

# UTMath: Math Evaluation with Unit Test via Reasoning-to-Coding Thoughts

Anonymous ACL submission

## Abstract

The evaluation of mathematical reasoning capabilities is essential for advancing Artificial General Intelligence (AGI). While Large Language Models (LLMs) have shown impressive performance in solving mathematical problems, existing benchmarks such as GSM8K and MATH present limitations, including narrow problem definitions with specific numbers and reliance on predetermined rules that hinder accurate assessments of reasoning and generality. This paper introduces the UTMath Benchmark, a robust evaluation framework designed to assess LLMs through extensive unit tests, with a focus on both the accuracy and generality of model responses. The benchmark comprises 1,053 cutting-edge problems spanning nine mathematical domains, with an average of 68 test cases per problem. Furthermore, we present the Reasoning-to-Coding of Thoughts (RCoT) approach, which encourages LLMs to engage in explicit reasoning prior to code generation, thereby facilitating the production of more sophisticated solutions and enhancing overall performance and efficiency. Furthermore, we also release the UTMath-Train training dataset (more than 70k samples), to support the community in further exploring mathematical reasoning. Our benchmark can be accessed via the following link: [UTMath](#)

## 1 Introduction

The pursuit of AGI necessitates strong mathematical reasoning capabilities, making the evaluation of such abilities a crucial area of research (Zhou et al., 2024a). Recent advancements in LLMs have demonstrated remarkable proficiency in solving complex mathematical problems, achieving amazing performance on various datasets of Math Word Problems (MWP), such as GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021), TheoremQA (Chen et al., 2023).

However, classic benchmarks exhibit several limitations that impede the accurate and comprehen-

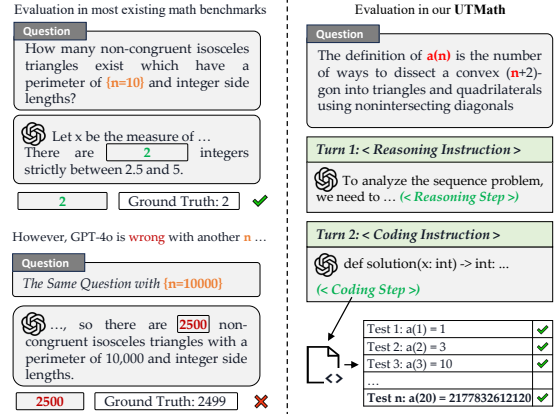


Figure 1: Comparison of UTMath with other benchmarks. On the left, GPT-4o easily solved one question but failed with a different numeric input. On the right, our benchmark is shown, where each problem includes multiple test cases, and a solution is correct only if all test cases are passed. We also propose a new prompting method RCoT in which the LLM first reasons through the problem and then generates code.

sive assessment of these models' capabilities (Ahn et al., 2024). First, these benchmarks test models on narrowly defined problems with some specific numbers, which do not adequately assess generality to similar but varied scenarios as shown in Fig. 1. Second, their evaluation relies on predetermined rules or the method of LLM-as-a-Judge (Dubois et al., 2024; Zheng et al., 2023) that usually failed with capricious responses of LLMs. For example, an accurate answer need to be extracted to exactly match the final answer in the dataset GSM8K, TheoremQA, and MATH dataset. While recent work has made great progress in developing new benchmarks, many of these approaches still fall short of addressing the fundamental limitations of earlier datasets. For instance, benchmarks like GSM-HARD (Gao et al., 2023), GSM-IC (Shi et al., 2023), GSM-Plus (Li et al., 2024a), MetaMath (Yu et al., 2023) have extended the dataset of GSM8K

or MATH with some perturbation such as substitution, reversing, distractor insertion. These efforts, while valuable, are characterized by limited coverage and high costs. In this context, our work seeks to bridge these gaps by proposing a solid and robust benchmark that accurately evaluates the mathematical capabilities of LLMs.

Drawing on evaluation methods from software development, we propose the design of a comprehensive set of unit tests for mathematical problems to rigorously assess the reasoning processes of LLMs. If a solution successfully passes all unit tests within a class of problems, it suggests that the reasoning underlying the solution is more reliable and trustworthy. Specifically, we introduce the **UTMath**, a novel benchmark derived from the On-Line Encyclopedia of Integer Sequences (OEIS) (OEIS Foundation Inc., 2024). The benchmark consists of 1,053 cutting-edge problems spanning 9 mathematical domains, such as Number Theory and Geometry. Each problem is accompanied by more than 68 test cases that provide a set of inputs and their corresponding outputs.

In terms of evaluation methodology, our benchmark requires models to derive a general solution for a class of problems, typically represented in the form of code. Compared to solving a problem defined by specific numbers, developing such a general solution is substantially more challenging, requiring higher levels of intelligence and reasoning ability. However, when we performed “Program of Thoughts (PoT)” (Chen et al., 2022), where the model is required to perform reasoning and coding in a single response, it consistently tends to produce simpler and straightforward solutions. We surmise that this tendency may be influenced by the distribution of coding data. To address this, we introduced the “**Reasoning-to-Coding of Thoughts (RCoT)**”, which requires the LLM to perform mathematical reasoning in the first turn without any coding instruction, then writing code based on the reasoning. Compared to PoT, RCoT shifts the code distribution towards mathematics in the first turn, prompting more reasoning steps.

We conducted a comprehensive study with 8 LLMs. Some of our key findings are summarized as follows: (1) The best model, GPT-4o, only solves 26.93% problem in our benchmark, demonstrate the difficulty of our benchmarks. (2) Modern LLMs perform poorly in Graph Theory, Group Theory, Geometry and Topology (Fig. 5). (3) With RCoT, all evaluated LLMs generated more effi-




Figure 2: RCoT significantly improves the efficiency and effectiveness of the solution. It indicates that our RCoT proves to be more effective, suggesting that it encourages the model to reason critically and find more efficient solutions.

cient solutions, with most models achieving higher scores (Fig. 2). (4) RCoT can significantly improve the pass@k performance of LLMs (§ 5.4). (5) The quality of reasoning significantly impacts the accuracy and efficiency of the model’s final solution (§ 5.5). More interesting findings can be found in § 5. We hope our findings contribute to a deeper understanding of current reasoning ability of LLMs and the further development of models.

## 2 Related Work

### 2.1 Benchmarks

With the rapid development of LLMs, evaluating and exploring the intelligence and limitations of these models has emerged as an urgent issue to address (Chang et al., 2024). Reasoning ability, as a crucial component of general intelligence, has garnered widespread attention since the advent of LLMs (Patel et al., 2021; Cobbe et al., 2021; Valmeekam et al., 2022; Perez et al., 2022; Gupta et al., 2022; Shakarian et al., 2023). Mathematical reasoning, due to its complex mathematical characteristics and rigorous logical relationships, is considered an abstract and high-difficulty task, playing a pivotal role in demonstrating a model’s reasoning capabilities.

To this end, researchers have proposed various benchmarks focused on mathematical reasoning.

A natural and mainstream approach is to evaluate LLMs as humans would take math exams, using human exam questions to test their reasoning abilities, categorized by required knowledge levels. Examples include GSM8K at elementary school level, Math and GaokaoBench-Math (Zhang et al., 2023) at high school level, College Math (Tang et al., 2024), TheoremQA (Chen et al., 2023), ARB (Sawada et al., 2023) at university level, and OlympiadBench (He et al., 2024), AGIeval-Math (Zhong et al., 2023) at competition level.

Besides, researchers have also introduced many others focused on evaluating various aspects of LLMs like the robustness. These include GSM8K-based variants: GSM-8K-Adv (Anantheswaran et al., 2024), GSM-Hard (Gao et al., 2023), GSM-Plus (Li et al., 2024a), GSM-IC (Shi et al., 2023), and several independent benchmarks: RobustMath (Zhou et al., 2024b), MetaMathQA (Yu et al., 2023), PROBLEMATHIC (Anantheswaran et al., 2024), MATHCHECK (Zhou et al., 2024a), as well as other benchmarks (Li et al., 2024b, 2023).

The distinctions between our proposed benchmark and existing ones are as follows. (1) Multiple Case Validation. Instead of using single cases that can be memorized, our questions are sequence-based, allowing numerous cases for validating true understanding. (2) True Reasoning Evaluation. Hard cases and runtime metrics help filter memorization and compare solution efficiency, precisely assessing reasoning abilities.

## 2.2 Building Methods

Constructing effective, high-quality datasets is a complex and labor-intensive process. The advent of LLMs offers an opportunity to change this scenario (Valmeekam et al., 2022; Drori et al., 2023; Perez et al., 2022; Chiang and Lee, 2023; Liu et al., 2023; Fu et al., 2023; Kocmi and Federmann, 2023; Li et al., 2024b). For instance, (Almoubayyed et al., 2023) employed GPT-4o to rewrite mathematics problems based on MATHia (Ritter et al., 2007) to aid students in improving their math performance, validating this with 12,374 students and demonstrating the effectiveness of using LLMs for data construction. These efforts provide a reliable foundation for utilizing LLMs in data processing.

In our study, we utilized GPT-4o to help us deal with data, such as by providing necessary background knowledge for questions and making them more understandable, with more information about the prompts used shown in the Appendix C. Sub-

sequently, human verification was performed to ensure consistency before and after LLM usage.

## 2.3 Prompting Methods

Considering the attributes of large models, they exhibit significant sensitivity to prompts, rendering prompt engineering a critical area of study.

The Chain-of-Thought (Wei et al., 2022) prompting technique encourages models to express reasoning steps in natural language before concluding. Similarly, the approach by (Kojima et al., 2022) uses the phrase "Let's think step by step" to effectively guide large language models through their reasoning. Inspired by CoT, several effective prompting methods have been developed, such as Tree-of-Thoughts (Yao et al., 2024), Graph-of-Thoughts (Besta et al., 2024). Program-of-Thought prompting (Chen et al., 2022): PoT generates programs as the intermediate steps and integrates external tools like a Python interpreter for precise calculations, as well as other prompting methods (Wang et al., 2023; Gao et al., 2023).

Our RCoT method stands out by dividing reasoning into two steps: reasoning and implementing based on reasoning. This segmentation provides deeper insights into LLM reasoning abilities. The advantages can be summarized as follows. (1) Code Output Evaluation. We require LLMs to output code, focusing on reasoning rather than direct answers, to better reflect their reasoning skills. (2) Observation of Reasoning Process. By mandating code implementation, we can validate the LLM's reasoning process, not just the final answer.

## 3 UTMATH Benchmark

### 3.1 Introduction for OEIS.

The OEIS was established to document integer sequences of interest to both professional and amateur mathematicians, and it has become widely cited in the mathematical community. Most sequences are derived or updated from academic papers, contributing to their cutting-edge level of difficulty. As of February 2024, it contains over 370,000 sequences (OEIS Foundation Inc., 2024). Each sequence is accompanied by an identification number, a brief description, some sample integers, generation rules, links to relevant literature, and, where possible, program code for computing the sequences. An example sequence is shown in Appendix A.




Figure 3: UTMATH generation pipeline. After downloading 23,238 Principle Sequences from OEIS and cleaning the data, 1,053 usable sequences were obtained. Descriptions were standardized by adding background information and improving readability (highlighted in green, also shown in Appendix B.2). Hard cases were introduced to enhance discriminative capability, including terms from later positions to prevent simplistic algorithms from passing.

### 3.2 Benchmark Construction.

UTMATH is a rigorous and expansive benchmark designed to more accurately assess the mathematical reasoning abilities of LLMs. It consists of 1053 math problems, with each problem having an average of 68 test cases. The benchmark covers 9 mathematical domains, including not only common topics like number theory but also graph theory, group theory, topology, and geometry. The difficulty of UTMATH is considered Cutting-Edge, with the majority of the sequences that form the problems having been studied in academic papers. UTMATH was obtained as follow (see also Fig.3).

**Data Crawling.** OEIS provides users with a list of principal sequences<sup>1</sup>, which are most important sequences defined by OEIS. OEIS categorizes these sequences into sections based on the first 2-3 letters of their content themes. By scraping the category tags within each section and the AIDs of their subordinate sequences, we obtained 23,238 principal sequences’ AIDs. OEIS provides an interface to request the JSON data of the HTML page for each sequence using its AID<sup>2</sup>. By passing the sequence AIDs to this interface, we acquired the JSON data for these 23,238 sequences.

**Data Cleaning.** We found that some of the sequences we collected did not meet our criteria and should be removed, with further details provided in the Appendix B. Here are several main situations.

- **Hard to solve, few terms are discoverable.**

A portion of the sequences retrieved are marked as “hard” in the keyword field of their entries in OEIS. According to OEIS, “Any sequence which can be extended only by new ideas, rather than more computation deserves keyword: hard. Simi-

larly, if computing a term of the sequence would probably merit a paper in a peer-reviewed journal (discussing the result, the algorithm, etc.)”<sup>3</sup> Another related keyword attribute is “fin” (finite), indicating sequences with limited length. For our purposes, sequences should be infinitely derivable.

- **Difficult to Generate Programmatically.** In OEIS, most sequences are provided with fields such as Mathematica, program, or formula, but not all sequences include these fields. Since OEIS does not specify the reasons for missing these fields, we assume that these sequences might be difficult to generate programmatically.

- **Simple Sequences.** Some sequences are too simple to require any reasoning. We use GPT-4o to determine if a sequence requires reasoning or just implementation; if mostly implementation, it’s excluded. For instance, A000178<sup>4</sup>: ‘Superfactorials: product of the first n factorials,’ a sequence requiring only implementation, will be excluded.

After addressing the aforementioned issues, we ultimately obtained 1053 sequences.

**Standardization of Question Statements.** As a specialized academic database in the field of mathematics, OEIS provides a wealth of useful information for each sequence. However, we have found that some sequences cannot be directly used with the descriptions provided by OEIS as problem statements, primarily for the following reasons: (1) Specialized Terminology. Some sequence descriptions use complex math terms that need examples or explanations to be clear. Using them directly as problems might test mathematical knowledge rather than reasoning skills. So, it is important to explain key concepts to focus on reasoning and

<sup>1</sup>[https://oeis.org/wiki/Index\\_to\\_OEIS](https://oeis.org/wiki/Index_to_OEIS)

<sup>2</sup>[https://oeis.org/wiki/JSON\\_Format](https://oeis.org/wiki/JSON_Format)

<sup>3</sup>[https://oeis.org/wiki/User:Charles\\_R\\_Greathouse\\_IV/Keywords/difficulty](https://oeis.org/wiki/User:Charles_R_Greathouse_IV/Keywords/difficulty)

<sup>4</sup><https://oeis.org/A000178>



Dataset	Size	Level	Multi-test	Efficiency	Metric	Output
College Math	2,818	University	✗	✗	Accuracy	Text
GSM8K	1,319	Elementary school	✗	✗	Accuracy	Text
MATH	5,000	High school	✗	✗	Accuracy	Text
RobustMath	300	High school	✗	✗	Accuracy	Text
OlympiadBench	8,476	Competition	✗	✗	Accuracy	Text
TheoremQA	800	University	✗	✗	Accuracy	Text
UTMath(ours)	1,053	Cutting-edge	✓	✓	Pass Rate	Code

Table 1: Comparison between UTMath and other benchmarks. UTMath offers a cutting-edge benchmark with a comprehensive set of 1,053 problems across multiple mathematical domains, providing a more accurate evaluation of LLMs’ mathematical reasoning capabilities.

reduce the extra knowledge needed. (2) Brevity and Ambiguity. Some sequence descriptions are excessively brief and lack a clear definition of what  $a(n)$  is. We used GPT-4o to standardize these by adding background info without revealing solutions and making the language smoother. The prompts we used are provided in the Appendix C and an example is shown in Appendix B.2. Despite our requirement that ensuring the consistency of meaning between the original and processed descriptions, hallucinations can still occur. To mitigate this, we performed a manual verification of the standardized problem statements to ensure they matched the original descriptions in meaning and were easy to comprehend.

**Hard Test Cases Mining.** The primary goal of this paper is to evaluate the reasoning ability of LLMs. Generally, more efficient solutions to a problem imply stronger reasoning capabilities. Therefore, we aim for our evaluation to distinguish whether a solution is efficient. However, in the OEIS, each sequence only lists the first few  $n$  terms, normally  $n < 100$ , which can be obtained without requiring particularly efficient methods. This limitation prevents the evaluation from effectively distinguishing between efficient and inefficient solutions. An obvious fact is that the difficulty of computing the first 10 terms of a sequence within a time limit is significantly different from computing terms starting from the  $10^6$ th term. Therefore, we aim to create more challenging test data to better assess the reasoning capabilities of LLMs.

Fortunately, many OEIS sequences include corresponding Mathematica code that can be regarded as the ground-truth solution for each problem. We extract this Mathematica code for each sequence, formalizing it to compute the first  $N$  terms,  $A_1, \dots, A_N$ , of the sequence. We determine the maximum value of  $N_{max}$  for which the code can

Category	# of Problems
Number Theory	159
Graph Theory	79
Group Theory	65
Discrete Mathematics	158
Combinatorial Mathematics	158
Geometry and Topology	70
Poly. and Series Expan.	151
Special Numbers	157
Formal Languages	56
<b>Total</b>	1053

Table 2: Categories and distribution of problems.

compute the sequence within 10 seconds, where we set  $10^6$  as the upper bound. Finally, we add the last 10 terms  $A_{N_{max}-9}, \dots, A_{N_{max}}$  into our benchmark as the hard test cases to evaluate the complexity of a solution. Our experiments demonstrate that these cases precisely differentiate more efficient and intelligent solutions.

### 3.3 Evaluation Metrics

We adopt the metric  $\text{pass}@k$  to evaluate the performance of LLMs. The metric  $\text{pass}@k$  is a classic metric in code generation, where a problem is solved if any of the  $k$  generated samples passes the unit tests. We use the stable method of calculation proposed by (Chen et al., 2021):

$$\text{pass}@k := \mathbb{E}_{\text{Problems}} \left[ 1 - \binom{n-c}{k} / \binom{n}{k} \right] \quad (1)$$

### 3.4 Dataset Statistics

The main statistics of UTMath are shown in Tab. 1. To gain a deeper understanding of the composition of the UTMath Benchmark, we identified nine mathematical fields and used GPT-4o to categorize each problem to these fields as shown in Tab. 2. Our analysis reveals that only 10 out of 1,053 problems have no references. The reference years span

Model	Pass@1 (%) ↑		Pass@5 (%) ↑		Avg. Run Time (s) ↓		
	PoT	RCoT	PoT	RCoT	PoT	RCoT	Efficiency
<i>closed-source models</i>							
GPT-4o	<b>25.53</b>	<b>26.93 (+1.40)</b>	<b>32.67</b>	<b>35.90 (+3.23)</b>	6.98	6.23	+12.04%
Gemini-1.5-Pro	19.70	19.43 (−0.27)	31.24	33.14 (+1.90)	6.30	6.22	+1.28%
Claude-3.5-Sonnet	18.58	19.11 (+0.53)	27.83	31.34 (+3.51)	6.44	5.32	+21.05%
GPT-3.5-Turbo	11.68	6.82 (−4.86)	17.09	13.30 (−3.79)	5.42	5.06	+7.11%
<i>open-source models</i>							
Qwen2.5-72B	<b>23.48</b>	<b>22.17 (−1.31)</b>	<b>31.05</b>	<b>33.33 (+2.28)</b>	5.88	4.31	+36.42%
DeepSeek-V2.5-236B	20.95	21.63 (+0.68)	30.10	31.72 (+1.62)	6.64	5.44	+22.06%
Qwen2.5-Math-72B	19.72	20.53 (+0.81)	26.69	28.11 (+1.42)	5.04	3.81	+24.40%
LLaMA-3.1-405B	15.76	16.09 (+0.33)	25.26	27.35 (+2.09)	5.73	5.12	+11.91%

Table 3: Pass Rate and Average Run Time of LLMs on UTMATH. We listed the performance of eight large models by the PoT or the RCoT methods across a range of metrics. The average run time is calculated based on the problems solved by both the PoT and the RCoT methods. The efficiency is calculated as: (Avg.RunTime(PoT) - Avg.RunTime(RCoT)) / Avg.RunTime(RCoT). Two qualitative cases are shown in Appendix D.

from 1950 to 2024, with the maximum number of references exceeding 6,000. These findings underscore the cutting-edge nature of our benchmark. More details can be found in Appendix B.

## 4 Reasoning-to-Coding of Thoughts

Compared to methods that simply check whether the outputs generated by LLMs are identical, the code-based evaluation approach enables more accurate assessment by using multiple test cases. It provides additional evaluation metrics, such as runtime, and allows for the observation of the reliability of LLMs.

Initially, we adopted the Program of Thought (PoT) method, where the LLM had to perform reasoning and implement it in one step. However, we noticed that the LLMs often resorted to simpler algorithms, which led to high time complexity or even failure with more complex problems due to limited reasoning depth. To improve this, we explored the **Reasoning-to-Coding of Thoughts (RCoT)** framework, which separates reasoning and implementation into different steps.

In the first round, LLM focuses only on reasoning about the problem. Compared to PoT, RCoT dedicate an entire round to reasoning allowing the LLM to generate a step-by-step, detailed logical reasoning chain, including mathematical theorems, formulas, and properties used. This deeper reasoning approach facilitates the creation of more efficient algorithms with lower time complexity.

In the second round, RCoT requires the LLM to implement the reasoning process generated in the first round. By converting the computation process into code, we can introduce new metrics beyond the

final pass rate. The solution generated by the LLM can be evaluated based on its runtime during testing, which indirectly reflects the solution’s time complexity. Additionally, this approach avoids errors caused by the limited computational capabilities of large models, allowing for more accurate and genuine insights into the LLM’s reasoning abilities. Clearly, the stronger the reasoning capabilities of the large model, the higher the overall pass rate and the lower the time complexity, which is reflected in the shorter runtime of the generated solution.

We present qualitative cases in Appendix D, where GPT-4o solves problems from UTMATH using PoT and RCoT, respectively. The case studies show that, with RCoT prompting, the model engages in deeper reasoning, significantly reducing solution complexity.

## 5 Experiment

### 5.1 Experimental Setup

Here, we consider the closed-source models, i.e., GPT-3.5-Turbo/GPT-4o from OpenAI (OpenAI, 2024), Claude-3.5-Sonnet (Claude, 2024), Gemini-1.5-Pro (Reid et al., 2024), as well as the open-source models, i.e., LLaMA-3.1 (Dubey et al., 2024), Qwen2.5 (Qwen, 2024a), Qwen2.5-Math (Qwen, 2024b), DeepSeek-V2.5 (Bi et al., 2024). The metric pass@1 is calculated as the average result over 5 run times. We run all evaluations in a laptop with CPU Intel(R) Core(TM) i7-10750H CPU @ 2.60GHz.

### 5.2 Evaluation on UTMATH

Here we evaluate both open-source and closed-source models using RCoT and PoT in Tab. 3. The

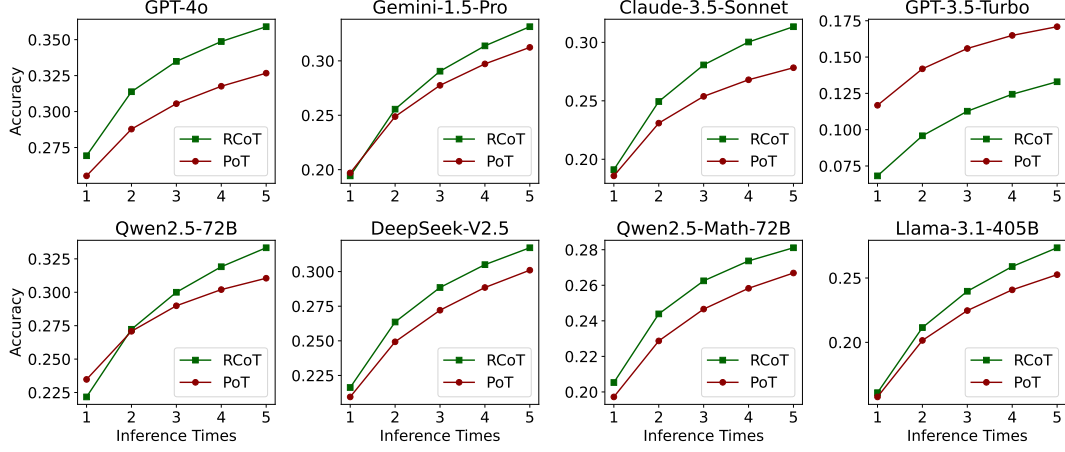


Figure 4: Performance comparison of models across PoT and RCoT tasks at different pass@k levels.

experimental results shows that all tested models performed poorly on our benchmark. The best model, GPT-4o, only solves 26.93% problem in our benchmark. Since our problems are sourced from the OEIS, they consist of sequences and solutions proposed by various mathematicians in the context of cutting-edge research. This suggests that our benchmark is challenging enough to help guide future directions for improving LLMs.

Compared to PoT, our method RCoT demonstrates superiority in two aspects. First, prompting with RCoT achieves higher pass@5 performance across 7 LLMs, with the best results observed on GPT-4o. Second, the solutions generated by RCoT for all LLMs demonstrate more efficient performance, particularly Qwen2.5-72B, where the RCoT approach achieves an efficiency improvement of over 36.42% compared to PoT, as shown in Tab. 3 and Fig. 2. It indicates that, RCoT prompting enables the model to engage in deeper reasoning, significantly reducing solution complexity and enhancing solution performance. However, some models experienced a decrease in pass@1 with RCoT. Specifically, the accuracies of Gemini-1.5-Pro, GPT-3.5-Turbo, and Qwen2.5-72B slightly dropped. Notably, while Gemini-1.5-Pro and Qwen2.5-72B experienced a drop in pass@1, their pass@5 performance improved. It indicates that RCoT brings more room in multiple inference times. The observed decrease in performance may stem from the fact that formulating more efficient solutions often requires higher-level reasoning, which can increase the difficulty of the task and make these models more susceptible to errors when attempting more sophisticated solutions.

	Model	Easy	Easy & Hard
closed	GPT-4o	<b>34.95</b>	<b>26.93</b>
	Gemini-1.5-Pro	23.84	19.43
	Claude-3.5-Sonnet	24.86	19.11
	GPT-3.5-Turbo	8.72	6.82
open	Qwen2.5-72B	<b>28.96</b>	<b>22.17</b>
	DeepSeek-V2.5	27.52	21.63
	Qwen2.5-Math-72B	24.60	20.53
	LLaMA-3.1-405B	22.09	16.09

Table 4: Performance (%) of different models on easy and hard test cases. Easy cases: The initial terms in OEIS. Hard cases: mined hard test cases (§ 3.2).

### 5.3 The Effectiveness of Hard Test Cases

As we mentioned in § 3.2, each sequence in the OEIS lists only the initial terms, which we refer to as “easy test cases”. To investigate the model’s ability to handle challenging cases, we evaluated whether it could predict values that appear later (i.e.,  $10^6$ ) in a sequence. These later values are typically underrepresented in pre-training data and often require more computation time and a more precise implementation to retrieve accurately. The experimental results, which depicted in Tab. 4, reveal that the model’s performance drops significantly when handling these hard cases. This indicates that introducing these cases can prevent simple solutions from passing all the test cases, thereby filtering for more advanced solutions.

### 5.4 Scaling of the Inference Times

We compared the performance difference between running the LLMs five times and reported the metric of pass@k. As shown in Fig. 4, all models improved their performance with an increasing

Model	NT	Graph T.	Group T.	DM	CM	GT	PSE	SN	FL	pass@1
<i>closed-source models</i>										
GPT-4o	<b>43.90</b>	<b>2.78</b>	<b>11.69</b>	<b>38.23</b>	<b>24.94</b>	3.43	<b>16.42</b>	<b>33.89</b>	<b>42.50</b>	<b>26.93</b>
Gemini-1.5-Pro	31.70	1.27	8.92	27.47	15.19	<b>5.71</b>	15.23	27.39	17.86	19.43
Claude-3.5-Sonnet	33.58	1.52	8.00	29.49	12.91	5.43	11.52	26.62	20.36	19.11
GPT-3.5-Turbo	13.08	0.00	1.85	11.39	3.29	0.29	2.78	10.96	8.93	6.82
<i>open source models</i>										
Qwen2.5-72B	36.86	<b>2.53</b>	<b>12.00</b>	30.63	15.95	<b>6.00</b>	<b>18.15</b>	29.43	<b>24.29</b>	<b>22.17</b>
DeepSeek-V2.5	<b>38.24</b>	1.27	8.92	<b>33.16</b>	<b>17.34</b>	2.29	12.45	<b>31.08</b>	20.00	21.63
Qwen2.5-Math-72B	35.35	1.27	8.62	28.73	14.81	4.00	17.48	28.15	20.00	20.53
LLaMA-3.1-405B	29.56	0.76	4.92	25.44	9.62	2.00	9.54	22.55	21.43	16.09

Table 5: Performance (%) on different problem categories. Categories are represented by abbreviations. NT: Number Theory; T.: Theory; DM: Discrete Mathematics; CM: Combinatorial Mathematics; GT: Geometry and Topology; PSE: Polynomial and Series Expansions; SN: Special Numbers; FL: Formal Languages.

number of inference times. For Qwen2.5-72B and Gemini-1.5-Pro, RCoT was slightly weaker than PoT in pass@1 but quickly approached and surpassed PoT in subsequent run times. We observed that with an increasing number of inference time, RCoT consistently demonstrated a growing advantage in performance across almost all models, except for GPT-3.5. However, it is worth noting that GPT-3.5 exhibited the lowest pass rate. This suggests that RCoT may perform better in models with stronger reasoning capabilities.

## 5.5 Importance of the Reasoning Step

GPT-4o has the best performance, while we are unclear whether this was due to its superior reasoning or coding ability. To investigate further, we tested using GPT-4o for the reasoning step while other models perform the coding step based on the reasoning result from GPT-4o. As depicted in Fig. 5, the results showed that the performance of models increased significantly when implementing coding based on GPT-4o’s reasoning output, suggesting that the reasoning quality is important and GPT-4o does produce higher-quality reasoning results.

## 5.6 Performance on Different Categories

Our benchmark comprehensively evaluates the LLMs’ ability across various categories of math problems. GPT-4o continued to demonstrate a clear advantage across most tasks, achieving the best performance in six domains. Notably, it outperformed the second-best model by 7.6% in combinatorial mathematics and by 18.21% in formal languages. All models performed very poorly in the categories of Graph Theory, Group Theory, and Geometry and Topology, with accuracy rates below 12%, highlighting the need for further exploration in these areas.




Figure 5: Performance comparison between self-reasoning and using GPT-4o reasoning for coding across different models. The results show that models perform better when relying on GPT-4o’s reasoning output.

## 6 Conclusion

In this work, we investigate how to more accurately and effectively evaluate the mathematical reasoning capabilities of LLMs. We propose a cutting-edge benchmark, UTMATH, which comprises 1,053 problems spanning nine mathematical domains, with an average of 68 test cases per problem. This benchmark presents significant challenges: GPT-4o, the best-performing model, successfully solves only 26.93% of the problems. Additionally, we introduce RCoT (Reasoning-to-Coding of Thought). Our study finds that, compared to PoT (Program-of-Thought), RCoT significantly improves the algorithmic efficiency and pass rates of most models by facilitating deeper reasoning steps prior to code generation. Overall, this research contributes to a deeper understanding of the current capabilities of LLMs in mathematical reasoning and lays the groundwork for the development of more advanced models in the future.



## Limitation

The primary limitation of UTMATH lies in the evaluation metrics: the performance of the evaluation machine affects the runtime of the generated code, making the absolute numerical results incomparable across different machines. We utilized an i7-10750H processor to execute the reference solutions and conduct evaluations, and we recommend using the same machine for testing and replication. There are two main limitations of RCoT. First, we only installed a set of common packages, such as sympy, in the standard testing environment. This avoids allowing LLMs to call highly integrated packages while also preventing the generation of potentially harmful code that could damage the evaluation system. Second, while our experiments demonstrate the critical role of reasoning quality in determining success rates, we have not further explored methods for enhancing reasoning quality, which remains an area for future investigation.

## Ethics Statements

The UTMATH Benchmark is designed to advance the evaluation of mathematical reasoning in LLMs. We recognize the potential ethical concerns associated with this work, particularly the risk of data misuse. To mitigate this, we strictly adhere to usage guidelines and licensing terms for the UTMATH-Train dataset, which is intended solely for academic and research purposes. While the UTMATH Benchmark evaluates model performance in terms of accuracy and generality, automated evaluations may introduce biases due to the nature of the datasets and evaluation algorithms. Additionally, while UTMATH covers a wide range of mathematical domains, it may not fully represent diverse cultural or educational perspectives. We encourage further development of benchmarks that incorporate a broader array of reasoning styles to ensure more inclusive evaluations. By releasing UTMATH, we aim to foster responsible AI development, promoting better, more generalizable mathematical reasoning systems.

## References

- Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. 2024. Large language models for mathematical reasoning: Progresses and challenges. *arXiv preprint arXiv:2402.00157*.
- Husni Almoubayyed, Rae Bastoni, Susan R Berman,

Sarah Galasso, Megan Jensen, Leila Lester, April Murphy, Mark Swartz, Kyle Weldon, Stephen E Fancsali, et al. 2023. Rewriting math word problems to improve learning outcomes for emerging readers: a randomized field trial in carnegie learning’s mathia. In *International Conference on Artificial Intelligence in Education*, pages 200–205. Springer.

Ujjwala Ananthaswaran, Himanshu Gupta, Kevin Scaria, Shreyas Verma, Chitta Baral, and Swaroop Mishra. 2024. Investigating the robustness of llms on math word problems. *arXiv preprint arXiv:2406.15444*.

Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, et al. 2024. Graph of thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17682–17690.

Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiusi Du, Zhe Fu, et al. 2024. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*.

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3):1–45.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2022. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.

Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. 2023. Theoremqa: A theorem-driven question answering dataset. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7889–7901.

Cheng-Han Chiang and Hung-yi Lee. 2023. Can large language models be an alternative to human evaluations? *arXiv preprint arXiv:2305.01937*.

Claude. 2024. Claude-3.5-sonnet. <https://www.anthropic.com/claude/>.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

656	Iddo Drori, Sarah J Zhang, Reece Shuttleworth, Sarah	Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap,	713
657	Zhang, Keith Tyser, Zad Chin, Pedro Lantigua,	Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and	714
658	Saisamrit Surbehera, Gregory Hunter, Derek Austin,	Ion Stoica. 2024b. From crowdsourced data to high-	715
659	et al. 2023. From human days to machine seconds:	quality benchmarks: Arena-hard and benchbuilder	716
660	Automatically answering and generating machine	pipeline. <i>arXiv preprint arXiv:2406.11939</i> .	717
661	learning final exams. In <i>Proceedings of the 29th</i>		
662	<i>ACM SIGKDD Conference on Knowledge Discovery</i>	Zekun Li, Baolin Peng, Pengcheng He, and Xifeng	718
663	<i>and Data Mining</i> , pages 3947–3955.	Yan. 2023. Do you really follow me? adversarial	719
664	Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey,	instructions for evaluating the robustness of large	720
665	Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman,	language models. <i>arXiv preprint arXiv:2308.10819</i> .	721
666	Akhil Mathur, Alan Schelten, Amy Yang, Angela		
667	Fan, et al. 2024. The llama 3 herd of models. <i>arXiv</i>	Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang,	722
668	<i>preprint arXiv:2407.21783</i> .	Ruochen Xu, and Chenguang Zhu. 2023. G-eval:	723
669	Yann Dubois, Balázs Galambosi, Percy Liang, and Tat-	Nlg evaluation using gpt-4 with better human align-	724
670	sunori B Hashimoto. 2024. Length-controlled al-	ment. <i>arXiv preprint arXiv:2303.16634</i> .	725
671	pacaeval: A simple way to debias automatic evalua-		
672	tors. <i>arXiv preprint arXiv:2404.04475</i> .	OEIS Foundation Inc. 2024. The on-line encyclope-	726
673	Yao Fu, Litu Ou, Mingyu Chen, Yuhao Wan, Hao	dia of integer sequences. Published electronically at	727
674	Peng, and Tushar Khot. 2023. Chain-of-thought	<a href="https://oeis.org">https://oeis.org</a> .	728
675	hub: A continuous effort to measure large language		
676	models’ reasoning performance. <i>arXiv preprint</i>	OpenAI. 2024. Hello gpt-4o, 2024. <a href="https://openai.com/index/hello-gpt-4o/">https://openai.com/index/hello-gpt-4o/</a> .	729
677	<i>arXiv:2305.17306</i> .		730
678	Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon,	Arkil Patel, Satwik Bhattamishra, and Navin Goyal.	731
679	Pengfei Liu, Yiming Yang, Jamie Callan, and Gra-	2021. Are nlp models really able to solve	732
680	ham Neubig. 2023. Pal: Program-aided language	simple math word problems? <i>arXiv preprint</i>	733
681	models. In <i>International Conference on Machine</i>	<i>arXiv:2103.07191</i> .	734
682	<i>Learning</i> , pages 10764–10799. PMLR.		
683	Himanshu Gupta, Neeraj Varshney, Swaroop Mishra,	Ethan Perez, Sam Ringer, Kamilė Lukošiuotė, Karina	735
684	Kuntal Kumar Pal, Saurabh Arjun Sawant, Kevin	Nguyen, Edwin Chen, Scott Heiner, Craig Pettit,	736
685	Scaria, Siddharth Goyal, and Chitta Baral. 2022. "	Catherine Olsson, Sandipan Kundu, Saurav Kada-	737
686	john is 50 years old, can his son be 65?" evaluat-	vath, et al. 2022. Discovering language model behav-	738
687	ing nlp models’ understanding of feasibility. <i>arXiv</i>	iors with model-written evaluations. <i>arXiv preprint</i>	739
688	<i>preprint arXiv:2210.07471</i> .	<i>arXiv:2212.09251</i> .	740
689	Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu,	Team Qwen. 2024a. Qwen2.5-72b-instruct.	741
690	Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han,	<a href="https://qwenlm.github.io/blog/qwen2.5/">https://qwenlm.github.io/blog/qwen2.5/</a> .	742
691	Yujie Huang, Yuxiang Zhang, et al. 2024. Olympiad-	Team Qwen. 2024b. Qwen2.5-math-72b-instruct.	743
692	bench: A challenging benchmark for promoting agi	<a href="https://qwenlm.github.io/blog/qwen2.5-math/">https://qwenlm.github.io/blog/qwen2.5-math/</a> .	744
693	with olympiad-level bilingual multimodal scientific		
694	problems. <i>arXiv preprint arXiv:2402.14008</i> .	Machel Reid, Nikolay Savinov, Denis Teplyashin,	745
695	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul	Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste	746
696	Arora, Steven Basart, Eric Tang, Dawn Song, and Ja-	Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Fi-	747
697	cob Steinhardt. 2021. Measuring mathematical prob-	rat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Un-	748
698	lem solving with the math dataset. <i>arXiv preprint</i>	locking multimodal understanding across millions of	749
699	<i>arXiv:2103.03874</i> .	tokens of context. <i>arXiv preprint arXiv:2403.05530</i> .	750
700	Tom Kocmi and Christian Federmann. 2023. Large	Steven Ritter, John R Anderson, Kenneth R Koedinger,	751
701	language models are state-of-the-art evaluators of	and Albert Corbett. 2007. Cognitive tutor: Applied	752
702	translation quality. <i>arXiv preprint arXiv:2302.14520</i> .	research in mathematics education. <i>Psychonomic</i>	753
703	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-	<i>bulletin &amp; review</i> , 14:249–255.	754
704	taka Matsuo, and Yusuke Iwasawa. 2022. Large lan-	Tomohiro Sawada, Daniel Paleka, Alexander Havrilla,	755
705	guage models are zero-shot reasoners. <i>Advances in</i>	Pranav Tadepalli, Paula Vidas, Alexander Kranias,	756
706	<i>neural information processing systems</i> , 35:22199–	John J Nay, Kshitij Gupta, and Aran Komatsuzaki.	757
707	22213.	2023. Arb: Advanced reasoning benchmark for large	758
708	Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng	language models. <i>arXiv preprint arXiv:2307.13692</i> .	759
709	Kong, and Wei Bi. 2024a. Gsm-plus: A compre-	Paulo Shakarian, Abhinav Koyyalamudi, Noel Ngu, and	760
710	hensive benchmark for evaluating the robustness of	Lakshmivihari Mareedu. 2023. An independent eval-	761
711	llms as mathematical problem solvers. <i>arXiv preprint</i>	uation of chatgpt on mathematical word problems	762
712	<i>arXiv:2402.19255</i> .	(mwp). <i>arXiv preprint arXiv:2302.13814</i> .	763

764	Freda Shi, Xinyun Chen, Kanishka Misra, Nathan	Zihao Zhou, Qiufeng Wang, Mingyu Jin, Jie Yao, Jianan	821
765	Scales, David Dohan, Ed H Chi, Nathanael Schärli,	Ye, Wei Liu, Wei Wang, Xiaowei Huang, and Kaizhu	822
766	and Denny Zhou. 2023. Large language models can	Huang. 2024b. Mathattack: Attacking large language	823
767	be easily distracted by irrelevant context. In <i>Inter-</i>	models towards math solving ability. In <i>Proceedings</i>	824
768	<i>national Conference on Machine Learning</i> , pages	<i>of the AAAI Conference on Artificial Intelligence</i> ,	825
769	31210–31227. PMLR.	volume 38, pages 19750–19758.	826
770	Zhengyang Tang, Xingxing Zhang, Benyou Wan, and		
771	Furu Wei. 2024. Mathscale: Scaling instruction		
772	tuning for mathematical reasoning. <i>arXiv preprint</i>		
773	<i>arXiv:2403.02884</i> .		
774	Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan,		
775	and Subbarao Kambhampati. 2022. Large language		
776	models still can’t plan (a benchmark for llms on plan-		
777	ning and reasoning about change). In <i>NeurIPS 2022</i>		
778	<i>Foundation Models for Decision Making Workshop</i> .		
779	Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi		
780	Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. 2023. Plan-		
781	and-solve prompting: Improving zero-shot chain-of-		
782	thought reasoning by large language models. <i>arXiv</i>		
783	<i>preprint arXiv:2305.04091</i> .		
784	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten		
785	Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,		
786	et al. 2022. Chain-of-thought prompting elicits rea-		
787	soning in large language models. <i>Advances in neural</i>		
788	<i>information processing systems</i> , 35:24824–24837.		
789	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran,		
790	Tom Griffiths, Yuan Cao, and Karthik Narasimhan.		
791	2024. Tree of thoughts: Deliberate problem solving		
792	with large language models. <i>Advances in Neural</i>		
793	<i>Information Processing Systems</i> , 36.		
794	Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu,		
795	Zhengying Liu, Yu Zhang, James T Kwok, Zhen-		
796	guo Li, Adrian Weller, and Weiyang Liu. 2023.		
797	Metamath: Bootstrap your own mathematical ques-		
798	tions for large language models. <i>arXiv preprint</i>		
799	<i>arXiv:2309.12284</i> .		
800	Xiaotian Zhang, Chunyang Li, Yi Zong, Zhengyu Ying,		
801	Liang He, and Xipeng Qiu. 2023. Evaluating the		
802	performance of large language models on gaokao		
803	benchmark. <i>arXiv preprint arXiv:2305.12474</i> .		
804	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan		
805	Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin,		
806	Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023.		
807	Judging llm-as-a-judge with mt-bench and chatbot		
808	arena. <i>Advances in Neural Information Processing</i>		
809	<i>Systems</i> , 36:46595–46623.		
810	Wanjuan Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang,		
811	Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen,		
812	and Nan Duan. 2023. Agieval: A human-centric		
813	benchmark for evaluating foundation models. <i>arXiv</i>		
814	<i>preprint arXiv:2304.06364</i> .		
815	Zihao Zhou, Shudong Liu, Maizhen Ning, Wei Liu,		
816	Jindong Wang, Derek F Wong, Xiaowei Huang, Qi-		
817	ufeng Wang, and Kaizhu Huang. 2024a. Is your		
818	model really a good math reasoner? evaluating math-		
819	ematical reasoning with checklist. <i>arXiv preprint</i>		
820	<i>arXiv:2407.08733</i> .		

0 1 3 6 2 7  
: 13  
: 20  
23 12  
10 22 11 21


THE ON-LINE ENCYCLOPEDIA  
OF INTEGER SEQUENCES®

The OEIS is supported by [the many generous donors to the OEIS Foundation](#).

founded in 1964 by N. J. A. Sloane

Year-end appeal: Please make a donation to the OEIS Foundation to support ongoing development and maintenance of the OEIS. We are now in our 61st year, we have over 378,000 sequences, and we've reached 11,000 citations (which often say "discovered thanks to the OEIS").

Donate



[Other ways to Give](#)

(Greetings from [The On-Line Encyclopedia of Integer Sequences!](#))

Search [Hints](#)

A007369Numbers n such that sigma(x) = n has no solution.  
(Formerly M1355)51

2, 5, 9, 10, 11, 16, 17, 19, 21, 22, 23, 25, 26, 27, 29, 33, 34, 35, 37, 41, 43, 45, 46, 47, 49, 50, 51, 52, 53, 55, 58, 59, 61, 64, 65, 66, 67, 69, 70, 71, 73, 75, 76, 77, 79, 81, 82, 83, 85, 86, 87, 88, 89, 92, 94, 95, 97, 99, 100, 101, 103, 105, 106, 107, 109, 111, 113  
[\(list; graph; refs; listen; history; text; internal format\)](#)

OFFSET1,1

COMMENTSWith an initial 1, may be constructed inductively in stages from the list L = {1,2,3,...} by the following sieve procedure. Stage 1. Add 1 as the first term of the sequence a(n) and strike off 1 from L. Stage n+1. Add the first (i.e. leftmost) term k of L as a new term of the sequence a(n) and strike off k, sigma(k), sigma(sigma(k)),... from L. - [Joseph L. Pe](#), May 08 2002  
This sieve is a special case of a more general sieve. Let D be a subset of N and let f be an injection on D satisfying f(n) > n. Define the sieve process as follows: 1. Start with the empty sequence S and let E = D. 2. Append the smallest element s of E to S. 3. Remove s, f(s), f(f(s)), f(f(f(s))), ... from E. 4. Go to step 2. After this sieving process, S = D - f(D). To get the current sequence, take f = sigma and D = {n | n >= 2}. - [Max Alekseyev](#), Aug 08 2005  
By analogy with the untouchable numbers ([A005114](#)), these numbers could be named "sigma-untouchable". - [Daniel Lignon](#), Mar 28 2014  
The asymptotic density of this sequence is 1 (Niven, 1951, Rao and Murty, 1979). - [Amiram Eldar](#), Jul 23 2020

REFERENCESM. Abramowitz and I. A. Stegun, eds., Handbook of Mathematical Functions, National Bureau of Standards Applied Math. Series 55, 1964 (and various reprintings), p. 840.  
N. J. A. Sloane and Simon Plouffe, The Encyclopedia of Integer Sequences, Academic Press, 1995 (includes this sequence).

LINKSM. F. Hasler, [Table of n, a\(n\) for n = 1..10000](#) (first 1000 terms from T. D. Noe)  
M. Abramowitz and I. A. Stegun, eds., [Handbook of Mathematical Functions](#), National Bureau of Standards, Applied Math. Series 55, Tenth Printing, 1972 [alternative scanned copy].  
Ivan Niven, [The asymptotic density of sequences](#), Bull. Amer. Math. Soc., Vol. 57 (1951), pp. 420-434.  
R. Sita Rama Chandra Rao and G. Sri Rama Chandra Murty, [On a theorem of Niven](#), Canadian Mathematical Bulletin, Vol 22, No. 1 (1979), pp. 113-115.  
R. G. Wilson, V, [Letter to N. J. A. Sloane, Jul. 1992](#)

FORMULA[A175192](#)(a(n)) = 0, [A054973](#)(a(n)) = 0. - [Jaroslav Krizek](#), Mar 01 2010  
a(n) < 2n + sqrt(8n). - [Charles R Greathouse IV](#), Oct 23 2015

EXAMPLEa(4) = 10 because there is no x < 10 whose sigma(x) = 10.

MATHEMATICAa = {}; Do[s = DivisorSigma[1, n]; a = Append[a, s], {n, 1, 115}]; Complement[Table[n, {n, 1, 115}], Union[a]]

PROG(PARI) list(lim)=my(v=List(), u=vectorsmall(lim\1, t); for(n=1, lim, t=sigma(n); if(t<=lim, u[t]=1)); for(n=2, lim, if(u[n]==0, listput(v, n))); Vec(v) \\ [Charles R Greathouse IV](#), Mar 09 2017  
(PARI) [A007369](#). list(LIM, m=0, L=List(), s)={for(n=2, LIM, (s=sigma(n-1))>LIM || bittest(m, s) || m+=1<<s; bittest(m, n)||listput(L, n)); L} \\ A bit slower, but bitmask requires less memory, avoiding stack overflow produced by the earlier code for lim = 1e6 with standard gp setup. - [M. F. Hasler](#), Mar 12 2018

CROSSREFSComplement of [A002191](#).  
See [A083532](#) for the gaps, i.e., first differences.  
See [A048995](#) for the missed sums of nontrivial divisors.  
Sequence in context: [A133508](#) [A125969](#) [A070240](#) \* [A307213](#) [A362398](#) [A295567](#)  
Adjacent sequences: [A007366](#) [A007367](#) [A007368](#) \* [A007370](#) [A007371](#) [A007372](#)

KEYWORDnonnn

AUTHORN. J. A. Sloane, Mira Bernstein, Robert G. Wilson v

EXTENSIONSMore terms from [David W. Wilson](#)

STATUSapproved

Figure 6: Sequence A007369 in OEIS. Its description: "Numbers n such that sigma(x) = n has no solution." (Clearly, without specific background knowledge, we cannot fully understand what the function sigma() represents, which is one of the reasons we perform standardization. §B.2) Next, OEIS shows the first 67 terms of this sequence, which we classify as easy cases. Below that, additional metadata is provided, including comments, references, links, formulas, examples, programs, author, status, and more. It is evident that this sequence has garnered significant attention from researchers, reflecting the Cutting-Edge difficulty of our benchmark. We used the Mathematica program included in the metadata to generate Hard cases, with detailed procedures provided in § 3.2. As a scientific database, each sequence submitted to OEIS undergoes a review process, and the status "approve" indicates that the sequence has been validated and approved by OEIS administrators.



## B Dataset Construction Details

This section primarily presents some details on the construction of UTMATH. B.1 discusses the issues encountered when observing data crawled from OEIS, along with the corresponding cleaning rules. UTMATH applies all 14 rules. Additionally, we crawled all sequences from OEIS and, for convenience, applied only the first 12 rules to create UTMATH\_Train, which contains over 70k sequences. B.2 outlines the process followed for standardizing the descriptions of problems in UTMATH, while B.3 explains the referencing of sequences within UTMATH, highlighting both the Cutting-Edge difficulty level of UTMATH and its scalability.

### B.1 Rules for Data Cleaning

1. Issues: The sequence is too difficult, requiring extensive background knowledge, or only a limited number of terms are found.  
Method: Remove sequences with keywords containing 'hard', 'fin' (finite).
2. Issues: The sequence is hard to generate with a program.  
Method: Check if it contains program, formula, or Mathematica fields in the sequence's json data.
3. Issues: The sequence is too simple with an explicit recurrence or closed formula.  
Method: Search if the description includes 'a(n) = '.
4. Issues: Solving the sequence requires information from other OEIS sequences.  
Method: Search if the sequence's description contains the AID of other sequences ('A' + six-digit number with leading zeros).
5. Issues: The sequence is decimal expansion of a certain number.  
Method: Search if the description includes 'decimal'.
6. Issues: The sequence consists of repetitions or a constant value.  
Method: Search if the description includes both 'repeat' and 'period' or 'constant sequence'.
7. Issues: The description is too vague.  
Method: Search if the description includes 'related to'.
8. Issues: Another version of a concept.  
Method: Search if the title includes 'another version', 'second kind', etc.
9. Issues: The sequence is formed by taking mod of a constant.  
Method: Search if the description includes 'module'.
10. Issues: The values in the sequence are too large, which might cause LLM tokenization errors.  
Method: Check if any term's length exceeds 18 digits (i.e., greater than  $1e18$ ), remove it.

11. Issues: Coefficient triangles or 'read by row' topics.  
Method: Search if the sequence's description includes 'read by row', 'triangle of coefficient'.
12. Issues: The description is too short, either purely implementation or lacks necessary information.  
Method: Check if the title length is below 5.
13. Issues: More like a reasoning puzzle.  
Method: Use GPT-4o to judge, with the prompt outlined in the Appendix C.
14. Issues: Non-mathematical topics.  
Method: Use GPT-4o to judge, with the prompt outlined in the Appendix C.

### B.2 Standardization of Problems' Description




Figure 7: Comparison between original and standardized problem description. The standardized version includes hints and explains the specific meaning of  $\sigma(x)$ .

### B.3 Dataset Statistics

To demonstrate that our benchmark is of cutting-edge level, we have analyzed the distribution of the publication years and the number of references included in the problems of the benchmark as shown in Fig. 8. Additionally, OEIS is a dynamic database. Over the past five years, more than 35,000 sequences in UTMATH\_Train have been further researched, and over 2,000 new sequences have been added. This ongoing development makes it possible to continuously update UTMATH\_Train and UTMATH, helping to address the challenges posed by data leakage.




Figure 8: Distribution of references in UTMATH.

C Prompts

<p><b>Prompt 1: the prompt for Program-of-Thoughts</b></p> <p>Please reason through the following sequence problem and implement your reasoning using code. You need to follow these requirements:</p> <ol style="list-style-type: none"><li>1. The code must use the Python language.</li><li>2. Use the function signature <code>def solution(x: int)</code>, and make sure the code part is in markdown format.</li><li>3. To ensure the code is runnable, please import any necessary libraries.</li><li>4. Provide the reasoning process first.</li><li>5. Use the solution with the lowest time complexity.</li></ol> <p>Question Statements: <i>{The statement of the question}</i></p> <p>Examples: <i>Solution(1) == a<sub>1</sub></i> <i>Solution(2) == a<sub>2</sub></i> <i>Solution(3) == a<sub>3</sub></i></p>	<p><b>Prompt 2: the prompt for Reasoning-to-Coding of Thoughts</b></p> <p>----- Turn 1: &lt; Reasoning Instruction &gt; -----</p> <p>Please analyze the following sequence problem and provide a detailed reasoning process for the sequence. You need to follow these requirements:</p> <ol style="list-style-type: none"><li>1. Use the solution with the lowest time complexity.</li><li>2. Not to implement the solution.</li></ol> <p>Question Statements: <i>{The statement of the question}</i></p> <p>Examples: <i>a(1) == a<sub>1</sub></i> <i>a(2) == a<sub>2</sub></i> <i>a(3) == a<sub>3</sub></i></p> <p>-----Turn 2: &lt; Coding Instruction &gt; -----</p> <p>Please implement the above solution using Python code, adhering to the following requirements:</p> <ol style="list-style-type: none"><li>1. The code must be written in Python.</li><li>2. Use the function signature <code>def solution(x: int)</code>, and ensure the code portion is in markdown format.</li><li>3. To ensure the code is runnable, please import any necessary libraries.</li><li>4. You do not need to provide any explanations or examples, just the implementation code.</li><li>5. test contains multiple test cases, each of which will call the solution function.</li></ol> <p>Examples: <i>Solution(1) == a<sub>1</sub></i> <i>Solution(2) == a<sub>2</sub></i> <i>Solution(3) == a<sub>3</sub></i></p>
<p><b>Prompt 3: the prompt used to determine reasoning or implementation questions</b></p> <p>This is a description of a sequence. Please judge whether solving this sequence requires more reasoning or implementation. You need to follow these rules:</p> <ol style="list-style-type: none"><li>1. If the problem statement already has a clear recurrence relation or explicit formula, the question should be considered an implementation question.</li><li>2. If the problem does not include a direct calculation formula and requires reasoning to derive it, the question should be considered a reasoning question.</li><li>3. Implementation questions usually just require translating the problem requirements directly into code without designing complex algorithms or using advanced data structures.</li><li>4. If the question requires more reasoning, answer "reasoning question"; otherwise, answer "implementation question."</li><li>5. Your answer should be in italics.</li></ol> <p>Question Statements: <i>{The statement of the question}</i></p>	<p><b>Prompt 5: the prompt used to standardize question statements</b></p> <p>I want to create a math problem based on a sequence from the OEIS. The output should be the first n terms of the sequence, but the original problem statement for the sequence may be vague or difficult to understand or might require additional background knowledge. I will provide you with the original problem statement and other information. Please use this information and your knowledge to complete the original problem statement without revealing the sequence's reasoning method or content, and you need to follow these rules:</p> <ol style="list-style-type: none"><li>1. If the original problem statement is sufficient to deduce the entire sequence, no additional information is needed.</li><li>2. You should retain as much of the original problem statement as possible.</li><li>3. The completed problem statement should include necessary background knowledge.</li><li>4. The completed problem statement should maintain the same meaning as the original.</li><li>5. The completed problem statement should not contain direct recurrence relations or explicit formulas.</li><li>6. The completed problem statement should remain in English.</li><li>7. The completed problem statement should begin with "The definition of a(n) is".</li><li>8. Your response should only include the completed problem statement without any explanations or examples.</li></ol> <p>Question Statements: <i>{The statement of the question}</i></p> <p>Other information: <i>{The information about the question}</i></p>
<p><b>Prompt 4: the prompt used to determine whether the question belongs to the field of mathematics :</b></p> <p>This is information about a sequence from OEIS (The On-Line Encyclopedia of Integer Sequences) and contains four types of information: 'name', 'data', 'comment' and 'formula'. Please use this information and your knowledge to judge the domain to which the sequence belongs. Please follow these rules:</p> <ol style="list-style-type: none"><li>1. Your response should only contain the answer, without any other explanations or examples.</li><li>2. Your answer can only be selected from these five options: {'math question', 'physics question', 'chemistry question', 'biology question', 'other question'}</li><li>3. Your answer should be in italics</li></ol> <p>Question Statements: <i>{The statement of the question}</i></p> <p>Data: <i>{The items of the sequence}</i></p> <p>Comment: <i>{The comment of the question}</i></p> <p>Formula: <i>{The formula of the question}</i></p>	

Figure 9: The main prompts we used.

<p><b>Task_Id:</b> UTMath_948 <b>OEIS_Id:</b> A019567 <b>Main_Category:</b> Special Numbers</p> <p><b>Problem Statement:</b> The definition of <math>a(n)</math> is the least number <math>m</math> for which either <math>2^m + 1</math> or <math>2^m - 1</math> is divisible by <math>4n + 1</math>, where <math>n</math> represents the number of pairs of cards in a Mongean shuffle.</p> <p><b>Case 1:</b> <math>a(0) = 1</math> <b>Case 2:</b> <math>a(1) = 2</math> ...</p> <p><b>Hard Case 1:</b> <math>a(1882) = 941</math> <b>Hard Case 2:</b> <math>a(1883) = 810</math> ...</p> <p><b>Hard Case 10:</b> <math>a(1891) = 88</math></p>	<p><b>The Responses of GPT-4o by Reasoning-To-Coding Thoughts of Thought Methods</b></p> <p>----- Turn 1: &lt; Reasoning &gt; -----</p> <p>To analyze the sequence problem, we need to determine the least number <math>(m)</math> such that either <math>(2^m + 1)</math> or <math>(2^m - 1)</math> is divisible by <math>(4n + 1)</math>. Let's break down the problem and find a solution with the lowest time complexity.</p> <p><b>Understanding the Divisibility Condition:</b> We need to find the smallest <math>(m)</math> such that either <math>2^m + 1 \equiv 0 \pmod{4n + 1}</math> or <math>2^m - 1 \equiv 0 \pmod{4n + 1}</math>. This means <math>2^m \equiv -1 \pmod{4n + 1}</math> or <math>2^m \equiv 1 \pmod{4n + 1}</math>.</p> <p><b>Order of 2 Modulo <math>(4n + 1)</math>:</b></p> <ul style="list-style-type: none"> <li>The order of 2 modulo <math>(4n + 1)</math> is the smallest positive integer <math>(m)</math> such that <math>2^m \equiv 1 \pmod{4n + 1}</math>.</li> <li>If <math>2^m \equiv -1 \pmod{4n + 1}</math>, then <math>2^{2m} \equiv 1 \pmod{4n + 1}</math> and <math>(m)</math> is half the order of 2 modulo <math>(4n + 1)</math>.</li> </ul> <p><b>Efficient Calculation:</b></p> <ul style="list-style-type: none"> <li>To find the order of 2 modulo <math>(4n + 1)</math>, we can use the fact that the order must divide <math>\phi(4n + 1)</math> where <math>\phi</math> is Euler's totient function.</li> <li>Calculate <math>\phi(4n + 1)</math> and check divisors of this value to find the smallest <math>(m)</math> such that <math>2^m \equiv 1 \pmod{4n + 1}</math> or <math>2^m \equiv -1 \pmod{4n + 1}</math>.</li> </ul> <p><b>Steps to Solve:</b></p> <ol style="list-style-type: none"> <li>Compute <math>(4n + 1)</math>: For a given <math>(n)</math>, compute <math>(4n + 1)</math>.</li> <li>Calculate <math>\phi(4n + 1)</math>: Use the prime factorization of <math>(4n + 1)</math> to compute <math>\phi(4n + 1)</math>.</li> <li>Find the Order of 2: Check each divisor <math>d</math> of <math>\phi(4n + 1)</math> to see if <math>2^d \equiv 1 \pmod{4n + 1}</math> or <math>2^d \equiv -1 \pmod{4n + 1}</math>. The smallest such <math>d</math> is the desired <math>(m)</math>.</li> </ol> <p><b>Time Complexity</b></p> <ul style="list-style-type: none"> <li>Prime Factorization: The time complexity for factorizing <math>(4n + 1)</math> is <math>O(\sqrt{4n + 1})</math>.</li> <li>Checking Divisors: For each divisor <math>d</math> of <math>\phi(4n + 1)</math>, compute <math>2^d \pmod{4n + 1}</math> using fast exponentiation, which is <math>O(\log d)</math>.</li> <li>Overall Complexity: The overall complexity is dominated by the factorization step, making it <math>O(\sqrt{4n + 1})</math>.</li> </ul> <p>This approach efficiently finds the least <math>(m)</math> with the lowest time complexity by leveraging number theory concepts such as modular arithmetic and Euler's totient function.</p>
<p><b>The Responses of GPT-4o by Program of Thought Methods</b></p> <p>To solve this problem, we need to find the smallest integer <math>(m)</math> such that either <math>2^m + 1</math> or <math>2^m - 1</math> is divisible by <math>(4n + 1)</math>. This means we need to check for the smallest <math>(m)</math> where:</p> <ol style="list-style-type: none"> <li><math>2^m + 1 \equiv 0 \pmod{4n + 1}</math> or</li> <li><math>2^m - 1 \equiv 0 \pmod{4n + 1}</math></li> </ol> <p><b>Given the examples:</b></p> <ul style="list-style-type: none"> <li>For <math>n = 0</math>, <math>4n + 1 = 1</math>. The smallest <math>(m)</math> is 1 because <math>2^1 - 1 = 1</math> which is divisible by 1.</li> <li>For <math>n = 1</math>, <math>4n + 1 = 5</math>. The smallest <math>(m)</math> is 2 because <math>2^2 - 1 = 3</math> and <math>2^2 + 1 = 5</math>, and 5 is divisible by 5.</li> <li>For <math>n = 2</math>, <math>4n + 1 = 9</math>. The smallest <math>(m)</math> is 3 because <math>2^3 + 1 = 9</math>, and 9 is divisible by 9.</li> </ul> <p>The approach is to iterate over increasing values of <math>(m)</math> starting from 1, and for each <math>(m)</math> check the divisibility conditions. The first <math>(m)</math> that satisfies either condition is the answer. Here's the implementation in Python:</p> <pre>def solution(x: int) -&gt; int:     n = x     divisor = 4 * n + 1      m = 1     while True:         if (2**m + 1) % divisor == 0 or (2**m - 1) % divisor == 0:             return m         m += 1</pre> <p>This code iterates over possible values of <math>(m)</math> starting from 1, checking the divisibility conditions for each <math>(m)</math>. It returns the first <math>(m)</math> that satisfies either condition, ensuring the solution is found with the lowest time complexity possible given the problem constraints.</p>	<p>----- Turn 2: &lt; Coding &gt; -----</p> <pre>def solution(x: int) -&gt; int:     import math     def euler_totient(n):         result = n         p = 2         while p * p &lt;= n:             if n % p == 0:                 while n % p == 0:                     n //= p                 result -= result // p             p += 1         if n &gt; 1:             result -= result // n         return result      def order_of_2_modulo(n):         phi = euler_totient(n)         for d in range(1, phi + 1):             if phi % d == 0:                 if pow(2, d, n) == 1 or pow(2, d, n) == n - 1:                     return phi                     return d      mod = 4 * x + 1     return order_of_2_modulo(mod)</pre>

Figure 10: GPT-4o solves UTMath\_948 by the PoT method, by the RCoT method, respectively. The input prompt is omitted here but can be found in Appendix C. PoT simply performs brute-force solving, while RCoT involves deeper reasoning through Case merging after a classification discussion and the application of Euler's formula, providing a solution with lower time complexity.

<p><b>Task_Id:</b> UTMath_629 <b>OEIS_Id:</b> A002309 <b>Main_Category:</b> Polynomial and Series Expansions</p> <p><b>Problem Statement:</b> The definition of <math>a(n)</math> is the sum of the fourth powers of the first <math>n</math> odd numbers.</p> <p><b>Case 1:</b> <math>a(1) = 1</math> <b>Case 2:</b> <math>a(2) = 82</math> ...</p> <p><b>Hard Case 1:</b> <math>a(63647) = 334225666620196206703263</math> <b>Hard Case 2:</b> <math>a(63648) = 334251823678794269653888</math> ...</p> <p><b>Hard Case 10:</b> <math>a(64656) = 3344619392264074754875880</math></p>	<p><b>The Responses of GPT-4o by Reasoning-To-Coding Thoughts of Thought Methods</b></p> <p>----- Turn 1: &lt; Reasoning &gt; -----</p> <p>To analyze the sequence problem where <math>a(n)</math> is defined as the sum of the fourth powers of the first <math>n</math> odd numbers, we need to derive a formula or method to compute <math>a(n)</math> efficiently.</p> <p><b>Step-by-Step Analysis:</b> <b>Understanding the Sequence:</b> The sequence involves summing the fourth powers of the first <math>n</math> odd numbers. The first <math>n</math> odd numbers are:</p> $1, 3, 5, 7, \dots, (2n - 1)$ <p><b>Expressing the Sequence:</b> The <math>n</math>th odd number can be expressed as <math>(2k - 1)</math> for <math>k = 1</math> to <math>n</math>. Therefore:</p> $a(n) = (1^4) + (3^4) + (5^4) + \dots + ((2n - 1)^4)$ <p><b>General Term:</b> The general term for the <math>k</math>th odd number is <math>(2k - 1)</math>. The fourth power of this term is:</p> $(2k - 1)^4$ <p><b>Summation:</b> We need to compute the sum:</p> $a(n) = \sum_{k=1}^n (2k - 1)^4$ <p><b>Simplifying the Expression:</b> Expand <math>(2k - 1)^4</math> using the binomial theorem:</p> $(2k - 1)^4 = (16k^4 - 32k^3 + 24k^2 - 8k + 1)$ <p>Therefore, the sum becomes:</p> $a(n) = \sum_{k=1}^n (16k^4 - 32k^3 + 24k^2 - 8k + 1)$ <p><b>Separate the Summation:</b> Break down the summation into individual components:</p> $a(n) = 16 \sum_{k=1}^n k^4 - 32 \sum_{k=1}^n k^3 + 24 \sum_{k=1}^n k^2 - 8 \sum_{k=1}^n k + \sum_{k=1}^n 1$ <p><b>Use Known Summation Formulas:</b> Use the formulas for the sum of powers of integers:</p> $\sum_{k=1}^n k = \frac{n(n+1)}{2}, \sum_{k=1}^n k^2 = \frac{n(n+1)(2n+1)}{6}, \sum_{k=1}^n k^3 = \left(\frac{n(n+1)}{2}\right)^2, \sum_{k=1}^n k^4 = \frac{n(n+1)(2n+1)(3n^2+3n+1)}{30}$ <p><b>Substitute and Simplify:</b> Substitute these formulas into the expression for <math>a(n)</math> and simplify to get a closed-form expression.</p> <p><b>Time Complexity:</b> The derived formula will allow us to compute <math>a(n)</math> in constant time, <math>O(1)</math>, since it involves only arithmetic operations on <math>n</math>. By following these steps, we can derive a closed-form expression for <math>a(n)</math> that allows us to compute the sum of the fourth powers of the first <math>n</math> odd numbers efficiently, without iterating through each term. This approach ensures the lowest time complexity, <math>O(1)</math>, for computing the sequence.</p>
<p><b>The Responses of GPT-4o by Program of Thought Methods</b></p> <p>To solve the problem of finding the sum of the fourth powers of the first <math>n</math> odd numbers, we need to follow these steps:</p> <p><b>Identify the Sequence of Odd Numbers:</b> The sequence of the first <math>n</math> odd numbers is <math>1, 3, 5, 7, \dots</math> up to the <math>n</math>-th odd number. The <math>k</math>-th odd number can be expressed as <math>2k - 1</math>.</p> <p><b>Calculate the Fourth Power:</b> For each odd number in the sequence, calculate its fourth power. The fourth power of a number <math>x</math> is <math>x^4</math>.</p> <p><b>Sum the Fourth Powers:</b> Sum up all the fourth powers calculated in the previous step to get the result for <math>a(n)</math>.</p> <p><b>Optimize the Calculation:</b> Instead of recalculating the odd numbers each time, we can directly compute the sum of the fourth powers using a loop that iterates <math>n</math> times, calculating <math>(2k - 1)^4</math> for each <math>k</math> from 1 to <math>n</math>.</p> <p>The time complexity of this approach is <math>O(n)</math> because we are iterating through the first <math>n</math> odd numbers and performing a constant amount of work (raising to the fourth power and summing) for each.</p> <p>Here is the implementation in Python:</p> <pre>def solution(x: int) -&gt; int:     total_sum = 0     for k in range(1, x + 1):         odd_number = 2 * k - 1         total_sum += odd_number ** 4     return total_sum</pre>	<p>----- Turn 2: &lt; Coding &gt; -----</p> <pre>def solution(x: int) -&gt; int:     # Using the derived formula for the sum of fourth powers of the first n odd numbers     n = x     sum_k = n * (n + 1) // 2     sum_k2 = n * (n + 1) * (2 * n + 1) // 6     sum_k3 = (n * (n + 1) // 2) ** 2     sum_k4 = n * (n + 1) * (2 * n + 1) * (3 * n ** 2 + 3 * n - 1) // 30      result = 16 * sum_k4 - 32 * sum_k3 + 24 * sum_k2 - 8 * sum_k + n     return result</pre>

Figure 11: GPT-4o solves UTMath\_629 by the PoT method, by the RCoT method, respectively. The input prompt is omitted here but can be found in Appendix C. PoT only performs brute-force traversal, whereas RCoT engages in deeper reasoning by associating the problem with the binomial theorem and using the power sum formula, yielding a closed-form expression with lower time complexity from  $O(n)$  to  $O(1)$ .