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20251026. <https://thinkingmachines.ai/blog/on-policy-distillation>. 12
- [43] Weidi Luo, Siyuan Ma, Xiaogeng Liu, Xiaoyu Guo, and Chaowei Xiao. Jailbreakv: A benchmark for assessing the robustness of multimodal large language models against jailbreak attacks. *arXiv preprint arXiv:2404.03027*, 2024. 8
- [44] Minh-Thang Luong, Dawsen Hwang, Hoang H Nguyen, Golnaz Ghiasi, Yuri Chervonyi, Insuk Seo, Junsu Kim, Garrett Bingham, Jonathan Lee, Swaroop Mishra, et al. Towards robust mathematical reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 35406–35430, 2025. 17, 20
- [45] Wenjie Ma, Andrei Cojocaru, Neel Kolhe, Bradley Louie, Robin Said Sharif, Haihan Zhang, Vincent Zhuang, Matei Zaharia, and Sewon Min. Reliable fine-grained evaluation of natural language math proofs. *arXiv preprint arXiv:2510.13888*, 2025. 6, 38
- [46] MAA. American Invitational Mathematics Examination - AIME 2025, 2025. 20
- [47] MAA. American Invitational Mathematics Examination - AIME 2026, 2026. 20
- [48] Mike A. Merrill, Alexander G. Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Yeon Shin, Thomas Walshe, E. Kelly Buchanan, Junhong Shen, Guanghao Ye, Haowei Lin, Jason Poulos, Maoyu Wang, Marianna Nezhurina, Jenia Jitsev, Di Lu, Orfeas Menis Mastromichalakis, Zhiwei Xu, Zizhao Chen, Yue Liu, Robert Zhang, Leon Liangyu Chen, Anurag Kashyap, Jan-Lucas Uslu, Jeffrey Li, Jianbo Wu, Minghao Yan, Song Bian, Vedang Sharma, Ke Sun, Steven Dillmann, Akshay Anand, Andrew Lanpouthakoun, Bardia Koopah, Changran Hu, Etash Guha, Gabriel H. S. Dreiman, Jiacheng Zhu, Karl Krauth, Li Zhong, Niklas Muennighoff, Robert Amanfu, Shangyin Tan, Shreyas Pimpalgaonkar, Tushar Aggarwal, Xiangning Lin, Xin Lan, Xuandong Zhao, Yiqing Liang, Yuanli Wang, Zilong Wang, Changzhi Zhou, David Heineman, Hange Liu, Harsh Trivedi, John Yang, Junhong Lin, Manish Shetty, Michael Yang, Nabil Omi, Negin Raoof, Shanda Li, Terry Yue Zhuo, Wuwei Lin, Yiwei Dai, Yuxin Wang, Wenhao Chai, Shang Zhou, Dariush Wahdany, Ziyu She, Jiaming Hu, Zhikang Dong, Yuxuan Zhu, Sasha Cui, Ahson Saiyed, Arinbjörn Kolbeinsson, Jesse Hu, Christopher Michael Rytting, Ryan Marten, Yixin Wang, Alex Dimakis, Andy Konwinski, and Ludwig Schmidt. Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces, 2026. URL <https://arxiv.org/abs/2601.11868>. 9, 23
- [49] NVIDIA. Nemo gym: An open source library for scaling reinforcement learning environments for llm. <https://github.com/NVIDIA-NeMo/Gym>, 2025. GitHub repository. 15
- [50] NVIDIA. NeMo RL: A Scalable and Efficient Post-Training Library. <https://github.com/NVIDIA-NeMo/RL>, 2025. GitHub repository. 10
- [51] NVIDIA. NeMo-Skills. <https://github.com/NVIDIA-NeMo/Skills>, 2025. GitHub repository. 20
- [52] NVIDIA. Nemotron-3-Nano-30B-A3B-Base-BF16, 2025. URL <https://huggingface.co/nvidia/NVIDIA-Nemotron-3-Nano-30B-A3B-Base-BF16>. 6
- [53] OpenAI. Introducing SWE-bench Verified, 2024. 23

- [54] OpenAI. Introducing GPT-5, 2025. [17](#)
- [55] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *NeurIPS*, 35, 2022. [4](#)
- [56] Jiayi Pan\*, Xingyao Wang\*, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. Training software engineering agents and verifiers with swe-gym. In *ICML*, 2025. URL <https://arxiv.org/abs/2412.21139>. [9](#), [16](#)
- [57] Shishir G. Patil, Huanzhi Mao, Charlie Cheng-Jie Ji, Fanjia Yan, Vishnu Suresh, Ion Stoica, and Joseph E. Gonzalez. The Berkeley Function Calling Leaderboard (BFCL): From Tool Use to Agentic Evaluation of Large Language Models. In *ICML*, 2025. [23](#)
- [58] Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, and et al. Humanity’s last exam, 2025. URL <https://arxiv.org/abs/2501.14249>. [21](#)
- [59] Renjie Pi, Grace Lam, Mohammad Shoeybi, Pooya Jannaty, Bryan Catanzaro, and Wei Ping. On data engineering for scaling llm terminal capabilities, 2026. URL <https://arxiv.org/abs/2602.21193>. [9](#)
- [60] Valentina Pyatkin, Saumya Malik, Victoria Graf, Hamish Ivison, Shengyi Huang, Pradeep Dasigi, Nathan Lambert, and Hannaneh Hajishirzi. Generalizing verifiable instruction following, 2025. URL [https://huggingface.co/datasets/allenai/IF\\_multi\\_constraints\\_upto5](https://huggingface.co/datasets/allenai/IF_multi_constraints_upto5). [11](#), [22](#)
- [61] Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026. URL <https://qwen.ai/blog?id=qwen3.5>. [5](#)
- [62] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024. [21](#)
- [63] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. DeepseekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. [10](#)
- [64] Zhihong Shao, Yuxiang Luo, Chengda Lu, ZZ Ren, Jiewen Hu, Tian Ye, Zhibin Gou, Shirong Ma, and Xiaokang Zhang. Deepseekmath-v2: Towards self-verifiable mathematical reasoning. *arXiv preprint arXiv:2511.22570*, 2025. [16](#), [17](#)
- [65] Atefeh Sohrabizadeh, Jialin Song, Mingjie Liu, Rajarshi Roy, Chankyu Lee, Jonathan Raiman, and Bryan Catanzaro. Nemotron-cortexa: Enhancing llm agents for software engineering tasks via improved localization and solution diversity. In *Forty-second International Conference on Machine Learning*, 2025. [16](#)
- [66] Artificial Analysis Team. Artificial analysis long context reasoning benchmark(lcr), 2025. [22](#)
- [67] Minyang Tian, Luyu Gao, Shizhuo D Zhang, Xinan Chen, Cunwei Fan, Xuefei Guo, Roland Haas, Pan Ji, Kittithat Krongchon, Yao Li, et al. Scicode: A research coding benchmark curated by scientists. *Advances in Neural Information Processing Systems*, 37:30624–30650, 2024. [21](#)
- [68] Boxin Wang, Chankyu Lee, Nayeon Lee, Sheng-Chieh Lin, Wenliang Dai, Yang Chen, Yangyi Chen, Zhuolin Yang, Zihan Liu, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Nemotron-Cascade: Scaling cascaded reinforcement learning for general-purpose reasoning models. *arXiv preprint arXiv:2512.13607*, 2025. [4](#), [6](#), [7](#), [8](#), [9](#), [10](#), [11](#), [14](#), [15](#), [17](#), [27](#)

- [69] Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. Openhands: An open platform for AI software developers as generalist agents. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=OJd3ayDDoF>. 9, 16, 23
- [70] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*, 2024. 21
- [71] Zhilin Wang, Jiaqi Zeng, Olivier Delalleau, Hoo-Chang Shin, Felipe Soares, Alexander Bukharin, Ellie Evans, Yi Dong, and Oleksii Kuchaiev. Helpsteer3-preference: Open human-annotated preference data across diverse tasks and languages, 2025. URL <https://arxiv.org/abs/2505.11475>. 14
- [72] Zhilin Wang, Jiaqi Zeng, Olivier Delalleau, Hoo-Chang Shin, Felipe Soares, Alexander Bukharin, Ellie Evans, Yi Dong, and Oleksii Kuchaiev. Helpsteer3-preference: Open human-annotated preference data across diverse tasks and languages, 2025. URL <https://arxiv.org/abs/2505.11475>. 14
- [73] Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, Lingming Zhang, Daniel Fried, Gabriel Synnaeve, Rishabh Singh, and Sida I Wang. Swe-rl: Advancing llm reasoning via reinforcement learning on open software evolution. *arXiv preprint arXiv:2502.18449*, 2025. 9
- [74] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992. 11
- [75] Bangjun Xiao, Bingquan Xia, Bo Yang, Bofei Gao, Bowen Shen, Chen Zhang, Chenhong He, Chiheng Lou, Fuli Luo, Gang Wang, et al. Mimo-v2-flash technical report. *arXiv preprint arXiv:2601.02780*, 2026. 4, 12
- [76] Peng Xu, Wei Ping, Xianchao Wu, Chejian Xu, Zihan Liu, Mohammad Shoeybi, and Bryan Catanzaro. Chatqa 2: Bridging the gap to proprietary llms in long context and rag capabilities. *arXiv preprint arXiv:2407.14482*, 2024. 8
- [77] Weihao Xuan, Rui Yang, Heli Qi, Qingcheng Zeng, Yunze Xiao, Aosong Feng, Dairui Liu, Yun Xing, Junjue Wang, Fan Gao, et al. Mmlu-prox: A multilingual benchmark for advanced large language model evaluation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 1513–1532, 2025. 24
- [78] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. 8, 9, 12, 14
- [79] John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 50528–50652. Curran Associates, Inc., 2024. doi: 10.52202/079017-1601. URL [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/5a7c947568c1b1328ccc5230172e1e7c-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/5a7c947568c1b1328ccc5230172e1e7c-Paper-Conference.pdf). 9
- [80] Zonghan Yang, Shengjie Wang, Kelin Fu, Wenyang He, Weimin Xiong, Yibo Liu, Yibo Miao, Bofei Gao, Yejie Wang, YINGWEI MA, Yanhao Li, Yue Liu, Zhenxing Hu, kaitai zhang, Shuyi Wang, Huarong Chen, Flood Sung, Yang Liu, Yang Gao, Zhilin Yang, and Tianyu Liu. Kimi-dev: Agentless training as skill prior for SWE-agents. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=tYppHuGhxJ>. 16

- [81] Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. DAPO: An open-source LLM reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025. 11
- [82] Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, Da Yin, Chendi Ge, Chengxing Xie, Cunxiang Wang, et al. Glm-5: from vibe coding to agentic engineering. *arXiv preprint arXiv:2602.15763*, 2026. 4, 12
- [83] Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. Wildchat: 1m chatgpt interaction logs in the wild. *arXiv preprint arXiv:2405.01470*, 2024. 8
- [84] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric. P Xing, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. Lmsys-chat-1m: A large-scale real-world llm conversation dataset, 2023. 8
- [85] Zihan Zheng, Zerui Cheng, Zeyu Shen, Shang Zhou, Kaiyuan Liu, Hansen He, Dongruixuan Li, Stanley Wei, Hangyi Hao, Jianzhu Yao, et al. Livecodebench pro: How do olympiad medalists judge llms in competitive programming? *arXiv preprint arXiv:2506.11928*, 2025. 18, 21, 27
- [86] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models, 2023. URL <https://arxiv.org/abs/2311.07911>. 22