
RealMath: A Continuous Benchmark for Evaluating Language Models on Research-Level Mathematics

Jie Zhang Cezara Petruī Kristina Nikolić Florian Tramèr

ETH Zurich

Abstract

Existing benchmarks for evaluating mathematical reasoning in large language models (LLMs) rely primarily on competition problems, formal proofs, or artificially challenging questions—failing to capture the nature of mathematics encountered in actual research environments. We introduce REALMATH, a novel benchmark derived directly from research papers and mathematical forums that assesses LLMs’ abilities on authentic mathematical tasks. Our approach addresses three critical challenges: sourcing diverse research-level content, enabling reliable automated evaluation through verifiable statements, and designing a continually refreshable dataset to mitigate contamination risks. Experimental results across multiple LLMs reveal surprising capabilities in handling research mathematics compared to competition problems, suggesting current models may already serve as valuable assistants for working mathematicians despite limitations on highly challenging problems. The code and dataset for REALMATH are publicly available¹.

1 Introduction

The mathematical capabilities of Large Language Models (LLMs) have become a critical lens for assessing their reasoning and knowledge retention abilities. Although considerable effort has been invested in evaluating LLMs in basic mathematics [5, 11], competition-level mathematics [8, 10, 17, 19, 20] and formal proof generation [23, 24, 25], these evaluations may not adequately reflect the potential utility of LLMs in real-world mathematical research contexts.

Current mathematical benchmarks fall predominantly into three categories: (1) those derived from course materials [5, 11] or competitive examinations [8] (e.g., IMO or AIME) that offer abundant problems with solutions; (2) those centered on formal theorem proving where verification can be automated [24, 25]; (3) those designed by mathematical experts to be as challenging as possible [9, 16]. However, these benchmarks capture only a narrow slice of mathematical practice. The mathematics encountered “in the wild”—particularly in research settings—differs substantially from competition problems in structure and topics, rarely relies on formal proofs, and considers statements and results that are not (exclusively) designed to be maximally challenging.

This disconnect raises a fundamental question:

How effective might LLMs be as assistants for practicing mathematicians today?

To address this question, we introduce a novel benchmark designed to evaluate LLMs on research-level mathematics extracted directly from the literature. Constructing such a benchmark presents three significant challenges:

¹Code available at: [Github](#); Dataset available at: [Huggingface](#).

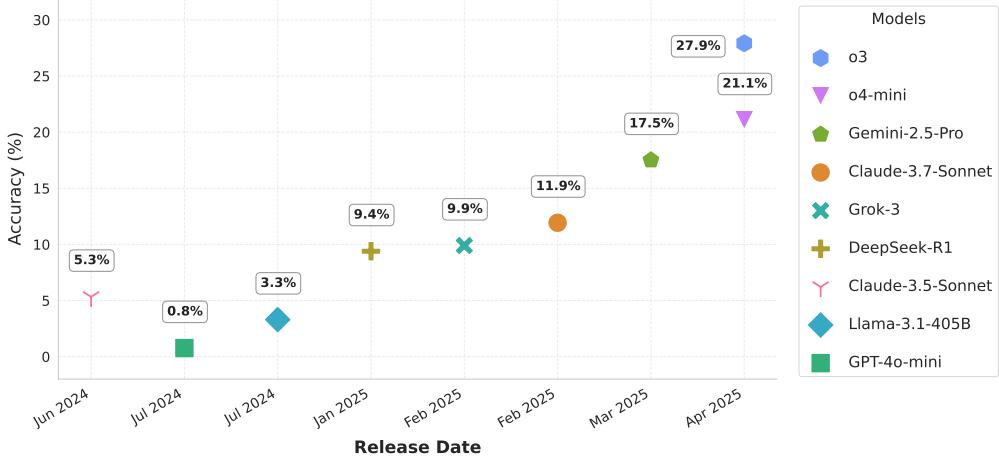


Figure 1: LLM performance on a hard subset of REALMATH from arXiv Mathematics papers.

First, we need to *source authentic content* that faithfully represents the diversity and complexity of contemporary mathematical research.

Second, we require a *reliable method for assessing correctness*. Unlike competition problems with standardized solutions or problems with formal proofs, research mathematics presents evaluation difficulties: human expert validation is resource-intensive and limits scalability, while using LLMs as judges introduces reliability concerns.

Third, we must address *test set contamination over time*, as mathematical content incorporated in benchmarks might be absorbed into training datasets for future models.

In this paper, we present REALMATH, a math benchmark that addresses these challenges through:

1. A data pipeline that extracts verifiable mathematical statements from research papers (e.g., arXiv) and mathematical forums (e.g., Stack Exchange), creating a rich corpus of research-level content.
2. An evaluation methodology focused on verifiable answers rather than proof assessment, allowing for automated correctness checking.
3. A continually refreshable dataset design that leverages the vast and growing body of mathematical literature, allowing for regular updates with new content to mitigate contamination concerns.

We systematically transform mathematical statements into question-answer pairs, preserving surrounding context necessary for comprehension. For example, a theorem stating that “*For $n \geq 1$, the number of shallow, 123-avoiding centrosymmetric permutations of length n is $\frac{n^2}{4} + 1$ when n is even*” [1] becomes a question, “*What is the number of 123-avoiding centrosymmetric permutations of length n when n is even?*”, with a verifiable answer, “ $\frac{n^2}{4} + 1$ ”, with relevant definitions and notation from the source material given as context.

From approximately 9,000 mathematics-related academic papers collected over nine months, our automated pipeline curated 633 high-quality samples and can generate over 70 new samples each month. Our evaluation of frontier LLMs reveals interesting patterns in mathematical capabilities that differ from those observed in other recent benchmarks. In particular, models demonstrate stronger performance on our research-mathematics benchmark (see Figure 1) than on deliberately challenging datasets such as math competitions or FrontierMath [9], suggesting that LLMs may already provide valuable assistance in research contexts, even though they cannot solve the most advanced problems.

Although our approach focuses on statement verification rather than proof generation or verification, it nonetheless provides a valuable signal about LLMs’ potential as mathematical assistants. Our findings suggest that current models may already serve as useful tools in mathematical research contexts, even as they continue to struggle with the most challenging mathematical problems.

To summarize, this work makes several contributions: it establishes a new paradigm for evaluating the mathematical capabilities of LLMs using organic research Mathematics; it provides a sustainable methodology for benchmark creation that resists contamination; and it offers insights into the relative strengths of current models on tasks representative of real mathematical practice.

2 Related Work

Math benchmarks sourced from exams and competitions. Recent years have seen rapid progress in evaluating LLMs on mathematical tasks. Early data sets were sourced from school materials or entry-level math competitions such as GSM8K [5], GHOSTS [7], or MATH [11], and are approaching saturation. In response, new benchmarks now source questions from advanced math competitions such as the IMO [2, 15, 19], AIME, or the Putnam competition [20].

Math benchmarks sourced from research experts. A growing line of work seeks to evaluate mathematical reasoning on research-level questions. FrontierMath [9] (and part of HLE [16]) evaluate LLMs on extremely challenging problems crafted by expert mathematicians. These benchmarks focus on a narrow slice of current mathematical practice (at the frontier of human expertise), and are extremely labor-intensive to curate. To guard against test set contamination, FrontierMath keeps the test set private, which introduces barriers to reproducibility. In contrast, REALMATH benchmark consists of questions spanning the full range of mathematical research practice, and can be automatically refreshed as new research is published.

Benchmarks for theorem proving. Orthogonal to our work are benchmarks focused on formal proof generation and machine-verifiable mathematics, such as LeanDojo [24] and MiniF2F [25]. These evaluate LLMs’ ability to produce formal proofs within interactive theorem provers, supporting automated verification of correctness. While such work is critical for advancing formal methods and proof automation, it is not representative of most mathematical research that is not fully formalized.

Benchmarks for LLM capabilities “in the wild”. Our work complements a growing body of research that aims to evaluate the capabilities of LLMs on real-world tasks rather than well-curated proxies. For example, instead of evaluating LLMs’ coding abilities on programming competitions [12], projects such as SWE-bench [13], Lancer [14], or BaxBench [21] focus on assessing LLMs in real-world software engineering tasks. A similar shift can be observed in the evaluation of cyber-offensive capabilities of LLMs, where a number of recent benchmarks focus on evaluating real exploit capabilities [4, 6] rather than the ability to solve curated capture-the-flag competitions [18].

3 Building a Research-Level Mathematical Benchmark

3.1 Design Criteria

To ensure high quality of our collected data, we establish the following design desiderata:

1. **Real-world application focus:** Our collection methodology prioritizes mathematical problems that genuinely represent those encountered in practical scenarios. We source problems from academic research publications (e.g., arXiv preprints), where researchers tackle mathematical challenges to advance scientific knowledge, and from educational platforms (e.g., Mathematics Stack Exchange), where learners engage with mathematical concepts to develop their understanding and problem-solving capabilities. Unlike competition-oriented benchmarks, we focus on representative mathematical tasks encountered in actual practice, emphasizing procedural and technique-driven aspects rather than contest-style ingenuity.
2. **Automated verification:** We generate *constructive* problems (as in [2]) with clear, unambiguous verification criteria, typically problems with a single, exact numerical or symbolic answer. This approach excludes problems that admit multiple solutions or involve qualitative assessments such as inequalities (e.g., lower bounds, upper bounds) or asymptotic relations, which would complicate the verification process and potentially introduce ambiguity. This choice also means that we omit statements where the main difficulty lies in finding a (non-constructive) proof (e.g., if a paper has a theorem that says “ $P \neq NP$ ”, then the interesting part is solely the proof and not the construction of the statement itself).
3. **Continuous acquisition:** Leading LLMs may be trained on internet data, so it is important to avoid data contamination. To ensure this, we collect data exclusively from publicly available sources such as arXiv and StackExchange, where automatic collection is feasible with minimal human intervention. After each new model release, we can automatically gather fresh data from the internet to evaluate the model, ensuring that the evaluation set remains uncontaminated and up-to-date.

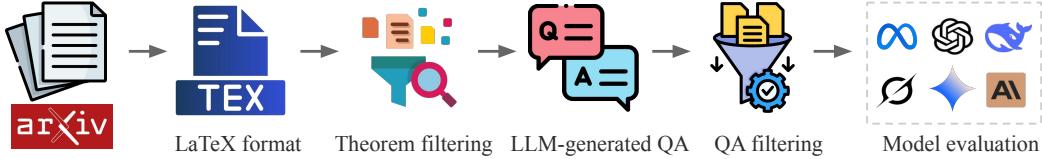


Figure 2: The data collection pipeline for arXiv papers. The core step is to ensure that each extracted theorem from arXiv papers has a single, exact answer. To maintain data quality, we apply filtering mechanisms, e.g., prompting an LLM to discard trivial samples that can be easily solved.

3.2 An Automatic and Refreshable Data Collection Pipeline

We present a detailed walkthrough of our data collection pipeline in Figure 2. Throughout the remainder of this paper, we primarily use data from arXiv papers to illustrate our methodology. A more detailed discussion of the Stack Exchange dataset is provided in Appendix A.5. As a concrete example, we use a subset of 4,000 arXiv papers selected from May to September 2022 to illustrate each stage of the data collection pipeline.

Step	# Usable Samples
Retrieve papers	05/2022 – 09/2022 (4,000 papers)
Extract LaTeX source	3,922 papers
Extract theorems	14,747 theorems
Identify constructive theorems with fixed answers	407 theorems
Generate question-answer (QA) pairs	401 QA pairs
Filter trivial questions	280 QA pairs

Table 1: An example of the pipeline steps and the number of usable samples at each stage.

Our pipeline consists of five stages:

- 1. Retrieve papers.** We collect all mathematics-related arXiv papers within a specified time window. This process involves querying the arXiv API for papers in categories such as Mathematics or Computer Science. To illustrate, we can acquire roughly 4,000 papers from Mathematics over 5 months, from May to September 2022.
- 2. Extract LaTeX source.** For each paper, we download and parse the original LaTeX source files to accurately preserve the mathematical notation. Due to occasional errors in the arXiv API responses or issues with file availability (e.g., missing or corrupted LaTeX sources), the number of successfully processed papers is often lower than the total initially retrieved. In our case, out of 4,000 papers, we obtained 3,922 usable samples.
- 3. Identifying constructive theorems with fixed answers.** We use an LLM² to identify theorems that involve the construction of a single exact answer. The LLM analyzes each extracted theorem and classifies it based on whether it presents a clear mathematical statement with an unambiguous result. We specifically exclude theorems that involve inequalities or those that admit multiple solutions, e.g., theorems of the type “*If condition A is satisfied, then relation B happens*” but where *condition A* is not the unique one for *B* to happen, and *relation B* is not an equality or fixed-answer numerical relation. From the 3,922 processed papers, we extracted 407 theorems that were classified as high-quality by our judge model, out of 14,747 theorems detected in total.
- 4. Generate question-answer (QA) pairs.** For each selected theorem, we use an LLM to convert it into a question–answer pair. LLMs will always attempt to generate a QA pair, even if the generated question–answer is easy to answer. For 407 theorems, we obtained 401 QA pairs.
- 5. Filtering trivial questions.** We also implement a post-processing stage where an LLM reviews each generated QA pair. This step filters out low-quality samples, e.g., those with easily guessable answers, or where the answer is obvious from the context fed to the LLM (it is common, for instance, for the paper’s introduction to state a theorem’s result informally). Finally, we obtained 280 QA pairs that are ready for use in the evaluation phase.

²By default, we use OpenAI’s o3-mini as the judge model. Detailed prompts are provided in Appendix A.1.

For each test sample, we provide the *context* (i.e., all relevant text preceding the theorem) along with an LLM-generated *question* to frontier LLMs. A response is considered correct only if it exactly matches the ground truth *answer*. We present some detailed examples in Figure 3.

Theorem to QA Conversion	
Original Theorem	Generated QA Pair
<p>(Theorem 1.1 in [3]) Let $q = p^{2t}$, where $p \equiv 3 \pmod{4}$ is a prime and t is a positive integer. Then, the number of cliques of order 3 in the Peisert graph $P^*(q)$ is given by $k_3(P^*(q)) = \frac{q(q-1)(q-5)}{48}$.</p>	<p>Q: Let $q = p^{2t}$, where $p \equiv 3 \pmod{4}$. What is the number of cliques of order 3 in the Peisert graph $P^*(q)$, expressed as a function of q? A: $\frac{q(q-1)(q-5)}{48}$</p>
<p>(Theorem 22 in [22]) The automorphism $\varphi_\alpha \circ (\Phi, \wedge d(\Phi))$, where $\alpha \in \mathbb{C}^*$, can be lifted to the Π-symmetric super-Grassmannian $\Pi\text{Gr}_{2k,k}$, where $k \geq 2$, if and only if $\alpha = \pm i$.</p>	<p>Q: Determine the value(s) of $\alpha \in \mathbb{C}^*$ for which the automorphism $\varphi_\alpha \circ (\Phi, \wedge d(\Phi))$ can be lifted to the super-Grassmannian $\Pi\text{Gr}_{2k,k}$ for $k \geq 2$. A: $\alpha = \pm i$</p>
Filtered Out Examples	
<p>Example 1: Let \mathbb{X} be a uniformly convex Banach space. Then $\Gamma(\mathbb{X}) < \frac{1}{2}$.</p> <p>► LLMs may respond with weaker but still valid bounds (e.g., $\Gamma(\mathbb{X}) < 1$), which are not equivalent. Stronger answers (e.g., < 0.4) require more advanced verification from the judge model.</p> <p>Example 2: For any convex body K in \mathbb{R}^n, there exists $\Phi \in \text{GL}(n)$ such that:</p> $\frac{S(\Phi K)^n}{ \Phi K ^{n-1}} \leq \frac{S(\Delta^n)^n}{ \Delta^n ^{n-1}} = \frac{n^{3n/2}(n+1)^{(n+1)/2}}{n!}.$ <p>► This is a non-constructive existence theorem. It cannot be converted into a fixed-answer QA that could easily be verified by an LLM judge.</p>	

Figure 3: Illustration of the theorem-to-QA conversion process and samples that are filtered out. The top panel shows examples of high-quality question-answer pairs generated from mathematical theorems that contain fixed, verifiable answers. The bottom panel provides examples of theorems that were filtered out due to ambiguity or the lack of a fixed answer.

This pipeline ensures that our benchmark remains current, uncontaminated, and representative of real mathematical challenges faced in research. By automating the collection process, we can continuously refresh the dataset with new mathematical problems as they emerge in the research community.

Discussion on label noise. The two primary sources of noise in our data pipeline are: (1) the inherent quality of the source materials; and (2) the reliability of using LLMs to evaluate and transform mathematical content. Samples from Mathematics Stack Exchange, in particular, tend to be of lower quality: Questions are often poorly formulated, contain mathematical inaccuracies, or lack answers on the forum. As a result, obtaining high-quality samples from this source requires filtering through a much larger volume of user-submitted content. Additionally, our pipeline relies on LLMs to assess the quality of extracted theorems and generate QA pairs. Thus, any limitations or biases in the LLM’s evaluation capabilities may introduce further inaccuracies into the dataset.

Overall, our pipeline yields more than 94% high-quality samples from arXiv papers, without relying on any human annotation. For example, in the Mathematics category, we processed more than 9,000

Table 3: Frontier LLMs achieve strong performance at answering mathematical questions extracted from research papers and math forums. See Figure 4 and Appendix B.2 for a more fine-grained evaluation based on the estimated difficulty of questions.

Dataset	o3	o4 mini	Gemini 2.5-pro	Deepseek R1	Grok 3	Claude 3.7-Sonnet	Claude 3.5-Sonnet	Llama 3.1-405B	GPT 4o-mini
Math.arXiv	49.1	43.4	32.5	30.5	29.5	34.1	18.3	16.4	12.5
CS.arXiv	44.1	42.3	25.2	31.5	25.2	31.5	16.2	15.3	7.2
Math.StackExchange	70.7	70.8	60.9	62.2	54.8	61.1	37.6	32.1	40.8

papers and extracted 633 theorems that were classified as high-quality by our judge model. We then manually reviewed each sample and found that approximately 6% did not meet our quality criteria; these were filtered out. These results suggest that our pipeline can reliably produce high-quality data with minimal human intervention, even when working with imperfect real-world sources (see Appendix A.3 for more details).

Although our data construction inherently assumes that the mathematical theorems sourced from arXiv papers are *correct*, this is, of course, not guaranteed. ArXiv papers have not necessarily been peer-reviewed, and errors or ambiguous claims may be present.³ On the one hand, incorrect statements could produce some noise in our results (which we hypothesize to be small). On the other hand, we view possible ambiguities in statements or notation as beneficial, since it is truly representative of mathematics “in the wild”, compared to the more polished and vetted content found in formal competition settings.

4 Experiments

4.1 Experimental Setup

We use the data pipeline described in Section 3 (and Appendix A) to extract over 1,200 QA pairs from three data sources, Math.arXiv, CS.arXiv, and Math.StackExchange (see Table 3 for summary statistics). We then evaluated multiple frontier models—OpenAI o3 and o4-mini, Claude 3.7 Sonnet, Gemini 2.5 Pro, Grok-3, and DeepSeek-R1—across these datasets. For comparison, we also include several earlier models, such as GPT-4o mini, Claude 3.5 Sonnet, and LLaMA-3.1-405B.

The input format varies slightly across datasets to match their structure. For CS.arXiv and Math.arXiv, each model receives the LLM-generated question along with the full relevant context preceding the target theorem, typically spanning from the introduction up to (but excluding) the theorem itself. In contrast, for Math.StackExchange, the input consists solely of the LLM-generated question, reflecting the format of the original forum posts.

4.2 Main Experimental Results

Results on frontier LLMs. Table 3 highlights the performance of the wide range of LLMs that we tested on our three datasets. The ten models under analysis show variation in accuracy, with o3 and o4-mini leading on all datasets. More specifically, o3 achieved the highest accuracy on both CS.arXiv and Math.arXiv, whereas o4-mini slightly outperformed it on the Math.StackExchange samples. Other models, such as Deepseek-R1 and Gemini-2.5-pro, showed mid-tier performance, with accuracy percentages ranging between 25.6% – 32.5% on the arXiv datasets, compared to higher accuracy levels of 60.9% – 62.2% on Math.StackExchange.

Table 2: Summary of final datasets used for evaluation.

Dataset	Time Span	# QA pairs
Math.arXiv	05/2022 – 09/2022	633
	12/2024 – 03/2025	
CS.arXiv	05/2022 – 10/2023	111
Math.StackExchange	04/2024 – 03/2025	542

³An alternative would be to limit the pipeline to peer-reviewed sources, but this would significantly reduce the diversity and coverage of available material.

We acknowledge that our full benchmark comes with relatively high initial performance, which may raise questions about its durability. In fact, a common trend in recent benchmarks has been to design problems that are as challenging as possible, resulting in near-zero performance [9, 16]. However, since our benchmark aims to track performance on *real, organic* mathematical research tasks, we do not view this as an issue. While some benchmarks track LLMs’ nascent ability to solve (a few) of the most challenging math problems, performance increases on REALMATH may be more indicative of the utility of LLMs on common mathematical research tasks. Nevertheless, in the following experiment, we also show that the difficulty of questions in our datasets is highly nonhomogeneous and that there is a “hard” subset on which current LLMs perform much worse.

Breakdown by difficulty levels. We analyze model performance across questions of varying difficulty levels. To determine these difficulty levels, we evaluated how older, weaker models performed on the dataset samples. Based on their performance, we categorized the samples into different difficulty levels, as detailed in Appendix B.2. As shown in Figure 4, accuracy consistently declines as question difficulty increases. For example, the top-performing model, o3, achieves 97.5% accuracy on easy questions, dropping to 81.4% on medium questions and further to just 27.9% on hard questions. A similar downward trend is observed for Gemini-2.5-Pro and DeepSeek-R1. This pattern highlights the significant challenge current models face in handling more complex and research-level mathematics.

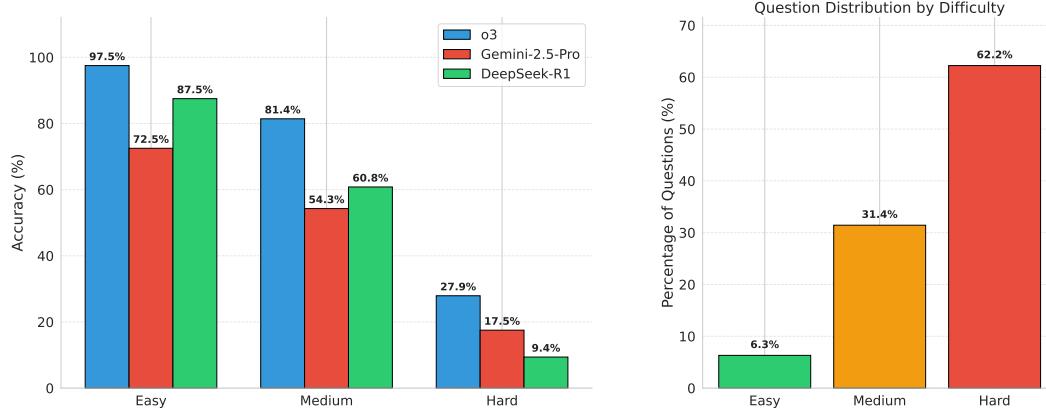


Figure 4: (a) Accuracy of the best-performing models across different difficulty levels of Math.arXiv and (b) distribution of difficulty levels on Math.arXiv. More details and results in Appendix B.2.

Performance analysis across research subcategories. The datasets we collect consist of mathematical problems in a wide variety of subjects. Using human-reported categories (arXiv categories, or StackExchange tags), we can get a more granular evaluation of LLMs’ capabilities in specific areas. Table 4 in Appendix B.4 shows the breakdown of categories for the Math.arXiv dataset, with a predominance of samples from *Number Theory (math.NT)* and *Combinatorics (math.CO)*, each comprising more than 20% of the dataset size. This distribution is consistent with other recent benchmarks, e.g., MathConstruct [2] and FrontierMath [9], and stems from a combination of factors (see Figure 9 in Appendix A.4): (1) a high number of papers on the topic; (2) a higher number of theorems per paper on average; (3) a higher fraction of constructive theorems with a fixed answer.

To illustrate LLM performance on Math.arXiv, Figure 5 shows the accuracy per category for two of the best models on this dataset, o3 and Gemini 2.5-pro. For o3, the subcategories with the highest accuracy are *Representation Theory (math.RT)*, *Number Theory (math.NT)* and *Analysis of PDEs (math.AP)*, each obtaining more than 60% accuracy. The subcategories with the poorest performance are *Optimization and Control (math.OC)*, *Discrete Mathematics (cs.DM)* and *Machine Learning (cs.LG)*, with the last having only 23.8% accuracy.

However, this capability profile is not uniform across models. Gemini 2.5-pro performs best on *Machine Learning (cs.LG)*, *Optimization and Control (math.OC)* and *Probability (math.PR)*, and worst on *Combinatorics (math.CO)*, *Representation Theory (math.RT)*, and *Group Theory (math.GR)*. Notably, the two best categories for Gemini 2.5-pro—*Machine Learning* and *Optimization and Control*—are the ones on which o3 obtained the lowest scores!

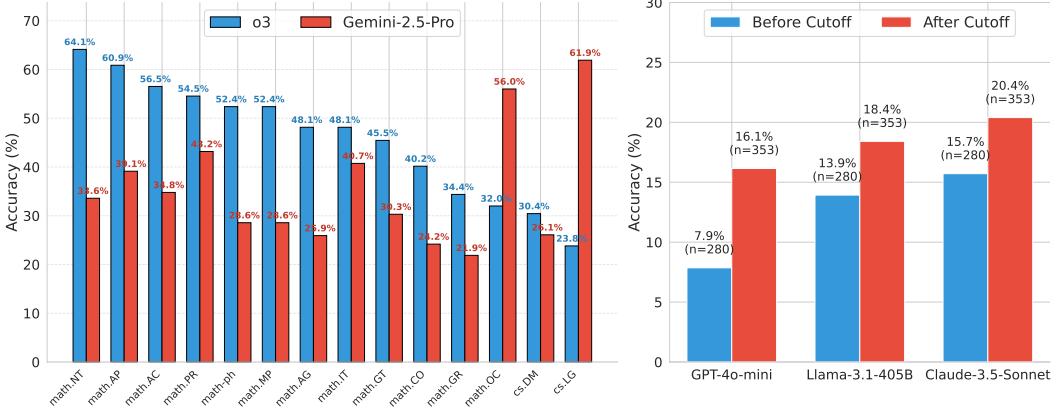


Figure 5: Evaluation of model performance on the Math.arXiv dataset across domain-specific and temporal dimensions. **Left:** Accuracy across mathematical domains for two models, showing significant variation by topic. **Right:** Accuracy of GPT-4o-mini, Llama-3.1-405B, and Claude-3.5-Sonnet on questions published before vs. after their training cutoff.

Overall, we observe that o3 performs well on highly theoretical tasks, such as the ones from *Representation Theory*, whereas Gemini 2.5-pro is best in domains with more practical applicability, such as in machine learning or optimization.

4.3 Additional Analysis

Analyzing LLM’s mistakes. To better assess model performance, we introduce a pipeline that compares the LLMs’ output with a human-authored solution, taken from the original paper or verified StackExchange answer. These are provided to the judge model o3-mini, which evaluates whether the LLMs’ reasoning aligns with the human solution. Common failures include arithmetic errors, incorrect arguments, missing critical insights, or logical inconsistencies. In Figure 6 we show the error breakdown for o3, Gemini 2.5-pro and DeepSeek R1 on the Math.arXiv dataset. The majority of errors come from flawed reasoning, followed by conceptual misunderstanding and missed insights. This suggests that even the best-performing models struggle with multistep logical coherence in a solution.

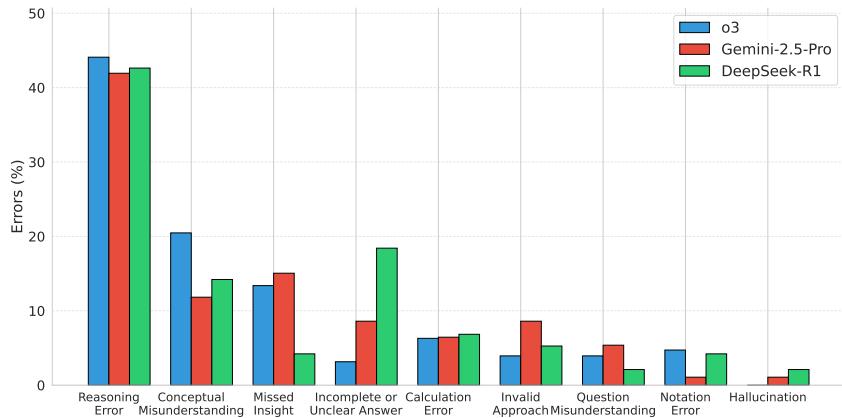


Figure 6: Error types for o3, Gemini 2.5-pro and DeepSeek R1 on the Math.arXiv dataset.

How important is a theorem’s context? Recall that for questions in Math.arXiv and CS.arXiv, we provide the LLM with the full context of the theorem, namely the entire research paper preceding the theorem statement. This context may be critical to understanding the notation and statement of the theorem. However, surprisingly, we find that LLMs still perform rather well even when this context is omitted. This suggests that many theorems use concepts or notation that can be understood or inferred in isolation. For example, in the CS.arXiv dataset, the o4-mini model achieves an accuracy

Question 1: Let \mathcal{F} be an $n \times m$ Ferrers diagram with $m \geq n$, and let $1 \leq d \leq n$ be an integer. Assume that the pair (\mathcal{F}, d) is MDS-constructible and let $\kappa = \kappa(\mathcal{F}, d)$. What is the value of $\lim_{q \rightarrow +\infty} \delta_q(\mathcal{F}, \kappa, d)$?

Question 2: Given parameters σ_1^2 , σ_2^2 , and $c \leq c(\sigma_1^2, \sigma_2^2)$, what is the value of $\lim_{n \rightarrow \infty} \frac{C_s(\sigma_1^2, \sigma_2^2, c\sqrt{n}, n)}{n}$ expressed in terms of σ_1^2 , σ_2^2 , and c ?

Question 3: What is the value of $\lim_{\text{odd } k \rightarrow \infty} \frac{\alpha(f_{\{(k+1)/2\}, k})}{\rho(f_{\{(k+1)/2\}, k})}$?

Figure 7: Questions solvable by LLMs without any context. Interestingly, LLMs can infer the meaning of C_s in Question 2 and of α, ρ, f in Question 3. Appendix B.5 contains examples of questions LLMs cannot solve without context.

of 21.6% despite the absence of any context (compared to 42.3% when the full context is provided). We present some examples of questions solvable without context in Figure 7.

Measuring the impact of data contamination. One of the core features of REALMATH is that it can be continuously updated with new organic samples, keeping the benchmark up to date with contemporary mathematics and preventing data contamination. This characteristic distinguishes REALMATH from other research-level benchmarks, which either require extensive manual curation by experts [9, 16], or which can only be extended by a few new samples per year [2, 19].

To test for data contamination, we split the Math.arXiv dataset into two batches, according to the cutoff dates of GPT-4o-mini, Llama-3.1-405B, and Claude-3.5-Sonnet. Figure 5 shows the accuracy for papers published before and after the respective cutoff dates (i.e., papers from 2022 vs. 2025). Interestingly, we find that models perform *better* on newer samples from 2025 than on older samples from 2022. One explanation could be that mathematical research became significantly easier (at least for LLMs) from 2022 to 2025, which seems unlikely. Rather, we believe a likely explanation is that current LLMs have been trained on a large amount of math data collected close to their cutoff date, and thus that LLMs may be more familiar with the type of topics studied in current math research.

We note that a nice application of REALMATH would be to study the evolution of mathematical research (and LLMs’ performance) over time, an analysis which we leave for future work.

Fine-tuning impact analysis. To evaluate whether REALMATH can be used to improve mathematical capabilities of LLMs, we fine-tuned GPT-4o-mini on 500 random samples from the Math.arXiv dataset and evaluated its performance on the remaining 133 samples. Surprisingly, fine-tuning did **not** lead to an improvement in accuracy (see Appendix B.1 for more details). This suggests that the benchmark’s difficulty does not stem from questions being out-of-distribution for current LLMs. Rather, current models appear to lack the specialized mathematical knowledge or skills necessary to solve certain problems, and these cannot be effectively acquired through simple fine-tuning.

5 Conclusion

Our work introduces REALMATH, a novel benchmark for evaluating LLMs in research-level mathematics. We show that current frontier models demonstrate surprisingly strong capabilities in research mathematics, with leading models achieving accuracy rates of 43–49% on Math.arXiv papers and up to 70% on Math.StackExchange questions. This suggests that these models may already serve as valuable assistants in mathematical research contexts, even as they continue to struggle with the most challenging problems.

A core feature of REALMATH is that it is not a *fixed* benchmark, but rather a refreshable data collection pipeline that reflects current mathematical practice while protecting against data contamination. We believe REALMATH can serve as a valuable tool for both evaluating and improving LLMs’ mathematical capabilities. Its design aligns with the actual needs of mathematicians, helping advance AI assistants for research. As models improve on our benchmark, LLMs may see broader use as collaborative tools in mathematical research and education.

6 Acknowledgment

J.Z. is funded by the Swiss National Science Foundation (SNSF) project grant 214838. K. N. is supported by an ETH AI Center Doctoral Fellowship. We sincerely thank Yani Zhang and Meng Ding for their insightful discussions during the early stages of our data collection pipeline. We also thank Javier Rando and Daniel Paleka for their help in ideating and experimenting with different approaches for classifying problem difficulty levels. Finally, we are grateful to the anonymous reviewers for their valuable feedback and constructive suggestions.

References

- [1] Kassie Archer, Aaron Geary, and Robert P. Laudone. Pattern-avoiding shallow permutations, 2025.
- [2] Mislav Balunović, Jasper Dekoninck, Nikola Jovanović, Ivo Petrov, and Martin Vechev. Math-construct: Challenging llm reasoning with constructive proofs. *arXiv preprint arXiv:2502.10197*, 2025.
- [3] Anwita Bhowmik and Rupam Barman. Number of complete subgraphs of peisert graphs and finite field hypergeometric functions. *Research in Number Theory*, 10(2):26, 2024.
- [4] Nicholas Carlini, Javier Rando, Edoardo Debenedetti, Milad Nasr, and Florian Tramèr. Autoad- vexbench: Benchmarking autonomous exploitation of adversarial example defenses, 2025.
- [5] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021.
- [6] Richard Fang, Rohan Bindu, Akul Gupta, and Daniel Kang. Llm agents can autonomously exploit one-day vulnerabilities, 2024.
- [7] Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Petersen, and Julius Berner. Mathematical capabilities of chatgpt. *Advances in neural information processing systems*, 36:27699–27744, 2023.
- [8] Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan, Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. Omni-math: A universal olympiad level mathematic benchmark for large language models. *CoRR*, abs/2410.07985, 2024.
- [9] Elliot Glazer, Ege Erdil, Tamay Besiroglu, Diego Chicharro, Evan Chen, Alex Gunning, Caroline Falkman Olsson, Jean-Stanislas Denain, Anson Ho, Emily de Oliveira Santos, Olli Järvinieniemi, Matthew Barnett, Robert Sandler, Matej Vrzala, Jaime Sevilla, Qiuyu Ren, Elizabeth Pratt, Lionel Levine, Grant Barkley, Natalie Stewart, Bogdan Grechuk, Tetiana Grechuk, Shreepranav Varma Enugandla, and Mark Wildon. Frontiermath: A benchmark for evaluating advanced mathematical reasoning in ai, 2024.
- [10] Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *CoRR*, abs/2402.14008, 2024.
- [11] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Proceedings of the NeurIPS Datasets and Benchmarks Track*, 2021.
- [12] Yiming Huang, Zhenghao Lin, Xiao Liu, Yeyun Gong, Shuai Lu, Fangyu Lei, Yaobo Liang, Yelong Shen, Chen Lin, Nan Duan, and Weizhu Chen. Competition-level problems are effective llm evaluators. *CoRR*, abs/2312.02143, 2023.

- [13] 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? *CoRR*, abs/2310.06770, 2023.
- [14] Samuel Miserendino, Michele Wang, Tejal Patwardhan, and Johannes Heidecke. Swe-lancer: Can frontier llms earn 1 million from real-world freelance software engineering?, 2025.
- [15] Ivo Petrov, Jasper Dekoninck, Lyuben Baltadzhiev, Maria Drencheva, Kristian Minchev, Mislav Balunović, Nikola Jovanović, and Martin Vechev. Proof or bluff? evaluating LLMs on 2025 USA math olympiad. *arXiv preprint arXiv:2503.21934*, 2025.
- [16] Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Sean Shi, Michael Choi, Anish Agrawal, Arnav Chopra, Adam Khoja, Ryan Kim, Jason Hausenloy, Oliver Zhang, Mantas Mazeika, Daron Anderson, Tung Nguyen, Mobeen Mahmood, Fiona Feng, Steven Y. Feng, Haoran Zhao, Michael Yu, Varun Gangal, Chelsea Zou, Zihan Wang, Jessica P. Wang, Pawan Kumar, Oleksandr Pokutnyi, Robert Gerbicz, Serguei Popov, John-Clark Levin, Mstyslav Kazakov, Johannes Schmitt, Geoff Galgon, Alvaro Sanchez, Yongki Lee, Will Yeadon, Scott Sauers, Marc Roth, Chidozie Agu, Søren Riis, Fabian Giska, Saiteja Utpala, Zachary Giboney, Gashaw M. Goshu, Joan of Arc Xavier, Sarah-Jane Crowson, Mohinder Maheshbhai Naiya, Noah Burns, Lennart Finke, Zerui Cheng, Hyunwoo Park, Francesco Fournier-Facio, John Wydallis, Mark Nandor, Ankit Singh, Tim Gehrunger, Jiaqi Cai, Ben McCarty, Darling Duclosel, Jungbae Nam, Jennifer Zampese, Ryan G. Hoerr, Aras Bacho, Gautier Abou Loumè, Abdallah Galal, Hangrui Cao, Alexis C. Garretson, Damien Sileo, Qiuyu Ren, Doru Cojoc, Pavel Arkhipov, Usman Qazi, Lianghui Li, Sumeet Motwani, Christian Schröder de Witt, Edwin Taylor, Johannes Veith, Eric Singer, Taylor D. Hartman, Paolo Rissone, Jaehyeok Jin, Jack Wei Lun Shi, Chris G. Willcocks, Joshua Robinson, Aleksandar Mikov, Ameya Prabhu, Longke Tang, Xavier Alapont, Justine Leon Uro, Kevin Zhou, Emily de Oliveira Santos, Andrey Pupasov Maksimov, Edward Vendrow, Kengo Zenitani, Julien Guillod, Yuqi Li, Joshua Vendrow, Vladyslav Kuchkin, and Ng Ze-An. Humanity’s last exam. *CoRR*, abs/2501.14249, January 2025.
- [17] Tomohiro Sawada, Daniel Paleka, Alexander Havrilla, Pranav Tadepalli, Paula Vidas, Alexander Perikles Kranias, John J Nay, Kshitij Gupta, and Aran Komatsuzaki. ARB: Advanced reasoning benchmark for large language models, 2024.
- [18] Minghao Shao, Sofija Jancheska, Meet Udeshi, Brendan Dolan-Gavitt, Haoran Xi, Kimberly Milner, Boyuan Chen, Max Yin, Siddharth Garg, Prashanth Krishnamurthy, Farshad Khorrami, Ramesh Karri, and Muhammad Shafique. Nyu ctf bench: A scalable open-source benchmark dataset for evaluating llms in offensive security. In *NeurIPS*, 2024.
- [19] Haoxiang Sun, Yingqian Min, Zhipeng Chen, Wayne Xin Zhao, Zheng Liu, Zhongyuan Wang, Lei Fang, and Ji-Rong Wen. Challenging the boundaries of reasoning: An olympiad-level math benchmark for large language models, 2025.
- [20] George Tsoukalas, Jasper Lee, John Jennings, Jimmy Xin, Michelle Ding, Michael Jennings, Amitayush Thakur, and Swarat Chaudhuri. Putnambench: A multilingual competition-mathematics benchmark for formal theorem-proving. In *AI for Math Workshop @ ICML 2024*, 2024.
- [21] Mark Vero, Niels Mündler, Victor Chibotaru, Veselin Raychev, Maximilian Baader, Nikola Jovanović, Jingxuan He, and Martin Vechev. Baxbench: Can llms generate correct and secure backends?, 2025.
- [22] Elizaveta Vishnyakova and Mikhail Borovoi. Automorphisms and real structures for a π -symmetric super-grassmannian. *Journal of Algebra*, 644:232–286, 2024.
- [23] Kaiyu Yang and Jia Deng. Learning to prove theorems via interacting with proof assistants. *CoRR*, abs/1905.09381, 2019.
- [24] Kaiyu Yang, Aidan M Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan Prenger, and Anima Anandkumar. Leandojo: Theorem proving with retrieval-augmented language models. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.

[25] Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. minif2f: a cross-system benchmark for formal olympiad-level mathematics. In *International Conference on Learning Representations*, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction describe our novel REALMATH benchmark for mathematical challenges encountered in research communities. We motivate the importance of such a benchmark and the main steps for the sample collection.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discussed the limitations in the quality of the source data and the performance of current frontier models, along with the reliability of LLMs as judges.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper provides empirical analysis and does not include theoretical results or other sorts of formal proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper describes all the used steps in the pipeline for dataset creation, filtering and model evaluation within Section 3. Moreover, we describe the setup and results in Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide access to all three datasets and the code used for their retrieval, as well as the prompts given to the LLMs for filtering and evaluation.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specify all the test details for comprehending the results in Section 3 and Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The benchmark results we provide are mainly to illustrate the difficulty composition of our dataset, and the performance trends of current models.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: Our experiments all involve querying APIs, which do not require computational resources. However, querying APIs comes with a cost.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We use publicly available datasets and we comply with all the specifications in the Code of Ethics of NeurIPS.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the importance of having the REALMATH benchmark for research-level mathematics and how it rigorously assesses how much LLMs can contribute to the research communities nowadays. This evaluation of large language models has impact for further improvements in the models' architecture or testing sets.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not use models or datasets with high risk of misuses.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Yes, we cited related work and papers.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide a detailed readme.md file.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This empirical study does not include any human subjects. We rely solely on samples from arXiv and Mathematics Stack Exchange forum, without any human intervention for any experimental part.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This empirical study does not include any human subjects. We rely solely on samples from arXiv and Mathematics Stack Exchange forum, without any human intervention for any experimental part.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [\[Yes\]](#)

Justification: We specify all the LLMs models used throughout the paper for sample collection, filtering, and evaluation steps. The prompts for LLMs are specified in Appendix A.1.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Data Collection and Evaluation

A.1 Prompts for Judge Models

We used the following two system prompts fed to an LLM (o3-mini) for: (1) filtering theorems in order to retain only the ones with unique answers; (2) generating high-quality question-answer pairs from theorems.

```
1 SYSTEM_PROMPT_THEOREM = r"""
2 You are an expert in mathematics and computer science. Your task is to
3 verify if a theorem has a single, numerical answer, easy to be
4 verified. The theorems should be at least graduate level.
5
6 The theorems should have a fixed numerical answer, not an
7 approximation. Some common examples:
8 - Necessary and Sufficient Conditions: e.g., "X holds if and only if
9 condition A holds" only when at least one of A and X is specific,
10 numerical quantity. We want results of the form "If condition A
11 holds, then condition X holds" ONLY WHEN X is a NUMERICAL VALUE.
12 We don't want "if some conditions are met, then the quantity
13 satisfies a particular equation, then we can get X" when X is not
14 a strict numerical value relation, because this does not have
15 fixed unique solutions. Please be very strict about these rules!
16 - Existence and Uniqueness Theorems: e.g., "There exists a unique X
17 that satisfies A.", but we don't want "There exists an X that
18 satisfies A", because the latter is not a fixed unique solution.
19 - Exact Formula Calculations: e.g., "The answer of formula (1) is 10",
20 or "The solution for formula (1) is X", then both are fixed
21 unique solutions.
22 - Unique Maximum/Minimum Points: e.g., "The maximum value of function
23 f is 10 at point x=1", but we don't want "The maximum value of
24 function f is at least 10", because the latter is not a fixed
25 unique solution.
26 - Exact Complexity Results in Computational Complexity: e.g., "The
27 time complexity of algorithm A is exactly $\\Theta(n^2)$" (not $\\
28 Omega(n^2)$ or $\\Omega(n^2)$, because big-O and big-omega are not exact
29 )."
30 - Explicit number of solutions of an equation: e.g. "X has a unique
31 solution y \\in Y" is accepted even if the numerical value of the
32 number of solutions is not specified because it can trivially be
33 deduced that the number of solutions is 1, which is a fixed answer
34 . We also accept "If X, there are no solutions y \\in Y" (implies 0
35 solutions). BUT we DON'T WANT the previous examples if the set Y
36 in which we look for answers is not clear.
37 - Equality of two numerical equations: e.g.,  $\sum_{k=1}^n k^2 = \frac{n(n-1)}{2}$  because we can assume the numerical fixed answer to be
38 the difference of the 2 which is 0. You MUST include these
39 equalities even if $n$ is not fixed but rather a variable. You
40 MUST also include equations of the form "limit of f(n) = integral
41 of g(x)"
42
43 Some examples of theorems that we don't want:
44 - We DON'T want the theorems that contain if and only if when neither
45 of the sides is numerical ($x \\in T$ does not represent a
46 numerical value), e.g. "A graph is bipartite if and only if it
47 contains no cycles of odd length."
48 - We DON'T want theorems of the type "A holds if and only if there
49 exists x such that X(x) holds", but we DO WANT "A holds if and
50 only if for all x, X(x) holds", where X(x) is a fixed numerical
51 value.
52 - We DON'T want the theorems that have any approximations, or any
53 inequalities, or any other non-deterministic statement. e.g. The
54 theorems for which the main result involves the Big-O notation, or
55 where the main result proven in the theorem is that a certain
```

```

relation holds "if and only if  $n \geq x$  or  $n \leq y$ " MUST be
rejected. We DO NOT consider any theorems where the answer is not
an equality or a fixed answer, i.e. results of the type " $n \geq 5$ "
should NOT be considered, so just SKIP these types of theorems.
17 - We DON'T want the theorems that state " $X \in$  complexity class  $Y$ "
since  $Y$  can belong to a bigger complexity class  $Z$ , so the answer
is not unique.
18 - We DON'T want the theorems that state " $X$  is isomorphic or
homomorphic with  $Y$ ", e.g., Chinese Remainder Theorem.
19
20
21 Important guidelines:
22 - if you cannot find a single, definitive answer, you should return an
empty result
23 - please be very strict about the theorem, if there is any ambiguity,
you should return a "false"
24 - Respond only in the specified JSON format
25
26 return in this exact JSON format:
27 {
28     "single_unique_answer": "true" if the theorem has a single,
definitive answer, otherwise "false"
29     "explanation": "explanation of if this theorem has a single,
definitive answer, otherwise an empty string",
30 }
31 """

```

```

1 SYSTEM_PROMPT_GENERATE_QA = r"""
2 You are a skilled problem setter for graduate-level mathematics and
theoretical computer science. You are provided with a set of
theorems (called theorems_dataset), each of which has already been
verified to contain a single, definitive, and numerical answer.
3
4 Your task is to convert each verified theorem into a precise ***
question-answer (QA) pair**. MAKE SURE TO NOT MENTION THE ANSWER
TO THE QUESTION IN THE QUESTION ITSELF.
5
6 Your outputs must follow these rules:
7 1. The **question** should be a well-posed mathematical or theoretical
problem that is **clearly understandable to a graduate-level
student**. Do not ask questions that are easy to answer without
any mathematical reasoning or easy to guess the answer. **You must
never begin your question with "Prove that"**
8 2. The **question** must be solvable in principle with a unique
numerical or mathematical answer**, based solely on the
information in the theorem.
9 3. The **answer** must be:
10    - Strictly and uniquely determined.
11    - Expressed as a number, closed-form expression, formula.
12 4. DO NOT introduce extra assumptions or background. Use only what is
stated or implied clearly by the theorem.
13 5. If a question naturally follows the structure of an identity (e.g.,
"What is the sum of ...?"), frame it that way.
14 6. All QA pairs must reflect **the exact scope of the theorem**. Do
not generalize, weaken, or strengthen its claim.
15 7. DO NOT generate a QA pair if the theorem is ambiguous. DO NOT
generate a QA pair for theorems where the main result to be proven
is an inequality or a Big-O notation. We MUST NOT include any
kind of inequalities, questions about the lower/upper bounds, or
any asymptotic running time of algorithms, i.e. do not generate QA
pairs for theorems where the main result is of the type " $n \geq
5$ ".
16 8. DO NOT include in the question the answer to it, e.g. if a theorem
states "The limit of  $X$  is equal to  $Y$ ," where  $Y$  is an expression of

```

some parameters defined earlier in the theorem, phrase the question in the manner "What is the limit of X in terms of the given parameters?", with the associated answer "The limit of X is Y". DO NOT formulate questions in the form "Prove the following relation..." since the answer will be already included in the question. Moreover, if the main result of the theorem inquires about the value of a parameter for which a relation holds, do not mention this result in the question itself but rather ask "What is the value of the parameter for which the relation holds?".

17 9. If you have a theorem where it says that a certain equation has a certain number of solutions (single/unique solution, no solutions, an infinity of solutions, etc.) but the actual value of the solution is not given, consider the Question-Answer pair to be of the form "What is the number of solutions to this equation?", i.e. even if the theorem does not have an explicit numerical expression for the answer, you can consider it to be a theorem with a fixed-answer, where the fixed-answer is the number of solutions of the equation. However, if the numerical or closed-form of the solution is mentioned in the theorem statement, it is preferred to formulate the question to be "What is the solution to the following equation...?" rather than to inquire about the number of solutions.

18 10. For theorems of the form "X has a certain property if and only if Y has a certain property", pose the question in such a way so that the answer is the side of the "if and only if" statement which indicates a clear, numerical expression and not an abstract definition, i.e. if a theorem states "X is Pareto optimal if and only if $\Phi(X) = 0$ ", consider the question to be "If X is Pareto optimal, what is the value of $\Phi(X)$?" and the answer to be " $\Phi(X) = 0$ ", since the relation $\Phi(X) = 0$ is a numerical one, whereas the other part of the "if and only if statement", namely "X is Pareto optimal" is an abstract property. Hence, you must NOT consider the question to be "How is X if $\Phi(X) = 0$?" since the answer "X is Pareto optimal" is not a unique one.

19 11. For theorems of the form "The following identity holds: $X = Y$ ", if the identity has both X and Y as mathematical expressions that are neither closed-form, nor fixed numerical values (i.e. if we have equality between two sums with complex formulas rather than an equality of the type " $X=5$ "), do not ask the question "What is the value of X?" and the answer to be "Y", since Y might not be the unique answer to the question. Rather, you must formulate the question to be "What is the value of $X - Y$?" and the fixed, clear answer to be " $X - Y = 0$ ". In this pool of theorems with this explicit QA pair, do not consider theorems where X is assumed to be a limit and Y the value of the limit, since this case should be treated as in Condition 8.

20 12. For theorems of the form "The following identity holds: $X = Y + ct$ ", if the identity has both X and Y as mathematical expressions that are neither closed-form, nor fixed numerical values, you should ask the question "What is the value of $X - Y$?" and the answer should be "ct" and not "What is the value of $X - Y - ct$?"

21 13. For theorems of the form "If X holds, then Y", formulate the QA pair to inquire about what happens with Y when relation X holds, and not under what conditions Y holds (since the condition X might not be unique if the theorem is not of the form if and only if).

22 14. DO NOT generate a QA pair for theorems where the main result is the belonging to a complexity class.

23 15. If the theorem states a result of the type " $Y = |X|$ ", where $|X|$ indicates the cardinality of the set X, formulate the question to be "What is the cardinality of X in terms of ...?" and the answer to be "Y", but not the other way around, i.e. DO NOT state "What is the value of Y?" and the answer to be "cardinality of X", since it does not make sense from a logical point of view to phrase the question about a numerical quantity rather than about the characteristic of a set that you must determine.

```
24
25 Return your output strictly in the following JSON format:
26 {
27     "question": "Clearly stated, unique-answer question derived from
28         the theorem. if the theorem is not good, return an empty
29         string",
30     "answer": "The single, unique, exact answer derived from the
31         theorem. if the theorem is not good, return an empty string",
32     "is_good_theorem": "true" if the theorem is good, otherwise "false
33         "
34 }
35 """
36
```

A.2 Prompts for Evaluation Model

We used the following system prompt and user prompt fed to GPT 4o for evaluating whether an LLM's answer to a mathematical question is correct or not.

```
1 system_prompt = r"""You are an expert mathematician tasked with
2 evaluating the correctness of an answer to a mathematical question
3 . Compare the generated answer to the ground truth answer and
4 determine whether the generated answer is mathematically correct
5 and equivalent to the ground truth.
6
7 Please be very strict and rigorous in your evaluation, mark the answer
8 as incorrect even if it is 80% or 90% correct. Ensure the
9 generated answer can be directly rendered in standard LaTeX
10 without requiring custom command definitions. Be precise and focus
11 on mathematical correctness, not formatting or style differences.
12 Your evaluation should be fair and consider that the same
13 mathematical content can be expressed in different ways."""
14
15 user_prompt = f"""QUESTION: {question} GROUND TRUTH ANSWER: {ground_truth} GENERATED ANSWER: {final_answer} Carefully evaluate
16 whether the generated answer is mathematically correct and
17 equivalent to the ground truth. Your response should only contain
18 a JSON object with the following fields:
19 {{{
20     "is_correct": boolean,
21     "explanation": "A concise explanation of why the answer is correct
22                 or incorrect, in a clean LaTeX format"
23 }}}
24 where is_correct is true if the answer is mathematically correct and
25 equivalent to the ground truth, and false if it isn't.
26 """
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
259
260
261
262
263
264
265
266
267
268
269
269
270
271
272
273
274
275
276
277
278
279
279
280
281
282
283
284
285
286
287
287
288
289
289
290
291
292
293
294
295
296
297
297
298
299
299
300
301
302
303
304
305
306
307
308
309
309
310
311
312
313
314
315
316
317
318
319
319
320
321
322
323
324
325
326
327
328
329
329
330
331
332
333
334
335
336
337
338
339
339
340
341
342
343
344
345
346
347
348
349
349
350
351
352
353
354
355
356
357
358
359
359
360
361
362
363
364
365
366
367
368
369
369
370
371
372
373
374
375
376
377
378
379
379
380
381
382
383
384
385
386
387
388
389
389
390
391
392
393
394
395
396
397
398
399
399
400
401
402
403
404
405
406
407
408
409
409
410
411
412
413
414
415
416
417
418
419
419
420
421
422
423
424
425
426
427
428
429
429
430
431
432
433
434
435
436
437
438
439
439
440
441
442
443
444
445
446
447
448
449
449
450
451
452
453
454
455
456
457
458
459
459
460
461
462
463
464
465
466
467
468
469
469
470
471
472
473
474
475
476
477
478
479
479
480
481
482
483
484
485
486
487
488
489
489
490
491
492
493
494
495
496
497
497
498
499
499
500
501
502
503
504
505
506
507
508
509
509
510
511
512
513
514
515
516
517
518
519
519
520
521
522
523
524
525
526
527
528
529
529
530
531
532
533
534
535
536
537
538
539
539
540
541
542
543
544
545
546
547
548
549
549
550
551
552
553
554
555
556
557
558
559
559
560
561
562
563
564
565
566
567
568
569
569
570
571
572
573
574
575
576
577
578
579
579
580
581
582
583
584
585
586
587
588
589
589
590
591
592
593
594
595
596
597
597
598
599
599
600
601
602
603
604
605
606
607
608
609
609
610
611
612
613
614
615
616
617
618
619
619
620
621
622
623
624
625
626
627
628
629
629
630
631
632
633
634
635
636
637
638
639
639
640
641
642
643
644
645
646
647
648
649
649
650
651
652
653
654
655
656
657
658
659
659
660
661
662
663
664
665
666
667
668
669
669
670
671
672
673
674
675
676
677
678
679
679
680
681
682
683
684
685
686
687
688
689
689
690
691
692
693
694
695
696
697
697
698
699
699
700
701
702
703
704
705
706
707
708
709
709
710
711
712
713
714
715
716
717
718
719
719
720
721
722
723
724
725
726
727
728
729
729
730
731
732
733
734
735
736
737
738
739
739
740
741
742
743
744
745
746
747
748
749
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
788
789
789
790
791
792
793
794
795
796
797
797
798
799
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
887
888
889
889
890
891
892
893
894
895
896
897
897
898
899
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
939
940
941
942
943
944
945
946
947
948
949
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
987
988
989
989
990
991
992
993
994
995
995
996
997
997
998
999
999
1000
1000
1001
1001
1002
1002
1003
1003
1004
1004
1005
1005
1006
1006
1007
1007
1008
1008
1009
1009
1010
1010
1011
1011
1012
1012
1013
1013
1014
1014
1015
1015
1016
1016
1017
1017
1018
1018
1019
1019
1020
1020
1021
1021
1022
1022
1023
1023
1024
1024
1025
1025
1026
1026
1027
1027
1028
1028
1029
1029
1030
1030
1031
1031
1032
1032
1033
1033
1034
1034
1035
1035
1036
1036
1037
1037
1038
1038
1039
1039
1040
1040
1041
1041
1042
1042
1043
1043
1044
1044
1045
1045
1046
1046
1047
1047
1048
1048
1049
1049
1050
1050
1051
1051
1052
1052
1053
1053
1054
1054
1055
1055
1056
1056
1057
1057
1058
1058
1059
1059
1060
1060
1061
1061
1062
1062
1063
1063
1064
1064
1065
1065
1066
1066
1067
1067
1068
1068
1069
1069
1070
1070
1071
1071
1072
1072
1073
1073
1074
1074
1075
1075
1076
1076
1077
1077
1078
1078
1079
1079
1080
1080
1081
1081
1082
1082
1083
1083
1084
1084
1085
1085
1086
1086
1087
1087
1088
1088
1089
1089
1090
1090
1091
1091
1092
1092
1093
1093
1094
1094
1095
1095
1096
1096
1097
1097
1098
1098
1099
1099
1100
1100
1101
1101
1102
1102
1103
1103
1104
1104
1105
1105
1106
1106
1107
1107
1108
1108
1109
1109
1110
1110
1111
1111
1112
1112
1113
1113
1114
1114
1115
1115
1116
1116
1117
1117
1118
1118
1119
1119
1120
1120
1121
1121
1122
1122
1123
1123
1124
1124
1125
1125
1126
1126
1127
1127
1128
1128
1129
1129
1130
1130
1131
1131
1132
1132
1133
1133
1134
1134
1135
1135
1136
1136
1137
1137
1138
1138
1139
1139
1140
1140
1141
1141
1142
1142
1143
1143
1144
1144
1145
1145
1146
1146
1147
1147
1148
1148
1149
1149
1150
1150
1151
1151
1152
1152
1153
1153
1154
1154
1155
1155
1156
1156
1157
1157
1158
1158
1159
1159
1160
1160
1161
1161
1162
1162
1163
1163
1164
1164
1165
1165
1166
1166
1167
1167
1168
1168
1169
1169
1170
1170
1171
1171
1172
1172
1173
1173
1174
1174
1175
1175
1176
1176
1177
1177
1178
1178
1179
1179
1180
1180
1181
1181
1182
1182
1183
1183
1184
1184
1185
1185
1186
1186
1187
1187
1188
1188
1189
1189
1190
1190
1191
1191
1192
1192
1193
1193
1194
1194
1195
1195
1196
1196
1197
1197
1198
1198
1199
1199
1200
1200
1201
1201
1202
1202
1203
1203
1204
1204
1205
1205
1206
1206
1207
1207
1208
1208
1209
1209
1210
1210
1211
1211
1212
1212
1213
1213
1214
1214
1215
1215
1216
1216
1217
1217
1218
1218
1219
1219
1220
1220
1221
1221
1222
1222
1223
1223
1224
1224
1225
1225
1226
1226
1227
1227
1228
1228
1229
1229
1230
1230
1231
1231
1232
1232
1233
1233
1234
1234
1235
1235
1236
1236
1237
1237
1238
1238
1239
1239
1240
1240
1241
1241
1242
1242
1243
1243
1244
1244
1245
1245
1246
1246
1247
1247
1248
1248
1249
1249
1250
1250
1251
1251
1252
1252
1253
1253
1254
1254
1255
1255
1256
1256
1257
1257
1258
1258
1259
1259
1260
1260
1261
1261
1262
1262
1263
1263
1264
1264
1265
1265
1266
1266
1267
1267
1268
1268
1269
1269
1270
1270
1271
1271
1272
1272
1273
1273
1274
1274
1275
1275
1276
1276
1277
1277
1278
1278
1279
1279
1280
1280
1281
1281
1282
1282
1283
1283
1284
1284
1285
1285
1286
1286
1287
1287
1288
1288
1289
1289
1290
1290
1291
1291
1292
1292
1293
1293
1294
1294
1295
1295
1296
1296
1297
1297
1298
1298
1299
1299
1300
1300
1301
1301
1302
1302
1303
1303
1304
1304
1305
1305
1306
1306
1307
1307
1308
1308
1309
1309
1310
1310
1311
1311
1312
1312
1313
1313
1314
1314
1315
1315
1316
1316
1317
1317
1318
1318
1319
1319
1320
1320
1321
1321
1322
1322
1323
1323
1324
1324
1325
1325
1326
1326
1327
1327
1328
1328
1329
1329
1330
1330
1331
1331
1332
1332
1333
1333
1334
1334
1335
1335
1336
1336
1337
1337
1338
1338
1339
1339
1340
1340
1341
1341
1342
1342
1343
1343
1344
1344
1345
1345
1346
1346
1347
1347
1348
1348
1349
1349
1350
1350
1351
1351
1352
1352
1353
1353
1354
1354
1355
1355
1356
1356
1357
1357
1358
1358
1359
1359
1360
1360
1361
1361
1362
1362
1363
1363
1364
1364
1365
1365
1366
1366
1367
1367
1368
1368
1369
1369
1370
1370
1371
1371
1372
1372
1373
1373
1374
1374
1375
1375
1376
1376
1377
1377
1378
1378
1379
1379
1380
1380
1381
1381
1382
1382
1383
1383
1384
1384
1385
1385
1386
1386
1387
1387
1388
1388
1389
1389
1390
1390
1391
1391
1392
1392
1393
1393
1394
1394
1395
1395
1396
1396
1397
1397
1398
1398
1399
1399
1400
1400
1401
1401
1402
1402
1403
1403
1404
1404
1405
1405
1406
1406
1407
1407
1408
1408
1409
1409
1410
1410
1411
1411
1412
1412
1413
1413
1414
1414
1415
1415
1416
1416
1417
1417
1418
1418
1419
1419
1420
1420
1421
1421
1422
1422
1423
1423
1424
1424
1425
1425
1426
1426
1427
1427
1428
1428
1429
1429
1430
1430
1431
1431
1432
1432
1433
1433
1434
1434
1435
1435
1436
1436
1437
1437
1438
1438
1439
1439
1440
1440
1441
1441
1442
1442
1443
1443
1444
1444
1445
1445
1446
1446
1447
1447
1448
1448
1449
1449
1450
1450
1451
1451
1452
1452
1453
1453
1454
1454
1455
1455
1456
1456
1457
1457
1458
1458
1459
1459
1460
1460
1461
1461
1462
1462
1463
1463
1464
1464
1465
1465
1466
1466
1467
1467
1468
1468
1469
1469
1470
1470
1471
1471
1472
1472
1473
1473
1474
1474
1475
1475
1476
1476
1477
1477
1478
1478
1479
1479
1480
1480
1481
1481
1482
1482
1483
1483
1484
1484
1485
1485
1486
1486
1487
1487
1488
1488
1489
1489
1490
1490
1491
1491
1492
1492
1493
1493
1494
1494
1495
1495
1496
1496
1497
1497
1498
1498
1499
1499
1500
1500
1501
1501
1502
1502
1503
1503
1504
1504
1505
1505
1506
1506
1507
1507
1508
1508
1509
1509
1510
1510
1511
1511
1512
1512
1513
1513
1514
1514
1515
1515
1516
1516
1517
1517
1518
1518
1519
1519
1520
1520
1521
1521
1522
1522
1523
1523
1524
1524
1525
1525
1526
1526
1527
1527
1528
1528
1529
1529
1530
1530
1531
1531
1532
1532
1533
1533
1534
1534
1535
1535
1536
1536
1537
1537
1538
1538
1539
1539
1540
1540
1541
1541
1542
1542
1543
1543
1544
1544
1545
1545
1546
1546
1547
1547
1548
1548
1549
1549
1550
1550
1551
1551
1552
1552
1553
1553
1554
1554
1555
1555
1556
1556
1557
1557
1558
1558
1559
1559
1560
1560
1561
1561
1562
1562
1563
1563
1564
1564
1565
1565
1566
1566
1567
1567
1568
1568
1569
1569
1570
1570
1571
1571
1572
1572
1573
1573
1574
1574
1575
1575
1576
1576
1577
1577
1578
1578
1579
1579
1580
1580
1581
1581
1582
1582
1583
1583
1584
1584
1585
1585
1586
1586
1587
1587
1588
1588
1589
1589
1590
1590
1591
1591
1592
1592
1593
1593
1594
1594
1595
1595
1596
1596
1597
1597
1598
1598
1599
1599
1600
1600
1601
1601
1602
1602
1603
1603
1604
1604
1605
1605
1606
1606
1607
1607
1608
1608
1609
1609
1610
1610
1611
1611
1612
1612
1613
1613
1614
1614
1615
1615
1616
1616
1617
1617
1618
1618
1619
1619
1620
1620
1621
1621
1622
1622
1623
1623
1624
1624
1625
1625
1626
1626
1627
1627
1628
1628
1629
1629
1630
1630
1631
1631
1632
1632
1633
1633
1634
1634
1635
1635
1636
1636
1637
1637
1638
1638
1639
1639
1640
1640
1641
1641
1642
1642
1643
1643
1644
1644
1645
1645
1646
1646
1647
1647
1648
1648
1649
1649
1650
1650
1651
1651
1652
1652
1653
1653
1654
1654
1655
1655
1656
1656
1657
1657
1658
1658
1659
1659
1660
1660
1661
1661
1662
1662
1663
1663
1664
1664
1665
1665
1666
1666
1667
1667
1668
1668
1669
1669
1670
1670
1671
1671
1672
1672
1673
1673
1674
1674
1675
1675
1676
1676
1677
1677
1678
1678
1679
1679
1680
1680
1681
1681
1682
1682
1683
1683
1684
1684
1685
1685
1686
1686
1687
1687
1688
1688
1689
1689
1690
1690
1691
1691
1692
1692
1693
1693
1694
1694
1695
1695
1696
1696
1697
1697
1698
1698
1699
1699
1700
1700
1701
1701
1702
1702
1703
1703
1704
1704
1705
1705
1706
1706
1707
1707
1708
1708
1709
1709
1710
1710
1711
1711
1712
1712
1713
1713
1714
1714
1715
1715
1716
1716
1717
1717
1718
1718
1719
1719
1720
1720
1721
1721
1722
1722
1723
1723
1724
1724
1725
1725
1726
1726
1727
1727
1728
1728
1729
1729
1730
1730
1731
1731
1732
1732
1733
1733
1734
1734
1735
1735
1736
1736
1737
1737
1738
1738
1739
1739
1740
1740
1741
1741
1742
1742
1743
1743
1744
1744
1745
1745
1746
1746
1747
1747
1748
1748
1749
1749
1750
1750
1751
1751
1752
1752
1753
1753
1754
1754
1755
1755
1756
1756
1757
1757
1758
1758
1759
1759
1760
1760
1761
1761
1762
1762
1763
1763
1764
1764
1765
1765
1766
1766
1767
1767
1768
1768
1769
1769
1770
1770
1771
1771
1772
1772
1773
1773
1774
1774
1775
1775
1776
1776
1777
1777
1778
1778
1779
1779
1780
1780
1781
1781
1782
1782
1783
1783
1784
1784
1785
1785
```

A.3 Manual Review

For the datasets we analyze in Section 4, we rely on a human-based review to ensure the quality of the samples. As noted in Section 3.2, for `Math.arXiv` we removed 5.2% bad samples, whereas for `CS.arXiv` we discarded 7.5% samples. Figure 8 shows the percentages of manually discarded items that fall into each category (35 samples in total for `Math.arXiv` and 9 samples for `CS.arXiv`). These samples were either theorems or QA pairs that did not have a unique answer but which LLMs failed to filter out in the previous stages of the automated pipeline.

A.4 Category Distribution Across Construction Stages

Questions in our datasets are grouped into human-labeled categories based on their domain (e.g., arXiv paper categories). Figure 9 shows how category distribution shifts throughout the dataset construction stages. Combinatorics (math.CO) becomes increasingly dominant in later stages, as

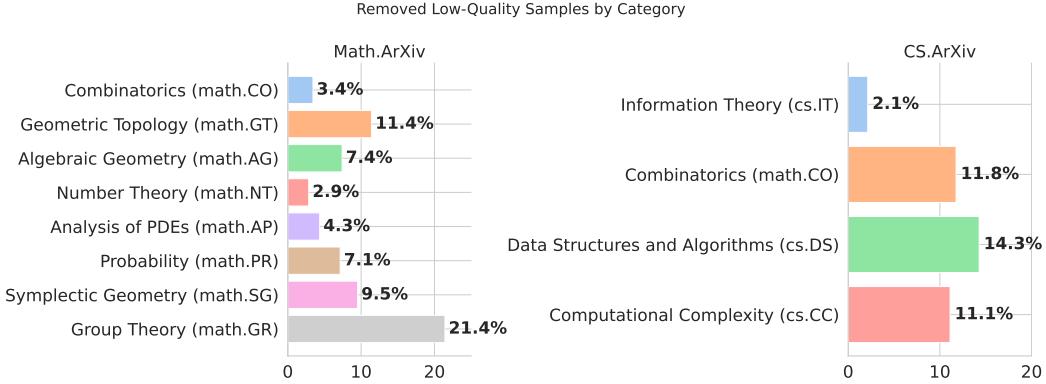


Figure 8: Samples eliminated in the final manual review phase from each category.

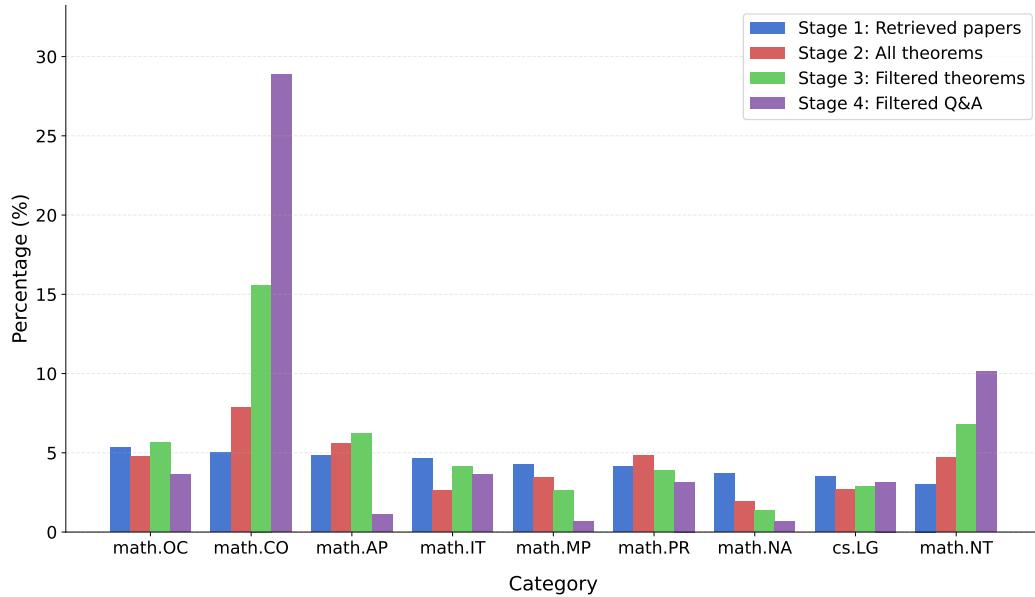


Figure 9: Category distribution across different stages of dataset construction for Math.arXiv (05/2022 – 09/2022). The x-axis shows subject categories (e.g., math.CO, math.NT, etc.), and the y-axis represents the percentage of the dataset each category comprises. Notably, combinatorics (math.CO) becomes increasingly dominant in later stages, particularly in the final Q&A dataset.

these papers tend to contain a large number of theorems, many of which are well-suited for the Q&A format.

A.5 Pipeline for Mathematics Stack Exchange

The pipeline we describe in Section 3 is designed for arXiv papers. To retrieve and process samples from Mathematics Stack Exchange, we need to adapt some steps. We first queried the corresponding API to retrieve users’ questions with tags such as `[limits]`, `[definite-integrals]`, `[integration]`—which were expected to have a clear numerical answer—as well as the top-ranked answer, if available. We extracted the HTML-formatted text and used an LLM to process the user’s post into a statement and proof format, mirroring the theorems extracted from arXiv papers. We then applied a manual filtering step to ensure statements have a fixed answer and that the answer is indeed correct (for this, we checked cited sources, as well as the top-rated answer posted by other users). Finally, we fed the formulated theorem to an LLM to generate a QA pair and filter trivial samples.

A.6 Optional LaTeX Conversion

To improve human readability of our datasets, we incorporated a LaTeX formatting step where we ask a LLM to convert each theorem into a self-contained and valid LaTeX code (this requires, for example, the expansion of any macros). We verified that this step does not negatively impact the language model’s ability to formulate a QA pair or filter out low-quality samples.

B Additional Results

B.1 Finetuning

In Figure 10, we use OpenAI’s fine-tuning service to fine-tune GPT-4o-mini on our dataset. We present the detailed training and test loss in Figure 11. To ensure reliability, we repeat the evaluation five times and report the mean and standard deviation of the results. Surprisingly, we observe no statistically significant improvement.

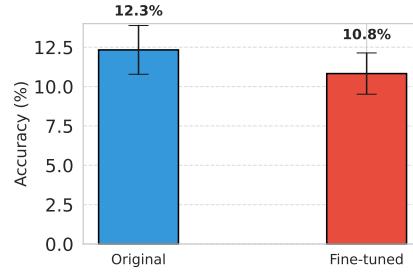


Figure 10: Accuracy of original vs. fine-tuned GPT-4o-mini on Math.arXiv.

Figure 11 shows the training and testing loss curves for OpenAI fine-tuning. The training loss decreases rapidly and remains consistently low, indicating that the model is fitting the training data well. On the other hand, the testing loss exhibits a fluctuating behavior and increases over time compared to the beginning. This pattern may justify the failure of generalization of the model and further reinforces the idea that fine-tuning does not lead to improved performance.

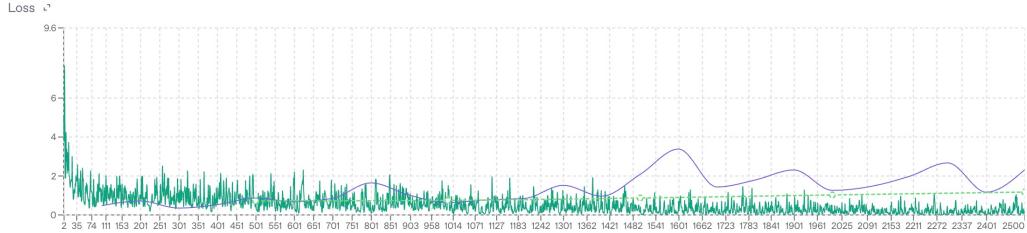


Figure 11: Training (green) and testing (purple) loss curves from the OpenAI fine-tuning service. Fine-tuning on our dataset does not show a clear improvement.

B.2 Difficulty Levels

We assessed the difficulty level of the samples in the datasets by using older, weaker models and checking their performance on the samples. More specifically, we evaluated the performance of four older models from distinct families—Qwen3-235B, Claude 3.5 Haiku, GPT-3.5 Turbo and Llama 3.3-70B.

Samples in the *Hard* category are the ones which have not been answered correctly by any of the aforementioned models, whereas samples in the *Medium* and *Easy* categories are the ones which have been answered correctly by less than 50% of the LLMs (i.e. 1-2 out of 4), and more than 50% of the models (i.e. 3-4 out of 4), respectively.

Figure 12 depicts the distribution of samples across categories in the CS.arXiv data collection, with 69.4% of the samples being considered *Hard*. This distribution suggests that indeed our dataset reflects a hardness level appropriate for the research communities.

Moreover, we show the accuracy per category for some of the best-performing models—o3, DeepSeek-R1, Gemini-2.5-Pro—on CS.arXiv. Each of the three LLMs (except o3) has accuracy level decreasing substantially as the difficulty level increases, e.g., DeepSeek-R1 has 100% accuracy on *Easy* samples and only 11.7% accuracy on *Hard* samples. This is an expected pattern which reinforces the correct assessment of our samples in these three categories.

The higher accuracy of o3 on *Medium* samples when compared to that of *Easy* samples—75% versus 71.4%—does not necessarily imply an abnormal behavior and can be justified by the difference in the number of samples from these two categories, where we have 18% of the samples considered *Medium* difficulty and only 12.6% of them labeled as *Easy*. Especially for the *Hard* category, we spot a discrepancy between the relatively high accuracy of o3, namely 31.2%, compared to DeepSeek-R1 and Gemini-2.5-Pro, who both achieve less than 12% accuracy.

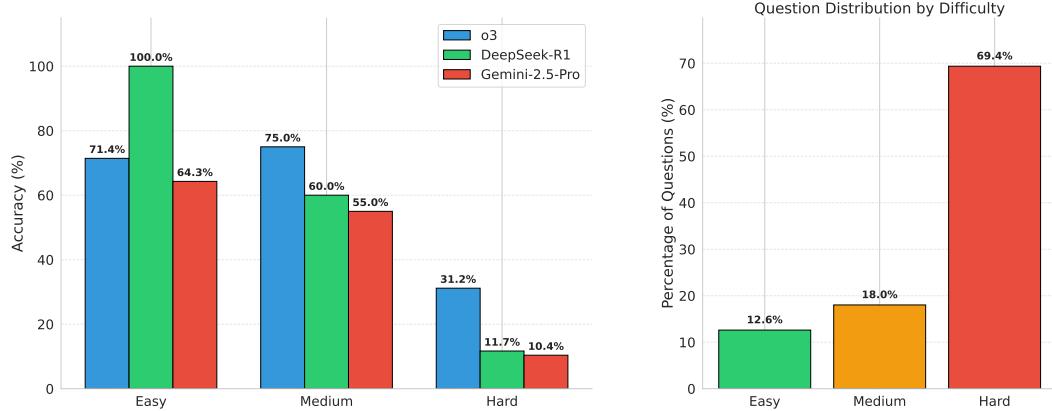


Figure 12: Accuracy across different difficulty levels for best-performing models on CS.arXiv.

Along the same line of thought, Figure 13 shows the distribution for the `Math.StackExchange` dataset. We remark that the distribution for this dataset is more spread, with the *Medium* difficulty questions being the ones that occur the most, as opposed to the `Math.arXiv` and `CS.arXiv` distributions, where the questions were dominantly classified as *Hard*. This reflects the nature of `Math.StackExchange`, which focuses on clarifications of undergraduate and graduate-level concepts and computations, rather than innovative, research-level inquiries.

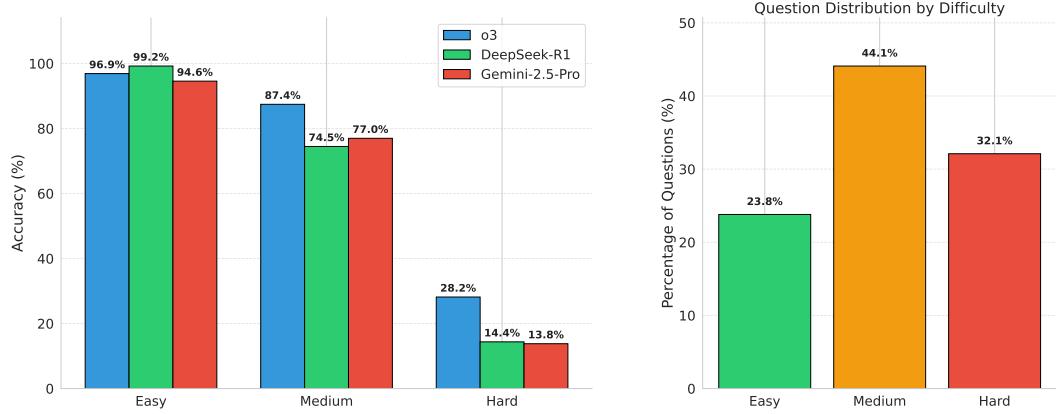


Figure 13: Accuracy across different difficulty levels for best-performing models on `math.stackExchange`.

B.3 Performance Across Categories

We have assessed the performance of the models on each subcategory (for `CS.arXiv`) and on each tag (for `Math.StackExchange`) in order to understand whether some models perform particularly well on a certain category of questions. In Figure 14, we show the accuracies of o3 and Gemini-2.5-Pro.

On one hand, we remark for `CS.arXiv` that, with the exception of *Logic in Computer Science (cs.LO)*, where the models have identical accuracies, o3 outperforms Gemini-2.5-Pro in every category. On the other hand, an important feature that can be observed in Figure 14(a) is that o3 is better at solving questions from applied categories such as *Discrete Mathematics (cs.DM)* and *Machine Learning (cs.LG)*, while Gemini-2.5-Pro gets its highest accuracies on theoretical topics such as *Computational Complexity (cs.CC)* and *Information Theory (cs.IT)*. This is consistent with the behavior that these two LLMs have for the `Math.arXiv` dataset.

For the `Math.StackExchange` dataset in Figure 14(b) we observe as well that o3 outperforms Gemini-2.5-Pro in each category. Moreover, the two LLMs follow a similar trend in the sense that their accuracy level per subcategory can be roughly ordered in the same hierarchy structure. Both models perform best on *Limits*, suggesting this is a relatively easier category where they can get fixed numerical answers correctly with high probability—more than 80%—, whereas their performance decreases on more advanced topics such as Complex Analysis.

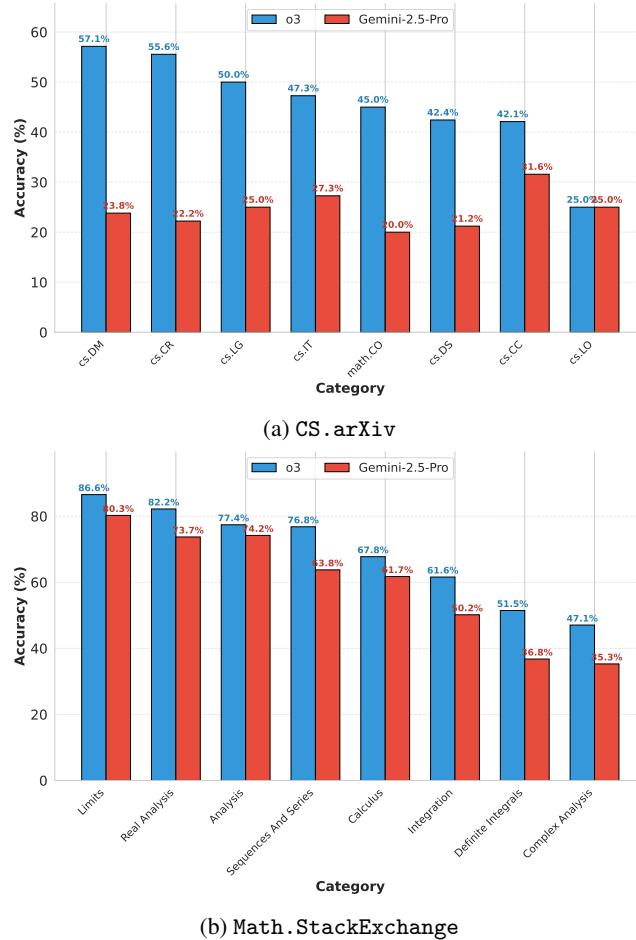


Figure 14: Performance across different math domains in (a) `CS.arXiv` and (b) `Math.StackExchange`.

B.4 Category Distribution of Questions

This subsection provides the statistics of how samples are distributed according to subdomains for each of the three datasets under consideration, clearly presented in the subsequent three tables. The subdomains are a good indicator for highlighting which mathematical domains are best for retrieving theorems along with question-answer pairs with a fixed, unique numerical answer or closed-form expression.

An interesting observation is that *Combinatorics* (*math.CO*) is among the top frequent domains for both *Math.arXiv* and *CS.arXiv* data collections. This suggests that problems in combinatorics are not only well-represented but also more likely to have fixed-answer solutions which satisfy our benchmarking criteria and make the automated evaluation process easy to be verified.

Category	arXiv Tag	Percentage (%)
Combinatorics	math.CO	28.89
Number Theory	math.NT	10.16
Algebraic Geometry	math.AG	5.19
Probability	math.PR	3.16
Geometric Topology	math.GT	4.51
Group Theory	math.GR	3.16
Information Theory	math.IT	7.22
Optimization and Control	math.OC	3.61
Discrete Mathematics	cs.DM	4.51
Analysis of PDEs	math.AP	1.13
Commutative Algebra	math.AC	1.35
Machine Learning	cs.LG	3.16
Representation Theory	math.RT	1.35
Mathematical Physics	math.MP	0.68
Rings and Algebras	math.RA	1.35
Symplectic Geometry	math.SG	3.61
Functional Analysis	math.FA	2.93
Classical Analysis and ODEs	math.CA	2.48
Differential Geometry	math.DG	1.58
Complex Variables	math.CV	1.58

Table 4: Percentage distribution of the top 20 most represented arXiv categories in the *Math.arXiv* (05/2022 – 09/2022) dataset. Categories are not mutually exclusive; papers with multiple tags are counted under each corresponding category. The predominance of questions in combinatorics and number theory is consistent with other recent math benchmarks (e.g., [2, 14]) and likely reflects an abundance of constructive statements in these topics.

Category	arXiv Tag	Percentage (%)
Information Theory	cs.IT	49.55
Combinatorics	math.CO	36.04
Data Structures and Algorithms	cs.DS	29.73
Discrete Mathematics	cs.DM	18.92
Computational Complexity	cs.CC	17.12
Machine Learning	cs.LG	10.81
Cryptography and Security	cs.CR	8.11
Logic in Computer Science	cs.LO	7.21
Computer Science and Game Theory	cs.GT	3.60
Optimization and Control	math.OC	3.60
Formal Languages and Automata Theory	cs.FL	2.70
Logic	math.LO	2.70
Databases	cs.DB	1.80
Algebraic Geometry	math.AG	1.80
Neural and Evolutionary Computing	cs.NE	1.80
Commutative Algebra	math.AC	1.80
Distributed, Parallel, and Cluster Computing	cs.DC	1.80
Systems and Control	cs.SY	0.90
Multiagent Systems	cs.MA	0.90
Artificial Intelligence	cs.AI	0.90
Information Retrieval	cs.IR	0.90
Networking and Internet Architecture	cs.NI	0.90
Numerical Analysis	cs.NA	0.90
Rings and Algebras	math.RA	0.90
Probability	math.PR	0.90

Table 5: Percentage distribution of arXiv categories in the CS.arXiv dataset. Categories are not mutually exclusive; papers with multiple tags are counted under each corresponding category.

Category	Percentage (%)
Integration	53.32
Limits	43.91
Calculus	33.76
Definite Integrals	25.09
Real Analysis	21.77
Sequences and Series	12.73
Complex Analysis	6.27
Analysis	5.72
Improper Integrals	5.35
Solution Verification	5.35
Multivariable Calculus	4.43
Derivatives	4.43
Limits Without Lhopital	4.06
Closed Form	3.32
Trigonometry	2.95
Convergence Divergence	2.77
Probability	2.77
Special Functions	2.77
Trigonometric Integrals	2.58
Contour Integration	2.03
Measure Theory	2.03
Indefinite Integrals	2.03
Complex Integration	1.85
Gamma Function	1.85
Epsilon Delta	1.66
Asymptotics	1.66

Table 6: Percentage distribution of the top 26 most represented tags in the Math.stackExchange dataset. Categories are not mutually exclusive; questions with multiple tags are counted under each corresponding category.

B.5 How Context Improves LLM Performance on Mathematical Problems

Figure 15 depicts some examples of questions from CS . arXiv dataset which are correctly solved if the relevant context is provided to the language model, but which are not answered properly when context is not provided.

Based on a manual verification of the samples, we observed that not providing context negatively impacts models in two directions:

- LLMs cannot answer the question since they do not understand the notations of certain quantities, which are only explained in the paper. In these cases, the context is crucial for the LLM since otherwise it either attempts to guess what a notation means, or they simply do not try to answer and specify that the quantities are not known to it.
- LLMs cannot answer the question because it is considered hard if we do not provide context. This situation comprises cases where the context presents some intermediate results or examples of proofs which are used for building up the current theorem or for its proof. Feeding the context to the language model improves its approach into correctly proving the question.

Question 1: For $r \geq 5$, what is the value of $\text{gcover}(\text{BF}(r))$ expressed in terms of r ?

► *LLMs fail to correctly respond because the notation for the quantity $\text{gcover}(\text{BF}(r))$ is ambiguous without prior definition.*

Question 2: What is the expression for the zero-rate error exponent $E^{(1)}(0)$ in terms of a maximization over $q \in \mathcal{P}(\mathcal{X})$ and the function $d_B(x, x', P)$?

► *LLMs do not have knowledge of the function $d_B(x, x', P)$ or the set $\mathcal{P}(\mathcal{X})$ if the context is not provided.*

Question 3: Let $n > 1$ be an integer and let $\mathcal{A} \subseteq \mathbb{Z}_2^n$ be a maximal anticode of diameter one. What is the maximum possible number of codewords in \mathcal{A} ?

► *The LLM knows what an anticode represents, but the question is too hard to be answered correctly.*

Question 4: Consider the complete bipartite graph $K_{m,n}$ with parameters satisfying $m < n$, n even, and $(m+n)|mn$. What is the ATN of $K_{m,n}$?

► *The notion of ATN (Alon-Tarsi Number) is a common abbreviation for language models, but the difficulty of the question is too high for them to respond correctly.*

Figure 15: Examples of questions in our CS . arXiv dataset that are not solvable by LLMs without the necessary context, but which become solvable when the context is given.