

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 VAR-MATH: PROBING TRUE MATHEMATICAL REAS- SONING IN LLMs VIA SYMBOLIC MULTI-INSTANCE BENCHMARKS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Recent advances in reinforcement learning (RL) have led to substantial improvements in the mathematical reasoning abilities of large language models (LLMs), as measured by standard benchmarks. Yet these gains often persist even when models are trained with flawed signals, such as random or inverted rewards. This raises a fundamental question: do such improvements reflect genuine reasoning, or are they merely artifacts of overfitting to benchmark-specific patterns? To answer this question, we adopt an evaluation-centric perspective and highlight two critical shortcomings in existing protocols. First, *benchmark contamination* arises because test problems are publicly available, thereby increasing the risk of data leakage. Second, *evaluation fragility* results from reliance on single-instance assessments, which are sensitive to stochastic outputs and fail to capture reasoning consistency. These limitations suggest the need for a new evaluation paradigm that can probe reasoning ability beyond memorization and one-off success. As response, we propose **VAR-MATH**, a symbolic evaluation framework that converts fixed numerical problems into parameterized templates and requires models to solve multiple instantiations of each. This design enforces consistency across structurally equivalent variants, mitigates contamination, and enhances robustness through bootstrapped metrics. We apply VAR-MATH to transform three popular benchmarks, AMC23, AIME24, and AIME25, into their symbolic counterparts, VAR-AMC23, VAR-AIME24, and VAR-AIME25. Experimental results show substantial performance drops for RL-trained models on these variabilized benchmarks, especially for smaller models, with average declines of 47.9% on AMC23, 58.8% on AIME24, and 72.9% on AIME25. These findings indicate that some existing RL methods rely on superficial heuristics and fail to generalize beyond specific numerical forms.

## 1 INTRODUCTION

Recent advances in large language models (LLMs) have led to remarkable improvements in mathematical reasoning tasks. Models such as OpenAI-01 (OpenAI, 2024), DeepSeek-R1 (Guo et al., 2025), and Kimi-k1.5 (Team et al., 2025) have achieved state-of-the-art results across a range of public benchmarks. A key contributor to this progress is the growing shift from conventional supervised fine-tuning (SFT) to reinforcement learning (RL), which has become a dominant strategy for aligning model outputs with desired reasoning behaviors. The impressive performance of models like DeepSeek-R1 has sparked a surge of research, which generally follows two directions. One focuses on improving data quality through filtering, deduplication, and verification pipelines (Meng et al., 2023; He et al., 2025b; Hu et al., 2025; Albalak et al., 2025). The other centers on refining RL algorithms themselves, including optimizations to PPO (Yuan et al., 2025b;a), extensions to GRPO variants (Yu et al., 2025; Liu et al., 2025; Zhang et al., 2025), entropy-regularized methods for exploration (Cui et al., 2025b; Yao et al., 2025; Wang et al., 2025), and alternative paradigms such as REINFORCE++ (Hu, 2025).

However, alongside this progress, a growing body of evidence has raised concerns about what these gains truly represent. Recent studies have shown that models trained with flawed or even adversarial reward signals can still achieve surprisingly strong results on standard mathematical bench-

marks (Shao et al., 2025). For example, rewards based purely on output format (e.g., the presence of expressions) can lead to improved scores regardless of correctness, and even models trained with random or inverted rewards have demonstrated non-trivial performance gains. These counterintuitive findings converge on a fundamental question: *Are RL-trained LLMs genuinely learning to reason, or are they merely exploiting superficial patterns embedded in benchmark datasets?* If benchmark success can be achieved without correctness, then current evaluation protocols may not be measuring true reasoning ability, which in turn calls into question the validity of benchmark-driven progress and highlights the need to reconsider what existing metrics actually assess.

At the core of this issue lies a structural limitation in how benchmarks are constructed. Most mathematical reasoning benchmarks present each problem as a single, fixed numerical instance. While this simplifies evaluation, it introduces two critical vulnerabilities. First, *benchmark contamination* is increasingly unavoidable. Many widely used datasets, such as AMC23 and AIME24&25, are sourced from public math competitions. Given the breadth of pretraining corpora, it is highly likely that some problems (or closely related variants) have appeared in training data, thereby confounding evaluations with memorization effects. Second, *evaluation instability* follows from the reliance on single-instance assessments. Since many competition-style math problems yield simple numeric answers (e.g., 0 or 1), models can often succeed through statistical priors, guesswork, or shallow heuristics rather than genuine reasoning. As a result, it becomes difficult to distinguish true problem-solving ability from superficial pattern exploitation.

To address these limitations, we propose **VAR-MATH**, a symbolic evaluation framework that probes true reasoning ability through *multi-instance verification*. As illustrated in Figure 1, the central idea is intuitive: *If a model genuinely understands a problem, it should solve not just one instance, but multiple variants that differ only in surface-level values while sharing the same underlying structure.* Concretely, VAR-MATH transforms fixed problems into symbolic templates by replacing constants with constrained variables. For example, the original question

“Calculate the area defined by  $||x| - 1| + ||y| - 1| \leq 1$ ”

can be generalized into

“Calculate the area defined by  $||x| - a| + ||y| - a| \leq a$ ”,

where  $a$  is sampled from a feasible domain. This symbolic multi-instantiation strategy shifts evaluation from one-shot correctness to structural consistency, thereby mitigating contamination, suppressing heuristic shortcuts, and enabling more faithful assessment of generalizable mathematical reasoning.

To quantify performance, VAR-MATH reports two complementary metrics. A *loose* score computes the average accuracy across sampled variants, while a *strict* score grants credit only if all variants of a problem are solved correctly. In addition, a bootstrapping procedure further stabilizes evaluation by reducing variance and yielding more reliable estimates.

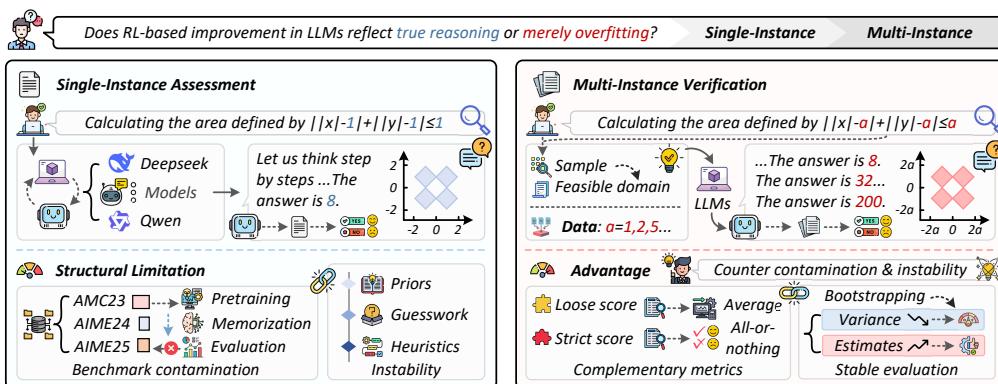


Figure 1: Multi-Instance Verification (VAR-MATH) vs. Single-Instance Assessment

108 Building on this protocol, we apply VAR-MATH to three widely used mathematical benchmarks,  
 109 AMC23 (MAA, 2023), AIME24, and AIME25 (MAA), generating their symbolic counterparts  
 110 VAR-AMC23, VAR-AIME24, and VAR-AIME25. When RL-finetuned models are re-evaluated on  
 111 these transformed benchmarks, their performance drops sharply. For instance, several 7B-parameter  
 112 models that previously achieved scores ranging from **36.9** to **78.6** on AMC23 drop to a range of **2.0**  
 113 to **57.0** on VAR-AMC23, with similar declines on VAR-AIME24 and VAR-AIME25.

114 To summarize, our contributions are:  
 115

- 116 1. We propose VAR-MATH, a symbolic evaluation framework that systematically variabilizes  
 117 three widely used benchmarks (AMC23, AIME24, and AIME25), enabling contamination-  
 118 robust and consistency-based assessment of mathematical reasoning.
- 119 2. We establish a principled evaluation protocol that combines *loose* and *strict* consistency  
 120 metrics with a bootstrapping procedure, ensuring statistically reliable comparison across  
 121 models.
- 122 3. We conduct an extensive empirical study on VAR-AMC23, VAR-AIME24, and VAR-  
 123 AIME25, revealing substantial performance declines in RL-finetuned models and exposing  
 124 the limitations of current RL strategies in cultivating genuine reasoning.

## 126 2 RELATED WORK

128 A wide range of benchmarks has been developed to evaluate the mathematical reasoning capabilities  
 129 of LLMs, spanning diverse difficulty levels, problem formats, and contamination risks. Existing  
 130 efforts can be broadly categorized into static and dynamic benchmarks.

132 **Static Benchmarks** The GSM8K dataset (Cobbe et al., 2021) contains  $8.5K$  grade-school math  
 133 word problems ( $7.5K$  train,  $1K$  test) targeting multi-step arithmetic reasoning. While founda-  
 134 tional, its limited numerical complexity reduces its utility for diagnosing advanced reasoning.  
 135 MATH500 (Hendrycks et al., 2021) offers 500 high-school-level problems covering algebra and  
 136 calculus. OlympiadBench (He et al., 2024) includes 8476 Olympiad-level problems from sources  
 137 such as the International Mathematical Olympiad and China’s Gaokao, featuring multimodal inputs  
 138 (e.g., diagrams) and step-by-step expert solutions for fine-grained evaluation in bilingual settings.  
 139 AMC23 (MAA, 2023) collects problems from the 2023 American Mathematics Competition, em-  
 140 phasizing functional equations and complex analysis; each requires an integer answer between 0  
 141 and 999. Because of its small size and public availability, repeated sampling is necessary to re-  
 142 duce variance, while contamination remains a concern. The AIME series (MAA) is drawn from  
 143 the American Invitational Mathematics Examination, with AIME24 containing 2024 contest prob-  
 144 lems and AIME25 adding novel problems curated in 2025. These increasingly challenging tasks  
 145 demand deeper combinatorial and geometric reasoning, yet their public accessibility leaves them  
 146 highly vulnerable to contamination as LLMs approach benchmark saturation.

147 **Dynamic Benchmarks** To mitigate contamination risks, recent work has shifted toward dynamic  
 148 evaluation (Chen et al., 2025). For instance, Srivastava et al. (2024) alleviates contamination by  
 149 creating functional variations of the MATH dataset, where new problems are generated by modi-  
 150 fying numeric parameters to yield distinct solutions. Similarly, Mirzadeh et al. (2024) introduces  
 151 an enhanced benchmark that generates diverse variants of GSM8K, while Gulati et al. (2024) alters  
 152 variables, constants, and phrasing in Putnam competition problems. There are also several success-  
 153 ful works that use symbolic variants (Xu et al., 2025; Li et al., 2024; Gao et al., 2022; Shi et al.,  
 154 2023). LiveBench (White et al., 2024) further advances this direction by sourcing fresh problems  
 155 monthly from arXiv papers, news, and contests, with rigorous contamination controls. However,  
 156 maintaining such a benchmark requires ongoing human curation, limiting scalability and introduc-  
 157 ing subjectivity.

158 Overall, these studies represent important progress and highlight the value of functional variation  
 159 in constructing dynamic benchmarks. Nevertheless, most existing efforts either concentrate on ele-  
 160 mentary problems (e.g., GSM8K, MATH) or remain restricted to Olympiad-style datasets. In con-  
 161 trast, our work focuses on advanced reasoning benchmarks that are widely adopted in RL-finetuning  
 evaluations, including AMC23, AIME24, and AIME25. Building on this foundation, we introduce

162 VAR-MATH, a symbolic evaluation framework that not only variabilizes these benchmarks but also  
 163 incorporates consistency-based metrics and bootstrapped stability analysis. This enables a more rig-  
 164 orous and reliable assessment of model reasoning robustness while further reducing susceptibility  
 165 to data contamination.

166

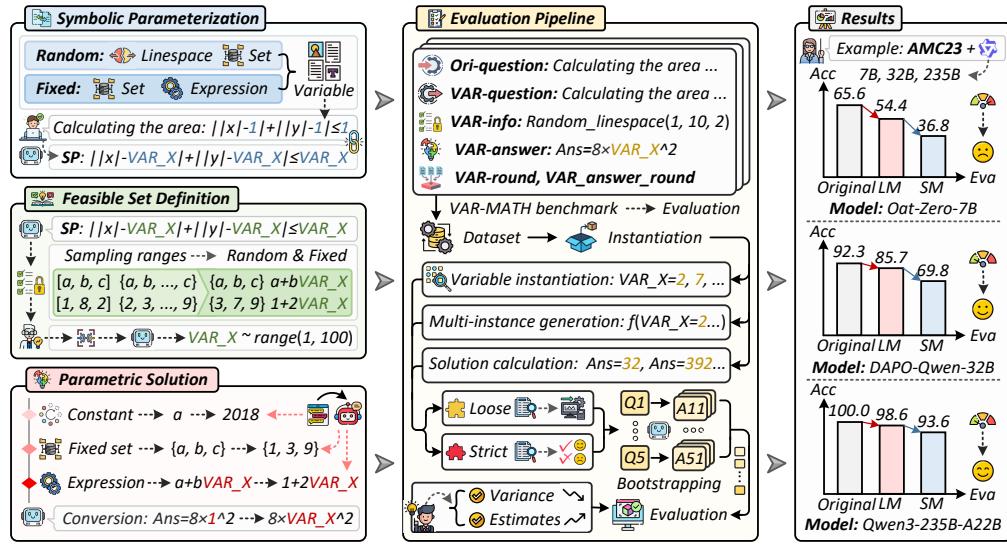
167

### 168 3 VAR-MATH

169

170 As shown in Figure 2, VAR-MATH consists of three core components: its design principles, the data  
 171 transformation process, and the evaluation protocol. We introduce each in turn below.

172



191

192

193

194

195

196

197

Figure 2: Overview of the VAR-MATH pipeline. The process consists of two stages: *preprocessing*, where original problems are symbolically abstracted by replacing constants with variables and defining feasible sampling ranges, and *evaluation*, where problems are instantiated into multiple concrete variants and assessed using loose (LM) and strict (SM) consistency metrics.

198

199

The core motivation behind VAR-MATH is to address two long-standing limitations in the evalua-  
 200 tion of mathematical reasoning: *benchmark contamination* and *evaluation fragility*. Traditional  
 201 benchmarks typically present problems as static numerical instances with fixed values, making them  
 202 vulnerable to memorization and shallow pattern exploitation. In such settings, models may succeed  
 203 by retrieving known solutions or leveraging statistical priors rather than performing genuine rea-  
 204 soning. These issues call for an evaluation paradigm that can separate true reasoning ability from  
 205 superficial success.

206

207

208

209

210

VAR-MATH introduces such a paradigm through a process we call *symbolic variabilization*, which  
 decouples problem structure from fixed numeric content. Instead of hardcoding specific constants,  
 problems are restructured into symbolic templates, where concrete values are dynamically instan-  
 tiated during evaluation. This abstraction allows models to be tested not on isolated instances, but  
 across families of structurally equivalent problems.

211

212

213

214

215

The key assumption is that a model that truly understands a mathematical problem should demon-  
 strate *reasoning consistency*, i.e., the ability to solve multiple variants of the same logical structure  
 regardless of specific numerical values. By systematically sampling from constrained parameter  
 spaces, VAR-MATH preserves the original semantics of each problem while introducing controlled  
 variation. This results in a more robust and contamination-resistant evaluation protocol, which is  
 capable of distinguishing genuine understanding from surface-level heuristics.

216 3.2 DATA PROCESSING  
217218 Building on the principle of symbolic variabilization, the data transformation pipeline systematically  
219 converts problems from established mathematical benchmarks into variabilized form. We focus on  
220 AMC23 and AIME24&25, which represent two distinct tiers of competition-level difficulty. Each  
221 selected problem undergoes symbolic abstraction through a structured four-step methodology:222 

- 223 • **Structural analysis.** Each problem is first solved independently by a mathematics expert.  
224 The expert works through the full reasoning process and cross-checks it against the official  
225 golden solution. This step identifies the algebraic structure of the problem, determines  
226 which quantities are essential constants, and isolates the core symbolic variables that drive  
227 the solution. We deliberately preserve the ratios or relationships that appear in the original  
228 derivation to maintain semantic fidelity.
- 229 • **Symbolic parameterization.** Key numerical constants are then replaced with symbolic  
230 variables. Feasible domains for each variable are chosen to stay close to the scale of the  
231 original values (e.g., an original value  $x = 5$  may become a variable ranging from 2 to 8),  
232 while ensuring that the mathematical meaning of the problem remains valid. In construct-  
233 ing these domains, we apply domain restrictions from the derivation (such as positivity,  
234 non-vanishing denominators, or geometric constraints). Both continuous ranges and dis-  
235 crete sets are supported, as summarized in Table 1.
- 236 • **Parametric solution formulation and verification.** The final answer is expressed as a  
237 symbolic function of the newly defined variables. This symbolic formula is derived man-  
238 ually by the expert and then verified in two stages: (i) *human verification*: another annotator  
239 solves each instantiated variant directly from the rewritten prompt to ensure correctness;  
240 and (ii) *model verification*: we run all variants through a frontier model (DeepSeek) as  
241 an additional sanity check. A problem is accepted only if all of its variants are correctly  
242 solved during this step; otherwise, it will go through a second round of verification by  
243 another expert.
- 244 • **Variant sampling and evaluation protocol.** For each symbolic template, we uniformly  
245 sample values from the predefined feasible domains to generate a fixed set of up to  $K = 5$   
246 concrete variants (with  $K = 2 \sim 4$  for a few problems to preserve difficulty alignment). All  
247 models are evaluated on *exactly* the same variants for a given problem, ensuring strict com-  
248 parability across models. For each sampled variant, the ground-truth answer is computed  
249 directly from the parametric solution, and all instantiated variants are evaluated using a  
250 standardized prompting strategy consistent with prior mathematical reasoning benchmarks.  
251 This results in approximately 430 instantiated questions across the entire benchmark. An  
252 example is provided in Appendix G.
- 253 • **Precision specification.** To ensure numerical stability, we apply consistent rounding rules  
254 and significant-digit constraints to both the instantiated variables and the computed an-  
255 swers.

256 In certain cases, special constants integral to the mathematical identity of a problem (e.g.,  $\pi$ ,  $e$ ,  
257 or fixed geometric parameters) are preserved without modification to maintain fidelity. The output  
258 of this pipeline is a set of variabilized benchmarks, namely **VAR-AMC23**, **VAR-AIME24**, and  
259 **VAR-AIME25**. Each problem is encoded as a structured object containing a symbolic expression,  
260 variable definitions with feasible sets, parametric answers, and metadata specifying its origin and  
261 difficulty. This unified representation enables efficient multi-instance instantiation and facilitates  
262 future benchmark extension and automation. Other details are provided in the Appendix.263 4 EXPERIMENTS  
264265 4.1 EXPERIMENTAL SETUP  
266267 We evaluate model performance on six benchmarks: the original **AMC23**, **AIME24**, and **AIME25**  
268 datasets, together with their variabilized counterparts **VAR-AMC23**, **VAR-AIME24**, and **VAR-**  
269 **AIME25** generated by the VAR-MATH framework.

270

271

Table 1: Variable and Answer Expression Formats

272 Variable Type (VAR_X)	273 Description
274 Random.linspace_[a, b, c]	275 Sampled from a linear space between $a$ and $b$ with $c$ intervals
275 Random_Set_{a, b, ..., c}	276 Sampled uniformly from the given discrete set
276 Fixed_Set_{a, b, c}	277 Must take one of the fixed values in the set
277 Expression_a · VAR_Y + b	278 Defined algebraically based on other variables
278 Answer Type	279 Description
280 Constant a	281 Answer is a constant value independent of input
281 Fixed_Set_{a, b, c}	282 Answer selected based on a fixed variable-to-output mapping
282 Expression_a · VAR_Y + b	283 Answer computed as a function of variable(s)

282

283

**7B-parameter and 32B-parameter models.** We benchmark a collection of open-source 7B models, including the base Qwen2.5-MATH-7B (Yang et al., 2024), and several RL-enhanced variants: Eurus-2-7B-PRIME (Cui et al., 2025a), Skywork-OR1-Math-7B (He et al., 2025a), SimpleRL-Zoo-7B (Zeng et al., 2025), Light-R1-7B-DS (Wen et al., 2025), and Oat-Zero-7B (Liu et al., 2025). These models cover a range of RL training pipelines and policy optimization techniques. We further evaluate three 32B-scale models: the base Qwen2.5-32B (Team, 2024), and two RL-finetuned variants, DAPO-Qwen-32B (Yu et al., 2025) and SRPO-Qwen-32B (Zhang et al., 2025), both trained with large-scale reinforcement learning systems. Our evaluation pipeline is based on the open-source Qwen2.5-MATH repository<sup>1</sup>, and employs **vLLM** (Kwon et al., 2023) for efficient decoding. All models are tested under consistent hardware and inference configurations on NVIDIA A6000 GPUs with `bfloat16` precision. Generation parameters are fixed at `temperature = 0.6` and `top-p = 1.0`, while batch sizes are adjusted for each model to maximize throughput without affecting reproducibility.

294

295

**High-Capacity Models** We further include several high-capacity state-of-the-art models, including DeepSeek-R1 (Guo et al., 2025), SEED-THINK (Seed et al., 2025), Qwen3-235B-A22B (Yang et al., 2025), and OpenAI-o4-mini-high (OpenAI, 2024). Evaluation for these models is conducted using a single inference pass per problem with default sampling configurations.

300

301

## 4.2 EVALUATION METRICS

302

We evaluate model performance using two complementary metrics: *loose* and *strict*. For each symbolic problem, up to five instantiated variants are generated by sampling values from the feasible domains of its parameters. Each variant is queried  $M$  times, producing  $M$  independent responses. The loose metric measures average correctness across all variants of a symbolic problem, while the strict metric enforces reasoning consistency: a problem is marked correct only if all of its variants are solved correctly. This all-or-nothing criterion emphasizes consistency across structurally equivalent problems rather than success on isolated instances.

309

310

To reduce variance and obtain statistically reliable estimates, we apply a bootstrap procedure with  $N$  resampling rounds. For a given symbolic problem  $Q$  with  $K$  variants  $\{Q_1, \dots, Q_K\}$ , and corresponding responses  $\{A_{kj}\}$  for variant  $Q_k$  under sample index  $j = 1, \dots, M$ , the dataset is

313

314

$$D = \{(Q_k, A_{kj}) \mid k = 1, \dots, K, j = 1, \dots, M\}.$$

315

316

In each bootstrap round  $i = 1, \dots, N$ , one response  $\hat{A}_{ki}$  is drawn uniformly from  $\{A_{kj}\}_{j=1}^M$  for every variant  $Q_k$ . The set  $\{\hat{A}_{ki}\}_{k=1}^K$  is then used to compute both loose and strict scores:

317

318

319

$$\text{score}_{\text{loose}} = \frac{1}{N} \sum_{i=1}^N \left( \frac{1}{K} \sum_{k=1}^K \mathbf{1}[\hat{A}_{ki} = \text{gt}_k] \right), \quad (1)$$

320

321

322

$$\text{score}_{\text{strict}} = \frac{1}{N} \sum_{i=1}^N \left( \prod_{k=1}^K \mathbf{1}[\hat{A}_{ki} = \text{gt}_k] \right), \quad (2)$$

323

<sup>1</sup><https://github.com/QwenLM/Qwen2.5-Math>

324 where  $gt_k$  denotes the ground-truth answer of  $Q_k$ . Final performance is reported as the mean of  
 325 these bootstrap estimates, with standard deviations serving as a measure of statistical stability. An  
 326 illustration of this procedure is provided in Appendix A.

328 **4.3 MAIN RESULTS**  
 329

331 **Table 2: Evaluation Results on AMC23 and VAR-AMC23.**

332 <b>Model</b>	333 <b>AMC23</b>	334 <b>(strict) VAR- AMC23</b>	335 <b>Drop</b>	336 <b>(loose) VAR- AMC23</b>	337 <b>Drop</b>
338 Qwen2.5-MATH-7B	339 36.9 (6.3)	340 2.0 (2.0)	341 <b>-94.5%</b>	342 22.7 (2.6)	343 <b>-38.5%</b>
344 Eurus-2-7B-PRIME	345 58.3 (4.3)	346 28.9 (3.7)	347 <b>-50.4%</b>	348 49.9 (2.5)	349 <b>-14.3%</b>
350 Skywork-OR1-Math-7B	351 73.9 (5.4)	352 57.0 (3.6)	353 <b>-22.9%</b>	354 72.0 (2.3)	355 <b>-2.6%</b>
356 SimpleRL-Zoo-7B	357 61.4 (4.8)	358 33.6 (4.0)	359 <b>-45.3%</b>	360 52.2 (2.3)	361 <b>-15.0%</b>
363 Light-R1-7B-DS	364 78.6 (6.3)	365 54.9 (4.6)	366 <b>-30.2%</b>	367 75.8 (2.3)	368 <b>-3.5%</b>
370 Oat-Zero-7B	371 65.6 (3.1)	372 36.8 (3.3)	373 <b>-43.9%</b>	374 54.4 (2.4)	375 <b>-17.0%</b>
378 Qwen2.5-32B	379 33.4 (4.5)	380 3.1 (2.5)	381 <b>-90.6%</b>	382 27.4 (2.8)	383 <b>-18.2%</b>
386 DAPO-Qwen-32B	387 92.3 (2.9)	388 69.8 (3.1)	389 <b>-24.4%</b>	390 85.7 (1.4)	391 <b>-7.2%</b>
395 SRPO-Qwen-32B	396 86.7 (3.7)	397 51.5 (4.5)	398 <b>-40.6%</b>	399 73.9 (2.6)	400 <b>-14.8%</b>
404 DeepSeek-R1-0528	405 100.0 (0.0)	406 96.4 (2.5)	407 <b>-3.6%</b>	408 99.2 (0.5)	409 <b>-0.8%</b>
413 Qwen3-235B-A22B	414 100.0 (0.0)	415 93.6 (3.1)	416 <b>-6.4%</b>	417 98.6 (0.7)	418 <b>-1.4%</b>
422 SEED-THINK-v1.6	423 100.0 (0.0)	424 98.8 (1.5)	425 <b>-1.2%</b>	426 99.8 (0.3)	427 <b>-0.2%</b>
431 OpenAI-o4-mini-high	432 100.0 (0.0)	433 93.4 (2.3)	434 <b>-6.6%</b>	435 98.2 (0.7)	436 <b>-1.8%</b>

353  
 354 **4.3.1 RESULTS ANALYSIS ON THE STRICT METRIC**  
 355

356 In this section, we focus on the strict metric, which emphasizes reasoning consistency across struc-  
 357 turally equivalent problem variants, and we recommend it as the primary evaluation measure.

358 **RL-tuned 7B models show fragile generalization.** Across all benchmarks, RL-optimized  
 359 7B models experience sharp drops in accuracy once problems are variabilized. For example,  
 360 Light-R1-7B-DS falls from 78.6 to 54.9 on AMC23, from 40.8 to 23.8 on AIME24, and from  
 361 32.7 to 17.1 on AIME25. Similar declines occur for Eurus-2-7B-PRIME and Oat-Zero-7B.  
 362 These results point to two issues: overfitting to specific numeric templates, possibly amplified by  
 363 contamination from public problem sets, and a lack of symbolic consistency, where solving one in-  
 364 stance does not transfer reliably to others with altered values. Such weaknesses remain hidden under  
 365 conventional single-instance benchmarks but are revealed by VAR-MATH.

366 **Scaling to 32B improves accuracy but inconsistency persists.** Larger 32B models achieve higher  
 367 raw accuracy, e.g., DAPO-Qwen-32B and SRPO-Qwen-32B exceed 85 on AMC23. Nevertheless,  
 368 they still suffer relative drops of more than 40% across the variabilized datasets, showing that scaling  
 369 enhances memorization and structural recognition but does not fully resolve the problem of symbolic  
 370 consistency.

371 **Frontier models are more robust yet still challenged by symbolic variation.** State-of-the-art  
 372 models such as DeepSeek-R1 and SEED-THINK maintain strong performance on AMC23, with  
 373 drops below 5%. This robustness likely stems from high-quality training data and sophisticated  
 374 alignment pipelines that mitigate shortcut learning. However, even these frontier systems experi-  
 375 ence notable degradation on the more difficult AIME24 and AIME25 variants, with relative drops  
 376 up to 28.1%. These findings indicate that symbolic variation remains a fundamental challenge, un-  
 377 derscoring the importance of evaluation protocols that move beyond surface-level accuracy toward  
 consistency-based reasoning assessment.

378 The score drop mentioned above primarily results from data contamination and evaluation fragility.  
 379 In the following text, we provide an in-depth analysis of data contamination and demonstrate how  
 380 VAR-Math enhances assessment stability.  
 381  
 382

383 Table 3: Evaluation Results on AIME24 and VAR-AIME24.

384 <b>Model</b>	385 <b>AIME24</b>	386 <b>(strict) VAR- 387 AIME24</b>	388 <b>Drop</b>	389 <b>(loose) VAR- 390 AIME24</b>	391 <b>Drop</b>
392 Qwen2.5-MATH-7B	393 10.8 (4.5)	394 3.2 (2.7)	395 <b>-70.0%</b>	396 7.9 (2.9)	397 <b>-27.1%</b>
398 Eurus-2-7B-PRIME	399 15.8 (4.8)	400 4.3 (2.9)	401 <b>-72.5%</b>	402 13.4 (2.7)	403 <b>-15.5%</b>
404 Skywork-OR1-Math-7B	405 41.5 (4.2)	406 23.9 (4.3)	407 <b>-42.3%</b>	408 39.0 (3.4)	409 <b>-6.0%</b>
410 SimpleRL-Zoo-7B	411 23.8 (5.9)	412 8.5 (3.7)	413 <b>-64.1%</b>	414 20.4 (3.5)	415 <b>-14.1%</b>
416 Light-R1-7B-DS	417 40.8 (5.1)	418 23.8 (4.8)	419 <b>-41.7%</b>	420 40.6 (3.3)	421 <b>-0.6%</b>
422 Oat-Zero-7B	423 34.0 (2.1)	424 12.8 (3.6)	425 <b>-62.3%</b>	426 22.3 (2.5)	427 <b>-34.3%</b>
428 Qwen2.5-32B	429 8.8 (4.4)	430 2.3 (2.3)	431 <b>-73.5%</b>	432 7.9 (2.6)	433 <b>-9.6%</b>
434 DAPO-Qwen-32B	435 51.7 (6.6)	436 29.8 (4.6)	437 <b>-42.4%</b>	438 50.9 (2.8)	439 <b>-1.5%</b>
440 SRPO-Qwen-32B	441 55.6 (5.0)	442 29.2 (4.2)	443 <b>-47.6%</b>	444 46.9 (2.9)	445 <b>-15.7%</b>
446 DeepSeek-R1-0528	447 86.8 (3.3)	448 73.7 (3.8)	449 <b>-15.1%</b>	450 82.3 (2.6)	451 <b>-5.1%</b>
454 Qwen3-235B-A22B	455 84.1 (3.2)	456 69.5 (3.4)	457 <b>-17.4%</b>	458 80.1 (1.9)	459 <b>-4.9%</b>
460 SEED-THINK-v1.6	461 87.5 (3.3)	462 73.4 (3.5)	463 <b>-16.1%</b>	464 82.7 (2.4)	465 <b>-5.5%</b>
466 OpenAI-o4-mini-high	467 91.8 (2.9)	468 78.1 (3.7)	469 <b>-14.9%</b>	470 89.0 (1.9)	471 <b>-2.7%</b>

## 405 4.3.2 DECOUPLING THE IMPACT OF DATA CONTAMINATION

406 To better diagnose the sources of performance degradation, we analyze results under the *loose metric*. Unlike the strict all-or-nothing criterion, this softer metric grants partial credit for solving sub-  
 407 sets of variants, thereby helping disentangle contamination-driven memorization from instability in  
 408 symbolic reasoning.

409 Results on AMC23 suggest that contamination exerts a substantial influence, especially on the  
 410 base models. For example, the base model Qwen2.5-MATH-7B shows a 38.5% decline, con-  
 411 sistent with heavy reliance on memorized patterns rather than generalizable reasoning. By contrast,  
 412 Skywork-OR1-Math-7B and DAPO-Qwen-32B record much smaller drops (2.6% and 7.2%,  
 413 respectively), indicating greater resistance to contamination and stronger abstraction of underlying  
 414 structures.

415 On the more challenging AIME24 and AIME25 benchmarks, degradation is more heterogeneous.  
 416 Models such as SRPO-Qwen-32B exhibit relatively mild drops (e.g., 4.6% on AIME25), sug-  
 417 gesting improved robustness across symbolic variants. Others, including Qwen2.5-32B and  
 418 DAPO-Qwen-32B, suffer sharp declines (21.2% and 13.5% on AIME25), reflecting persistent  
 419 fragility when faced with minor symbolic perturbations.

420 Together with the strict-metric results in Section 4.3.1, these findings point to two intertwined fac-  
 421 tors underlying symbolic degradation: benchmark-specific overfitting amplified by contamination,  
 422 and instability in applying reasoning consistently across variants. While RL can improve scores on  
 423 conventional benchmarks, it also risks reinforcing memorization and narrow heuristics. This under-  
 424 scores the necessity of evaluation frameworks that are both contamination-resistant and sensitive to  
 425 reasoning stability.

## 426 4.3.3 ENHANCING EVALUATION STABILITY VIA VAR-MATH

427 In Appendix B, Figure 4 reports the distribution of standard deviations of the scores for 7B and  
 428 32B models on both the original and variabilized benchmarks. The result shows that VAR-MATH

432

433

Table 4: Evaluation Results on AIME25 and VAR-AIME25.

434

435

Model	AIME25	(strict) VAR-AIME25	Drop	(loose) VAR-AIME25	Drop
Qwen2.5-MATH-7B	4.8 (3.1)	0.0 (0.0)	-100.0%	3.2 (1.3)	-34.2%
Eurus-2-7B-PRIME	10.0 (3.1)	1.2 (1.7)	-87.8%	7.4 (1.4)	-26.0%
Skywork-OR1-Math-7B	24.0 (3.8)	15.0 (2.5)	-37.3%	23.4 (1.6)	-2.4%
SimpleRL-Zoo-7B	12.5 (3.4)	2.9 (2.4)	-76.9%	11.5 (1.6)	-7.9%
Light-R1-7B-DS	32.7 (3.8)	17.1 (3.2)	-47.7%	30.3 (1.8)	-7.3%
Oat-Zero-7B	9.2 (3.2)	1.2 (1.8)	-87.4%	8.4 (1.4)	-7.9%
Qwen2.5-32B	3.5 (3.4)	0.0 (0.0)	-100.0%	2.8 (1.2)	-21.2%
DAPO-Qwen-32B	37.3 (5.3)	21.2 (2.4)	-43.2%	32.2 (2.3)	-13.5%
SRPO-Qwen-32B	26.5 (5.2)	14.5 (3.0)	-45.2%	25.2 (1.7)	-4.6%
DeepSeek-R1-0528	81.5 (3.3)	61.3 (4.4)	-24.8%	75.2 (2.5)	-7.9%
Qwen3-235B-A22B	82.6 (3.1)	61.6 (4.6)	-25.4%	75.7 (2.0)	-8.1%
SEED-THINK-v1.6	81.7 (3.0)	58.8 (4.0)	-28.1%	75.3 (2.6)	-7.7%
OpenAI-o4-mini-high	93.4 (2.4)	76.7 (3.5)	-17.8%	87.0 (1.9)	-6.8%

454

455

456 consistently reduces output variance, with the effect most evident on the more challenging AIME25  
 457 benchmark, where conventional single-instance evaluation is highly susceptible to sampling noise.  
 458

459 This improvement derives from VAR-MATH’s core design. By instantiating each symbolic prob-  
 460 lem multiple times and aggregating performance across variants, the framework dampens stochastic  
 461 artifacts and outlier completions. Such ensemble-style averaging yields a more faithful estimate  
 462 of reasoning ability and, with the incorporation of bootstrap methods to further stabilize estimates,  
 463 provides a stable, interpretable signal of a model’s true mathematical competence.  
 464

465

466

## 5 CONCLUSION

467

468 We introduced VAR-MATH, a systematic framework for evaluating mathematical reasoning in large  
 469 language models. Targeting widely used benchmarks (AMC23, AIME24, and AIME25), VAR-  
 470 MATH converts fixed problems into parameterized, multi-instance variants, enabling evaluation that  
 471 is both contamination-resistant and consistency-based. To measure performance, we proposed com-  
 472 plementary *loose* and *strict* metrics together with a bootstrap resampling procedure for stable statis-  
 473 tical estimation.

474 Empirical results demonstrate that many RL-finetuned models, despite strong performance on con-  
 475 ventional benchmarks, exhibit substantial degradation under VAR-MATH, exposing their reliance  
 476 on dataset-specific artifacts and their limited generalization ability. These findings underscore the  
 477 importance of principled dataset design and evaluation methodology in assessing reasoning com-  
 478 petence. By enforcing consistency across variants and stabilizing measurement, VAR-MATH provides  
 479 a more reliable indicator of genuine reasoning.

480 While our study centers on AMC23 and AIME24&25, the core methodology of VAR-MATH is  
 481 broadly applicable. Future work includes extending this framework to richer mathematical domains  
 482 and other reasoning-intensive tasks, such as program synthesis, formal logic, and decision-making.  
 483 Such extensions hold promise for establishing more rigorous and generalizable evaluation standards.  
 484 [Another direction is to include strong SFT-only baselines \(e.g., OpenThinker-3 Guha et al. \(2025\)\)](#)  
 485 side-by-side to better quantify how much of the observed drop is due to RL-specific behavior versus  
 general SFT math-tuning.

486 ETHICS STATEMENT  
487488 We confirm that all authors of this submission have read and agree to abide by the ICLR Code of  
489 Ethics.  
490491 REPRODUCIBILITY STATEMENT  
492493 The dataset and code are provided in the supplementary material to enable the reproduction of our  
494 results.  
495496 REFERENCES  
497498 Alon Albalak, Duy Phung, Nathan Lile, Rafael Rafailov, Kanishk Gandhi, Louis Castricato, Anikait  
499 Singh, Chase Blagden, Violet Xiang, Dakota Mahan, et al. Big-math: A large-scale, high-quality  
500 math dataset for reinforcement learning in language models. *arXiv preprint arXiv:2502.17387*,  
501 2025.502 Simin Chen, Yiming Chen, Zexin Li, Yifan Jiang, Zhongwei Wan, Yixin He, Dezhi Ran, Tianle Gu,  
503 Haizhou Li, Tao Xie, et al. Recent advances in large langauge model benchmarks against data  
504 contamination: From static to dynamic evaluation. *arXiv preprint arXiv:2502.17521*, 2025.505 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
506 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
507 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,  
508 2021.509 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu  
510 Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint*  
511 *arXiv:2502.01456*, 2025a.512 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen  
513 Fan, Huayu Chen, Weize Chen, et al. The entropy mechanism of reinforcement learning for  
514 reasoning language models. *arXiv preprint arXiv:2505.22617*, 2025b.515 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and  
516 Graham Neubig. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*, 2022.517 Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna  
518 Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, Ashima Suvarna, Benjamin Feuer, Liangyu  
519 Chen, Zaid Khan, Eric Frankel, Sachin Grover, Caroline Choi, Niklas Muenmighoff, Shiye Su,  
520 Wanjia Zhao, John Yang, Shreyas Pimpalgaonkar, Kartik Sharma, Charlie Cheng-Jie Ji, Yichuan  
521 Deng, Sarah Pratt, Vivek Ramanujan, Jon Saad-Falcon, Jeffrey Li, Achal Dave, Alon Albalak,  
522 Kushal Arora, Blake Wulfe, Chinmay Hegde, Greg Durrett, Sewoong Oh, Mohit Bansal, Saadia  
523 Gabriel, Aditya Grover, Kai-Wei Chang, Vaishaal Shankar, Aaron Gokaslan, Mike A. Merrill,  
524 Tatsunori Hashimoto, Yejin Choi, Jenia Jitsev, Reinhard Heckel, Maheswaran Sathiamoorthy,  
525 Alexandros G. Dimakis, and Ludwig Schmidt. Openthoughts: Data recipes for reasoning models,  
526 2025. URL <https://arxiv.org/abs/2506.04178>.527 Aryan Gulati, Brando Miranda, Eric Chen, Emily Xia, Kai Fronsdal, Bruno de Moraes Dumont,  
528 and Sanmi Koyejo. Putnam-axiom: A functional & static benchmark for measuring higher level  
529 mathematical reasoning in llms. In *Forty-second International Conference on Machine Learning*,  
530 2024.531 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
532 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
533 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.534 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,  
535 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for  
536 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*  
537 *arXiv:2402.14008*, 2024.

540 Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang  
 541 Zhang, Jiacheng Xu, Wei Shen, et al. Skywork open reasoner 1 technical report. *arXiv preprint*  
 542 *arXiv:2505.22312*, 2025a.

543 Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian  
 544 Yu, Zhenwen Liang, Wenxuan Wang, et al. Deepmath-103k: A large-scale, challenging, de-  
 545 contaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv preprint*  
 546 *arXiv:2504.11456*, 2025b.

547 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
 548 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*  
 549 *preprint arXiv:2103.03874*, 2021.

550 Jian Hu. Reinforce++: A simple and efficient approach for aligning large language models. *arXiv*  
 551 *preprint arXiv:2501.03262*, 2025.

552 Jingcheng Hu, Yinmin Zhang, Qi Han, Dixin Jiang, Xiangyu Zhang, and Heung-Yeung Shum.  
 553 Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base  
 554 model, 2025. URL <https://arxiv.org/abs/2503.24290>.

555 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.  
 556 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model  
 557 serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating*  
 558 *Systems Principles*, 2023.

559 Zenan Li, Zhi Zhou, Yuan Yao, Xian Zhang, Yu-Feng Li, Chun Cao, Fan Yang, and Xiaoxing Ma.  
 560 Neuro-symbolic data generation for math reasoning. *Advances in Neural Information Processing*  
 561 *Systems*, 37:23488–23515, 2024.

562 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee,  
 563 and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint*  
 564 *arXiv:2503.20783*, 2025.

565 MAA. American invitational mathematics examination - aime. In *American Invitational*  
 566 *Mathematics Examination - AIME*. URL <https://maa.org/math-competitions/american-invitational-mathematics-examination-aime>.

567 MAA. American mathematics competitions. In *American Mathematics Competitions*, 2023.

568 Chunyang Meng, Shijie Song, Haogang Tong, Maolin Pan, and Yang Yu. Deepscaler: Holistic  
 569 autoscaling for microservices based on spatiotemporal gnn with adaptive graph learning. In *2023*  
 570 *38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pp. 53–65.  
 571 IEEE, 2023.

572 Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad  
 573 Farajtabar. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large  
 574 language models. *arXiv preprint arXiv:2410.05229*, 2024.

575 OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms>, 2024.

576 ByteDance Seed, Jiaze Chen, Tiantian Fan, Xin Liu, Lingjun Liu, Zhiqi Lin, Mingxuan Wang,  
 577 Chengyi Wang, Xiangpeng Wei, Wenyuan Xu, et al. Seed1. 5-thinking: Advancing superb rea-  
 578 soning models with reinforcement learning. *arXiv preprint arXiv:2504.13914*, 2025.

579 Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei  
 580 Du, Nathan Lambert, Sewon Min, Ranjay Krishna, et al. Spurious rewards: Rethinking training  
 581 signals in rlvr. *arXiv preprint arXiv:2506.10947*, 2025.

582 Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael  
 583 Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context.  
 584 In *International Conference on Machine Learning*, pp. 31210–31227. PMLR, 2023.

594 Saurabh Srivastava, Anto PV, Shashank Menon, Ajay Sukumar, Alan Philipose, Stevin Prince,  
 595 Sooraj Thomas, et al. Functional benchmarks for robust evaluation of reasoning performance,  
 596 and the reasoning gap. *arXiv preprint arXiv:2402.19450*, 2024.

597

598 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun  
 599 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1.5: Scaling reinforcement learning with  
 600 llms. *arXiv preprint arXiv:2501.12599*, 2025.

601 Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.

602

603 Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,  
 604 Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive  
 605 effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025.

606

607 Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu  
 608 Tang, Xiaowei Lv, et al. Light-r1: Curriculum sft, dpo and rl for long cot from scratch and beyond.  
 609 *arXiv preprint arXiv:2503.10460*, 2025.

610

611 Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-  
 612 Ziv, Neel Jain, Khalid Saifullah, Siddartha Naidu, et al. Livebench: A challenging, contamination-  
 613 free llm benchmark. *arXiv preprint arXiv:2406.19314*, 4, 2024.

614

615 Xinnuo Xu, Rachel Lawrence, Kshitij Dubey, Atharva Pandey, Risa Ueno, Fabian Falck, Aditya V  
 616 Nori, Rahul Sharma, Amit Sharma, and Javier Gonzalez. Re-imagine: Symbolic benchmark  
 617 synthesis for reasoning evaluation. *arXiv preprint arXiv:2506.15455*, 2025.

618

619 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu,  
 620 Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu,  
 621 Xingzhang Ren, and Zhenru Zhang. Qwen2.5-math technical report: Toward mathematical ex-  
 622 pert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.

623

624 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,  
 625 Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint  
 626 arXiv:2505.09388*, 2025.

627

628 Jian Yao, Ran Cheng, Xingyu Wu, Jibin Wu, and Kay Chen Tan. Diversity-aware policy optimization  
 629 for large language model reasoning. *arXiv preprint arXiv:2505.23433*, 2025.

630

631 Qiyi Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian  
 632 Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system  
 633 at scale. *arXiv preprint arXiv:2503.14476*, 2025.

634

635 Yufeng Yuan, Qiyi Yu, Xiaochen Zuo, Ruofei Zhu, Wenyuan Xu, Jiaze Chen, Chengyi Wang,  
 636 TianTian Fan, Zhengyin Du, Xiangpeng Wei, et al. Vapo: Efficient and reliable reinforcement  
 637 learning for advanced reasoning tasks. *arXiv preprint arXiv:2504.05118*, 2025a.

638

639 Yufeng Yuan, Yu Yue, Ruofei Zhu, Tiantian Fan, and Lin Yan. What's behind ppo's collapse in  
 640 long-cot? value optimization holds the secret. *arXiv preprint arXiv:2503.01491*, 2025b.

641

642 Weiha Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-  
 643 zoo: Investigating and taming zero reinforcement learning for open base models in the wild, 2025.  
 644 URL <https://arxiv.org/abs/2503.18892>.

645

646 Xiaojiang Zhang, Jinghui Wang, Zifei Cheng, Wenhao Zhuang, Zheng Lin, Minglei Zhang, Shaojie  
 647 Wang, Yinghan Cui, Chao Wang, Junyi Peng, et al. Srpo: A cross-domain implementation of  
 648 large-scale reinforcement learning on llm. *arXiv preprint arXiv:2504.14286*, 2025.

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810

811

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000



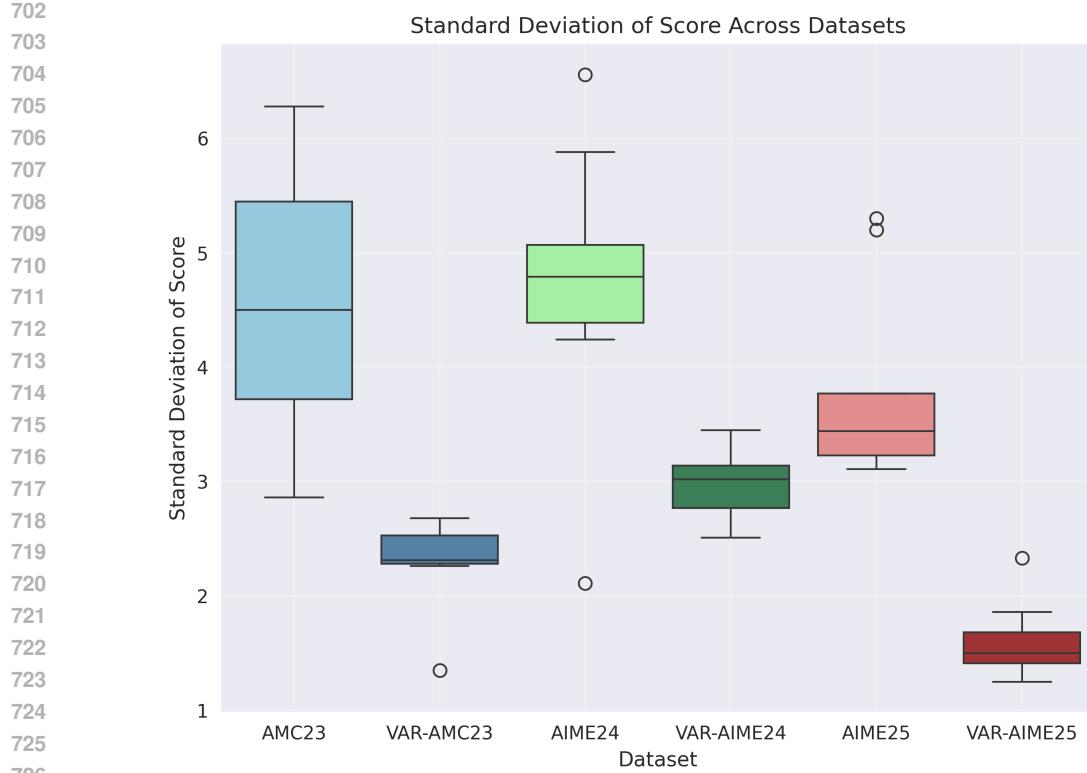


Figure 4: Standard deviation of model scores. VAR-MATH significantly reduces output variance across AMC23, AIME24, and AIME25.

731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
Table 5: Statistics of the original and variabilized benchmark datasets.

Dataset	Original Questions	Symbolizable Questions	Variant Questions
VAR-AMC23	40	37	183
VAR-AIME24	30	24	126
VAR-AIME25	30	25	130

## E EVALUATION DETAILS

### E.1 DATASETS AND TESTING ENVIRONMENT

We evaluate model performance on six mathematical reasoning benchmarks: the original **AMC23**, **AIME24**, and **AIME25**, together with their variabilized counterparts **VAR-AMC23**, **VAR-AIME24**, and **VAR-AIME25**, constructed using the symbolic multi-instantiation pipeline described in Section 3. The original AMC23<sup>2</sup>, AIME24<sup>3</sup>, and AIME25<sup>4</sup> datasets are sourced from Hugging Face.

The evaluation framework is based on the open-source Qwen2.5-MATH repository<sup>5</sup>, and uses PyTorch (v2.3.0), Transformers (v4.51.3), and vLLM (v0.5.1) for efficient decoding. All experiments are conducted on NVIDIA RTX A6000 GPUs with bf16 precision.

<sup>2</sup><https://huggingface.co/datasets/zwhe99/amc23>

<sup>3</sup>[https://huggingface.co/datasets/Maxwell-Jia/AIME\\_2024](https://huggingface.co/datasets/Maxwell-Jia/AIME_2024)

<sup>4</sup><https://huggingface.co/datasets/math-ai/aime25>

<sup>5</sup><https://github.com/QwenLM/Qwen2.5-Math>

	orl_answer	orl_question	VAR_question	VAR_info	VAR_round	VAR_answer	VAR_answer_0 und
756							
757	27	Cities \$A\\$ and \$B\\$ are \$45\\$ miles apart. Alicia lives in \$A\\$ and Beth lives in \$B\\$. Alicia bikes towards \$B\\$ at 18 miles per hour. Leaving at the same time, Beth bikes toward \$A\\$ at 12 miles per hour. How many miles from City \$A\\$ will they be when they meet?	Cities \$A\\$ and \$B\\$ are \$45\\$ miles apart. Alicia lives in \$A\\$ and Beth lives in \$B\\$. Alicia bikes towards \$B\\$ at \$VAR\_X\$ miles per hour. Leaving at the same time, Beth bikes toward \$A\\$ at \$VAR\_Y\$ miles per hour. How many miles from City \$A\\$ will they be when they meet?	VAR_X=Random_linspace_[10,20,2] VAR_Y=Expression_30-VAR_X	0	Expression_45*VAR_X/(VAR_X+VAR_Y)	0
758	36	Positive real numbers \$x\\$ and \$y\\$ satisfy \$y^3=x^2\\$ and \$y-x^2=4y-25\$. What is \$x+y\\$?	Positive real numbers \$x\\$ and \$y\\$ satisfy \$y^3=x^2\\$ and \$y-x^2=4y-25\$. What is \$x+y\\$?	VAR_X=Random_Set_[4,9,16,25,36]	0	Expression_(VAR_X**0.5+1)**2*(VAR_X**0.5+2)	0
759	45	What is the degree measure of the acute angle formed by lines with slopes \$2\\$ and \$1/\frac{1}{3}\\$?	What is the degree measure of the acute angle formed by lines with slopes \$VAR\_Y\\$ and \$VAR\_X\\$?	VAR_X=Random_Set_[10,12,6,9/5] VAR_Y=Expression_(1+VAR_X)/(1-VAR_X)	-1	Expression_45	0
760	3159	What is the value of \$ 2^2-3-1^2+3+4^2-3^2+6^2-5^2+10^2+18^2-17^2-9^2 \$	What is the value of \$ 2^2-3-1^2+3+4^2-3^2+6^2-5^2+10^2+18^2-17^2-9^2 \$	VAR_X=Random_linspace_[18,38,2] VAR_Y=Expression_VAR_X-1	0	Expression_(VAR_X+2)*VAR_X**2-3*VAR_X**2	0
761	7	How many complex numbers satisfy the equation \$z^2=\overline{z}\$, where \$\overline{z}\$ is the conjugate of the complex number \$z\$?	How many complex numbers satisfy the equation \$z^2=\overline{z}\$, where \$\overline{z}\$ is the conjugate of the complex number \$z\$?	VAR_X=Random_linspace_[5,25,2]	0	Expression_VAR_X+2	0
762	21	Consider the set of complex numbers \$z\\$ satisfying \$ 1+z+z^2 =4\$. The maximum value of the imaginary part of \$z\\$ can be written in the form \$\frac{m}{n}\\$ \{m,n\} \\$, where \$m\\$ and \$n\\$ are relatively prime integers. What is \$m+n\\$?	Consider the set of complex numbers \$z\\$ satisfying \$ 1+z+z^2 =4\$. The maximum value of the imaginary part of \$z\\$ can be written in the form \$\frac{m}{n}\\$ \{m,n\} \\$, where \$m\\$ and \$n\\$ are relatively prime integers. What is \$m+n\\$?	VAR_X=Random_Set_[2,3,4,5,7]	0	Expression_5+VAR_X*4	0
763	3	Flora the frog starts at 0 on the number line and makes a sequence of jumps to the right. In any one jump, independent of previous jumps, Flora leaps a positive integer distance \$m\\$ with probability \$\frac{1}{m}\\$.	Flora the frog starts at 0 on the number line and makes a sequence of jumps to the right. In any one jump, independent of previous jumps, Flora leaps a positive integer distance \$m\\$ with probability \$\frac{1}{m}\\$.	VAR_X=Random_linspace_[10,200,13]	0	Expression_3	0
764	96	Let \$f\\$ be the unique function defined on the positive integers such that \$\lfloor \sum_{d m} \frac{1}{d} \rfloor \cdot \text{cdot} \lfloor \frac{1}{f(m)} \rfloor \cdot \text{cdot} \lfloor \frac{1}{f(f(m))} \rfloor = 1\$ for all positive integers \$m\\$. What is \$f(2023)\\$?	Let \$f\\$ be the unique function defined on the positive integers such that \$\lfloor \sum_{d m} \frac{1}{d} \rfloor \cdot \text{cdot} \lfloor \frac{1}{f(m)} \rfloor \cdot \text{cdot} \lfloor \frac{1}{f(f(m))} \rfloor = 1\$ for all positive integers \$m\\$. What is \$f(2023)\\$?	VAR_X=Fixed_Set_[2019,2021,2023,2027,2029]	0	d_Set_[1344,1932,96,-206,-20]	0
765	1	How many ordered pairs of positive real numbers \$a,b\\$ satisfy the equation \$ (1+2a)(2+2b)(2a+b) =32ab\\$?	How many ordered pairs of positive real numbers \$a,b\\$ satisfy the equation \$ (1+2a)(2+2b)(2a+b) =32ab\\$?	VAR_X=Random_linspace_[1,10,1] VAR_Y=Random_linspace_[2,12,2]	0	Expression_1	0
766	8	How many positive perfect squares less than \$82023\\$ are divisible by \$5\\$?	How many positive perfect squares less than \$82023\\$ are divisible by \$5\\$?	VAR_X=Random_Set_[3,5,7,11,13] VAR_Z=Random_Set_[3,4,5,6,7] VAR_Y=Random_Set_[3,5,7,9,11] VAR_Y=Expression_3*VAR_X-VAR_Z	0	Expression_int(int(2023**0.5)/VAR_X)	0
767	18	How many digits are in the base-ten representation of \$8^8 \cdot 5 \cdot 10^5 \cdot 15^5\\$?	How many digits are in the base-ten representation of \$8^8 \cdot VAR\_X \cdot \text{cdot} 5^5 \cdot (VAR\_Y) \cdot \text{cdot} 15^5 \cdot VAR\_Z\\$?	VAR_X=Random_linspace_[5,25,2]	0	Expression_int(np.log(10)*VAR_Z)+1+3*VAR_X	0

Figure 5: Illustrative examples of symbolic abstraction and metadata in VAR-MATH.

**Question:** Consider the set of complex numbers \$z\$ satisfying \$|1 + z + z^2| = 4\$. The maximum value of the imaginary part of \$z\$ can be written in the form \$\frac{\sqrt{m}}{n}\$, where \$m\$ and \$n\$ are relatively prime positive integers. What is \$m + n\$?

**Answer:** 21

**Symbolic Question:** Consider the set of complex numbers \$z\$ satisfying \$|1 + z + z^2| = VAR\\_X\$. The maximum value of the imaginary part of \$z\$ can be written in the form \$\frac{\sqrt{m}}{n}\$, where \$m\$ and \$n\$ are relatively prime positive integers. What is \$m + n\$?

**Feasible Set:** \$VAR\\_X \sim \{2,3,4,5, \dots\}

**Answer:** \$5 + VAR\\_X \* 4

Figure 6: Example of original and symbolic variants from AMC23 and VAR-AMC23.

**Question:** Jen enters a lottery by picking 4 distinct numbers from \$S = \{1,2,3, \dots, 9,10\}\$. 4 numbers are randomly chosen from \$S\$. She wins a prize if at least two of her numbers were 2 of the randomly chosen numbers, and wins the grand prize if all four of her numbers were the randomly chosen numbers. The probability of her winning the grand prize given that she won a prize is \$\frac{\sqrt{m}}{n}\$, where \$m\$ and \$n\$ are relatively prime integers. Find \$m + n\$.

**Answer:** 116

**Symbolic Question:** Jen enters a lottery by picking 4 distinct numbers from \$S = \{1,2,3, \dots, 9, VAR\\_X\}\$. 4 numbers are randomly chosen from \$S\$. She wins a prize if at least two of her numbers were 2 of the randomly chosen numbers, and wins the grand prize if all four of her numbers were the randomly chosen numbers. The probability of her winning the grand prize given that she won a prize is \$\frac{\sqrt{m}}{n}\$, where \$m\$ and \$n\$ are relatively prime integers. Find \$m + n\$.

**Feasible Set:** \$VAR\\_X \sim \{10,11, \dots, 20\}

**Answer:** \$(3 \* VAR\\_X - 11) \* (VAR\\_X - 4) + 2

Figure 7: Example of original and symbolic variants from AIME24 and VAR-AIME24.

810  
811  
812  
813  
814  
815

**Question:** Six points  $A, B, C, D, E$  and  $F$  lie in a straight line in that order. Suppose that  $G$  is a point not on the line and that  $AC = 26$ ,  $BD = 22$ ,  $CE = 31$ ,  $DF = 33$ ,  $AF = 73$ ,  $CG = 40$ , and  $DG = 30$ . Find the area of  $\triangle BGE$ .

**Answer:** 468

816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827

**Symbolic Question:** Six points  $A, B, C, D, E$  and  $F$  lie in a straight line in that order. Suppose that  $G$  is a point not on the line and that  $AC = VAR\_Y$ ,  $BD = VAR\_Z$ ,  $CE = 31$ ,  $DF = 33$ ,  $AF = VAR\_U$ ,  $CG = 40$ , and  $DG = 30$ . Find the area of  $\triangle BGE$ .

**Feasible Set:**

$$\begin{aligned} VAR\_X &\sim \{4, 5, \dots, 12\} \\ VAR\_Y &= VAR\_X + 18 \\ VAR\_Z &= VAR\_X + 14 \\ VAR\_U &= VAR\_X + 65 \end{aligned}$$

**Answer:**  $12 * (VAR\_X + 31)$

Figure 8: Example of original and symbolic variants from AIME25 and VAR-AIME25.

## E.2 GENERATION CONFIGURATION

For 7B and 32B models, we adopt the system prompts and decoding configurations from their official implementations. The decoding hyperparameters are summarized in Table 6. High-capacity models are accessed via official APIs and evaluated using their default generation settings, without modification or additional prompts.

Table 6: Decoding and runtime configurations for model evaluation.

Hyperparameter	Value
<i>General settings</i>	
Temperature	0.6
Number of generations	16
Top- $p$	1.0
Use vLLM	True
GPU	NVIDIA RTX A6000
<i>7B-parameter models</i>	
Max tokens per call	8192
GPUs used per model	2
M	16
N	1000
<i>32B-parameter models</i>	
Max tokens per call	32768
GPUs used per model	4
M	16
N	1000
<i>Frontier models</i>	
M	4
N	1000

## F MORE DISCUSSION

### F.1 STRICT DROP APPLES-TO-ORANGES METRIC

A natural concern is that the performance drop between the original AMC/AIME datasets and their variabilized counterparts may partially arise from a mismatch in evaluation granularity: original scores are computed using pass@1, whereas strict VAR-AMC/AIME uses a consistency require-

864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917  
ment across  $K$  variants. This raises the question of whether the observed decline is an artifact of comparing pass@1 to a “ $K/K$  strict” metric.

To address this, we introduce a *strict-AMC/AIME* metric that mirrors strict VAR-AMC/AIME. For each original problem, we perform  $K$  independent inference runs and count the item as correct only if all  $K$  runs are correct, using the same  $K$ , sampling strategy, and bootstrap procedure as in the variabilized evaluation. As shown in Tables 7–9, strict-AMC/AIME results remain close to the original pass@1 scores, whereas the drop from strict-AMC/AIME to strict-VAR-AMC/AIME remains large and consistent across models (23%/18%/31% on AMC, AIME24, and AIME25, respectively). This demonstrates that the strict consistency requirement itself does not account for the degradation.

Table 7: Strict-Metric Performance on AMC23 and VAR-AMC23. (Avg. Drop –23.25%)

Model	(strict) AMC23	(strict) VAR-AMC23	Drop
Qwen2.5-MATH-7B	12.2 (4.0)	2.0 (2.0)	<b>-83.6%</b>
Eurus-2-7B-PRIME	40.8 (3.6)	28.9 (3.7)	<b>-29.2%</b>
Skywork-OR1-Math-7B	63.4 (3.2)	57.0 (3.6)	<b>-10.1%</b>
SimpleRL-Zoo-7B	42.1 (4.3)	33.6 (4.0)	<b>-20.2%</b>
Light-R1-7B-DS	59.9 (4.2)	54.9 (4.6)	<b>-8.3%</b>
Oat-Zero-7B	55.0 (3.1)	36.8 (3.3)	<b>-33.1%</b>
Qwen2.5-32B	6.5 (3.5)	3.1 (2.5)	<b>-52.3%</b>
DAPO-Qwen-32B	85.9 (3.4)	69.8 (3.1)	<b>-18.7%</b>
SRPO-Qwen-32B	72.4 (4.1)	51.5 (4.5)	<b>-28.9%</b>
DeepSeek-R1-0528	100.0 (0.0)	96.4 (2.5)	<b>-3.6%</b>
Qwen3-235B-A22B	100.0 (0.0)	93.6 (3.1)	<b>-6.4%</b>
SEED-THINK-v1.6	100.0 (0.0)	98.8 (1.5)	<b>-1.2%</b>
OpenAI-o4-mini-high	100.0 (0.0)	93.4 (2.3)	<b>-6.6%</b>

Table 8: Strict-Metric Performance on AIME24 and VAR-AIME24. (Avg. Drop –17.87%)

Model	(strict) AIME24	(strict) VAR-AIME24	Drop
Qwen2.5-MATH-7B	3.4 (2.8)	3.2 (2.7)	<b>-5.9%</b>
Eurus-2-7B-PRIME	6.7 (2.6)	4.3 (2.9)	<b>-35.8%</b>
Skywork-OR1-Math-7B	27.1 (3.7)	23.9 (4.3)	<b>-11.8%</b>
SimpleRL-Zoo-7B	11.3 (4.1)	8.5 (3.7)	<b>-24.8%</b>
Light-R1-7B-DS	24.0 (4.7)	23.8 (4.8)	<b>-0.8%</b>
Oat-Zero-7B	25.1 (3.2)	12.8 (3.6)	<b>-49.0%</b>
Qwen2.5-32B	2.7 (2.4)	2.3 (2.3)	<b>-14.8%</b>
DAPO-Qwen-32B	36.0 (4.1)	29.8 (4.6)	<b>-17.2%</b>
SRPO-Qwen-32B	39.0 (4.5)	29.2 (4.2)	<b>-25.1%</b>
DeepSeek-R1-0528	81.8 (2.8)	73.7 (3.8)	<b>-9.9%</b>
Qwen3-235B-A22B	80.8 (2.6)	69.5 (3.4)	<b>-14.0%</b>
SEED-THINK-v1.6	82.8 (2.3)	73.4 (3.5)	<b>-11.4%</b>
OpenAI-o4-mini-high	88.5 (1.9)	78.1 (3.7)	<b>-11.8%</b>

918

919

Table 9: Strict-Metric Performance on AIME25 and VAR-AIME25. (Avg. Drop  $-31.12\%$ )

920

921

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

## F.2 STATISTICAL SIGNIFICANCE CHECK

To ensure that the observed differences are statistically reliable, we perform a one-sided  $t$ -test on the  $M = 16$  independent inference runs for each open-weight model. For each model–benchmark pair, we test the null hypothesis

$$H_0 : \mu_{\text{loose}} = \mu_{\text{orig}},$$

i.e., the loose VAR-AMC/AIME score is equal to the original score, against the alternative hypothesis

$$H_1 : \mu_{\text{loose}} < \mu_{\text{orig}}.$$

This directly evaluates whether the variabilized versions lead to a statistically significant decline in performance under matched sampling conditions.

As shown in Tables 10–12, 18/27 model–benchmark pairs yield  $p < 0.05$ , allowing us to reject  $H_0$  with at least 95% confidence. This confirms that, for the majority of settings, loose scores are significantly lower than original scores. For the remaining cases with  $p \geq 0.05$ , most correspond to models whose original accuracy is already very low (approximately  $3.5 \sim 33$ ), leaving limited room for further decline and therefore a weaker statistical signal.

**Remark.** We chose not to run the  $t$ -test directly on the bootstrap replicates because these are resamples from the empirical distribution induced by the original  $M = 16$  runs. Applying a  $t$ -test on the bootstrap draws would effectively compare two empirical distributions generated from the same finite sample, which in our experiments leads to uniformly tiny  $p$ -values (often  $< 0.01$ ) and reflects the resampling procedure more than the underlying inference variability. To avoid overstating significance, we therefore conduct the  $t$ -test on the original  $M$  independent runs, and use the bootstrap only to stabilize point estimates and confidence intervals.

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972  
973

974 Table 10: Significance Check on AMC23 v.s. (loose) VAR-AMC23.

Model	AMC23	(loose) VAR-AMC23	Drop	p-value
Qwen2.5-MATH-7B	36.9 (6.3)	22.7 (2.6)	-38.5%	< 0.01**
Eurus-2-7B-PRIME	58.3 (4.3)	49.9 (2.5)	-14.3%	< 0.01**
Skywork-OR1-Math-7B	73.9 (5.4)	72.0 (2.3)	-2.6%	0.12
SimpleRL-Zoo-7B	61.4 (4.8)	52.2 (2.3)	-15.0%	< 0.01**
Light-R1-7B-DS	78.6 (6.3)	75.8 (2.3)	-3.5%	0.05*
Oat-Zero-7B	65.6 (3.1)	54.4 (2.4)	-17.0%	< 0.01**
Qwen2.5-32B	33.4 (4.5)	27.4 (2.8)	-18.2%	< 0.01**
DAPO-Qwen-32B	92.3 (2.9)	85.7 (1.4)	-7.2%	< 0.01**
SRPO-Qwen-32B	86.7 (3.7)	73.9 (2.6)	-14.8%	< 0.01**

989  
990  
991

992 Table 11: Significance Check on AIME24 v.s. (loose) VAR-AIME24.

Model	AIME24	(loose) VAR-AIME24	Drop	p-value
Qwen2.5-MATH-7B	10.8 (4.5)	7.9 (2.9)	-27.1%	0.02*
Eurus-2-7B-PRIME	15.8 (4.8)	13.4 (2.7)	-15.5%	0.04*
Skywork-OR1-Math-7B	41.5 (4.2)	39.0 (3.4)	-6.0%	0.03*
SimpleRL-Zoo-7B	23.8 (5.9)	20.4 (3.5)	-14.1%	< 0.01**
Light-R1-7B-DS	40.8 (5.1)	40.6 (3.3)	-0.6%	0.41
Oat-Zero-7B	34.0 (2.1)	22.3 (2.5)	-34.3%	< 0.01**
Qwen2.5-32B	8.8 (4.4)	7.9 (2.6)	-9.6%	0.27
DAPO-Qwen-32B	51.7 (6.6)	50.9 (2.8)	-1.5%	0.34
SRPO-Qwen-32B	55.6 (5.0)	46.9 (2.9)	-15.7%	< 0.01**

1008  
10091010  
1011

1012 Table 12: Significance Check on AIME25 v.s. (loose) VAR-AIME25.

Model	AIME25	(loose) VAR-AIME25	Drop	p-value
Qwen2.5-MATH-7B	4.8 (3.1)	3.2 (1.3)	-34.2%	0.04*
Eurus-2-7B-PRIME	10.0 (3.1)	7.4 (1.4)	-26.0%	< 0.01**
Skywork-OR1-Math-7B	24.0 (3.8)	23.4 (1.6)	-2.4%	0.31
SimpleRL-Zoo-7B	12.5 (3.4)	11.5 (1.6)	-7.9%	0.16
Light-R1-7B-DS	32.7 (3.8)	30.3 (1.8)	-7.3%	0.02*
Oat-Zero-7B	9.2 (3.2)	8.4 (1.4)	-7.9%	0.24
Qwen2.5-32B	3.5 (3.4)	2.8 (1.2)	-21.2%	0.25
DAPO-Qwen-32B	37.3 (5.3)	32.2 (2.3)	-13.5%	< 0.01**
SRPO-Qwen-32B	26.5 (5.2)	25.2 (1.7)	-4.6%	0.18

1024  
1025

1026 **G DETAILED EXAMPLE IN CONSTRUCTING THE PROBLEM**  
10271028 In this section, we provide two detailed examples (one easy and one hard) to illustrate the full  
1029 construction pipeline.  
10301031 **Example Conversion Pipeline (AMC 2023, Problem 11)**  
10321033 **Original Problem.** What is the degree measure of the acute angle formed by lines with slopes  
1034 2 and  $\frac{1}{3}$ ?  
10351036 **1. Structural Analysis.** A mathematics expert solves the problem and identifies the key structure:  
1037

1038 
$$\tan(\theta) = \left| \frac{m_1 - m_2}{1 + m_1 m_2} \right|.$$
  
1039

1040 Essential constants are  $m_1 = 2$  and  $m_2 = \frac{1}{3}$ , yielding  $\theta = 45^\circ$ . This determines which  
1041 quantities can be symbolized while preserving semantic fidelity.  
10421043 **2. Symbolic Parameterization.** Key constants are replaced with variables:  
1044

1045 
$$m_1 = \text{VAR\_Y}, \quad m_2 = \text{VAR\_X}.$$
  
1046

1047 Feasible domains:  
1048

1049 
$$\text{VAR\_X} \in \left\{ \frac{1}{3}, \frac{1}{2}, \frac{3}{5} \right\}, \quad \text{VAR\_Y} = \frac{1 + \text{VAR\_X}}{1 - \text{VAR\_X}},$$
  
1050

1051 chosen to stay close to the original scale while ensuring validity (nonzero denominator, positive  
1052 slopes) and similar difficulty (share the final solution step  $\arctan 1 = 45^\circ$ ).  
10531054 **3. Parametric Solution Formulation.**  
1055

1056 
$$\tan(\theta) = \left| \frac{\text{VAR\_Y} - \text{VAR\_X}}{1 + \text{VAR\_Y} \cdot \text{VAR\_X}} \right| = 1 \Rightarrow \theta = 45^\circ.$$
  
1057

1058 Thus the symbolic answer is:  
1059

1060 
$$\text{Answer} = 45^\circ.$$
  
1061

1062 **4. Verification.**  
10631064 

- **Human verification:** another annotator solves each instantiated variant to confirm  
1065 correctness and comparable difficulty.
- **Model verification:** all variants are run through a frontier model (DeepSeek). Any  
1066 inconsistency triggers re-checking by an expert. Geometry problems are additionally  
1067 validated using drawing tools.

1068 **5. Variant Sampling and Evaluation.** Variants are instantiated by sampling from the feasible  
1069 domains:  
10701071 

- Variant 1:  $\text{VAR\_X} = \frac{1}{3}$ ,  $\text{VAR\_Y} = 2$ . Slopes: 2 and  $\frac{1}{3}$ . Ground truth:  $45^\circ$ .
- Variant 2:  $\text{VAR\_X} = \frac{1}{2}$ ,  $\text{VAR\_Y} = 3$ . Slopes: 3 and  $\frac{1}{2}$ . Ground truth:  $45^\circ$ .
- Variant 3:  $\text{VAR\_X} = \frac{3}{5}$ ,  $\text{VAR\_Y} = 4$  (invalid). Slopes: 4 and  $\frac{3}{5}$ . Ground truth:  $45^\circ$ .

1076 Valid variants (typically  $K = 3$  for this problem) are fixed and shared by all models.  
1077 Loose/strict scores are computed using standardized prompting.  
1078

1079

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

## Example Conversion Pipeline (AIME 2025 II, Problem 1)

**Original Problem.** Six points  $A, B, C, D, E, F$  lie on a line in that order. A point  $G$  is not on the line, and the distances satisfy

$$AC = 26, \quad BD = 22, \quad CE = 31, \quad DF = 33, \quad AF = 73, \quad CG = 40, \quad DG = 30.$$

Find the area of  $\triangle BGE$ .

**1. Structural Analysis.** Following the official solution, set

$$AB = a, \quad BC = b, \quad CD = c, \quad DE = d, \quad EF = e.$$

Then

$$\begin{aligned} a + b + c + d + e &= AF = 73, \\ a + b &= AC = 26, \\ b + c &= BD = 22, \\ c + d &= CE = 31, \\ d + e &= DF = 33. \end{aligned}$$

From these equations we deduce

$$c = 14, \quad a + e = 34, \quad b + c + d = 39.$$

Using Heron's formula on  $\triangle CGD$  with side lengths  $CG = 40$ ,  $DG = 30$ , and  $CD = c = 14$  gives

$$[CGD] = \sqrt{42 \cdot 2 \cdot 12 \cdot 28} = 168.$$

Since  $BE = b + c + d = 39$  and  $CD = c = 14$  lie on the same line with the same altitude from  $G$ ,

$$\frac{[BGE]}{[CGD]} = \frac{BE}{CD} = \frac{39}{14},$$

so

$$[BGE] = 168 \cdot \frac{39}{14} = 468.$$

This structure (solving for  $c$  and  $b + c + d$ , then scaling areas by base ratio) is what we preserve in the symbolic version.

**2. Symbolic Parameterization.** We vary the lengths  $AC$ ,  $BD$ , and  $AF$  while keeping the configuration valid. Introduce a shift parameter

$$\text{VAR\_X} \in \{4, 5, 6, 7, 8, 9, 10, 11, 12\},$$

and define

$$\begin{aligned} AC &= \text{VAR\_Y} = \text{VAR\_X} + 18, \\ BD &= \text{VAR\_Z} = \text{VAR\_X} + 14, \\ AF &= \text{VAR\_U} = \text{VAR\_X} + 65, \end{aligned}$$

while keeping

$$CE = 31, \quad DF = 33, \quad CG = 40, \quad DG = 30$$

unchanged. The points remain ordered  $A, B, C, D, E, F$  on the line and  $G$  stays off the line.

**3. Parametric Solution Formulation.** With the same notation  $AB = a$ ,  $BC = b$ ,  $CD = c$ ,  $DE = d$ ,  $EF = e$ , the constraints become

$$\begin{aligned} a + b + c + d + e &= AF = \text{VAR\_X} + 65, \\ a + b &= AC = \text{VAR\_X} + 18, \\ b + c &= BD = \text{VAR\_X} + 14, \\ c + d &= CE = 31, \\ d + e &= DF = 33. \end{aligned}$$

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

From these equations we obtain, exactly as in the original case,

$$c = 14, \quad d = 17, \quad b = \text{VAR\_X}, \quad b + c + d = \text{VAR\_X} + 31.$$

Thus  $CD$  remains 14, so  $\triangle CGD$  still has side lengths 40, 30, and 14, and its area is always

$$[CGD] = 168.$$

Meanwhile, the base of  $\triangle BGE$  is

$$BE = b + c + d = \text{VAR\_X} + 31,$$

so with the same altitude from  $G$ ,

$$\frac{[BGE]}{[CGD]} = \frac{BE}{CD} = \frac{\text{VAR\_X} + 31}{14},$$

which yields the parametric area

$$[BGE] = 168 \cdot \frac{\text{VAR\_X} + 31}{14} = 12(\text{VAR\_X} + 31).$$

Hence the symbolic answer is

$$\text{Answer} = 12(\text{VAR\_X} + 31).$$

#### 4. Verification.

- **Human verification:** Another expert re-solves several instantiated variants using this symbolic derivation
- **Model verification:** All variants are also checked by a frontier model (DeepSeek) as a sanity check; any discrepancy triggers a second expert review to confirm correctness and comparable difficulty.

#### 5. Variant Sampling and Evaluation.

We sample  $\text{VAR\_X}$  from its feasible set, for example:

- $\text{VAR\_X} = 4$ : area =  $12(4 + 31) = 420$ .
- $\text{VAR\_X} = 7$ : area =  $12(7 + 31) = 456$ .
- $\text{VAR\_X} = 11$ : area =  $12(11 + 31) = 504$ .
- ...

These variants are fixed and shared across all models. Loose/strict scores are computed using standardized prompting.

## H THE USE OF LLMS

The authors utilized LLMs to assist with writing tasks such as text polishing and language refinement. All substantive intellectual content, research findings, and technical contributions are original to the authors. The LLM served only as a writing assistance tool under human supervision, and all output was critically evaluated and modified by the authors.