

MSCR: Exploring the Vulnerability of LLMs’ Mathematical Reasoning Abilities Using Multi-Source Candidate Replacement

Zhishen Sun¹ Guang Dai² Haishan Ye^{1,2†}
¹Xi’an Jiaotong University ²SGIT AI Lab

Abstract

LLMs demonstrate performance comparable to human abilities in complex tasks such as mathematical reasoning, but their robustness in mathematical reasoning under minor input perturbations still lacks systematic investigation. Existing methods generally suffer from limited scalability, weak semantic preservation, and high costs. Therefore, we propose MSCR, an automated adversarial attack method based on multi-source candidate replacement. By combining three information sources including cosine similarity in the embedding space of LLMs, the WordNet dictionary, and contextual predictions from a masked language model, we generate for each word in the input question a set of semantically similar candidates, which are then filtered and substituted one by one to carry out the attack. We conduct large-scale experiments on LLMs using the GSM8K and MATH500 benchmarks. The results show that even a slight perturbation involving only a single word can significantly reduce the accuracy of all models, with the maximum drop reaching 49.89% on GSM8K and 35.40% on MATH500, while preserving the high semantic consistency of the perturbed questions. Further analysis reveals that perturbations not only lead to incorrect outputs but also substantially increase the average response length, which results in more redundant reasoning paths and higher computational resource consumption. These findings highlight the robustness deficiencies and efficiency bottlenecks of current LLMs in mathematical reasoning tasks.

1 Introduction

Large language models (LLMs) have achieved breakthrough progress in natural language processing, with the continuous expansion of model parameter sizes and the widespread use of high-quality training data. Furthermore, by relying on the deep integration of technologies like Chain-of-Thought (CoT) (Wei et al., 2022; Zhou et al., 2022) and reinforcement learning (Schulman et al., 2017; Shao et al., 2024; Yu et al., 2025; Zheng et al., 2025), LLMs have demonstrated capabilities close to human-level performance in complex tasks such as code generation, especially mathematical reasoning and calculation (Yang et al., 2025; Guo et al., 2025; OpenAI, 2025; Comanici et al., 2025). This development has attracted widespread attention across society.

However, while LLMs have demonstrated remarkable capabilities in various complex tasks and achieved nearly perfect performance on multiple mathematical benchmarks, such as GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021), the vulnerability of their mathematical reasoning capabilities still lack systematic and in-depth research. Existing text adversarial attack methods for exploring language model vulnerability, such as TextFooler (Jin et al., 2020) and BERT-ATTACK (Li et al., 2020), primarily focus on classification models and are difficult to directly transfer to currently mainstream generative LLMs. Consequently, these methods suffer from significant limitations in practicality and scalability. Therefore, researches have gradually shifted to methods such as jailbreak attack

[†]Corresponding to: yehaishan@xjtu.edu.cn

Original question: Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?	
Here's how to break down the problem step-by-step:	
1. Josh spent \$80,000 (purchase price) + \$50,000 (repairs) = \$130,000;	
2. The value increased by 150% of the initial cost, so the increase is $\$80,000 \times 1.50 = \$120,000$;	
3. The new value is the initial cost plus the increase, which is $\$80,000 + \$120,000 = \$200,000$;	
4. Profit is the new value minus the total cost: $200,000 - 130,000 = \$70,000$;	
So, the profit is \boxed{70000}.	✓
Perturbation question: he decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?	
Here's how to break down the problem step-by-step:	
1. The initial cost of the house is \$80,000;	
2. He puts in \$50,000 in repairs, so the total cost is $\$80,000 + \$50,000 = \$130,000$;	
3. The repairs increased the value of the house by 150%. This means the new value of the house is 150% of the original value, which is \$80,000;	
4. To find the new value, multiply the original value by 150% : $\$80,000 \times 1.5 = \$120,000$;	
5. However, this is not the final value of the house. The repairs added value to the house, so the new value is the original value plus the added value: $\$120,000 + \$50,000 = \$170,000$;	
6. The profit is the difference between the selling price and the total cost: $\$170,000 - \$130,000 = \$40,000$;	
So, the profit is \boxed{40000}.	✗

Figure 1: Examples of adversarial samples generated by the MSCR algorithm on the GSM8K benchmark, answered using Meta-Llama-3-70B-Instruct. More examples are presented in Appendix D.

(Yong et al., 2023; Wei et al., 2023) to reveal their vulnerability. Most existing work focuses on general tasks, with little systematic exploration of the critical capability of mathematical reasoning. Furthermore, these attack methods often rely on manual design and repeated debugging of prompts, which have significant cost challenges. Meanwhile, some studies have perturbed or reconstructed existing datasets to generate new evaluation datasets to test mathematical reasoning capabilities of LLMs (Mirzadeh et al., 2024; Huang et al., 2025). These methods often require manual construction of perturbation rules and lack automated and scalable attack processes.

To address the issues of the previous methods and systematically evaluate the vulnerability of LLMs in mathematical reasoning tasks, we propose MSCR, an automated adversarial attack method based on multi-source candidate replacement. The core idea of this method is: for each word in the input question, we first generate a set of semantically similar and grammatically reasonable candidate replacement words using three complementary information sources: the model’s own knowledge, the WordNet dictionary, and knowledge from an external masked language model (MLM). These candidates are then filtered to ensure their quality. These candidates are sequentially substituted into the original question and solved using the target model. A successful attack is considered when the modified question causes the target model to produce an incorrect answer. If the candidate words for the current word are exhausted and still unsuccessful, the attack continues with the next word until the attack succeeds or all words fail to trigger a model error. Our method not only automatically generates high-quality replacement candidates but can also effectively attack LLMs by modifying just a single word without any human intervention (As shown in Figure 1, simply changing “Josh” to “he” in the question causes LLMs to produce an incorrect answer due to an error in its reasoning process), thereby revealing their potential vulnerabilities in mathematical reasoning tasks.

We successfully apply our algorithmic framework to two mathematical task benchmarks of varying difficulty, and conduct systematic attack experiments on twelve open-source LLMs, while also applying the generated adversarial examples to two commercial LLMs. The experimental results show that this framework can significantly reduce the accuracy of almost all large language models on different benchmark tasks with only a slight perturbation of a single word. As shown in Figure 2, on GSM8K, the accuracy of nearly all open-source models drops by more than 20%, with Qwen2.5-Math-1.5B-Instruct experiencing the largest decrease of 49.89%. On MATH500, the accuracy of all open-source models decreases by over 10%, with gemma-2-9b-it showing the largest drop of 35.40%. Furthermore, when the generated adversarial examples are transferred to the commercial LLMs OpenAI-o3 and GPT-4o, both models also exhibit significant accuracy reductions: on GSM8K, their accuracy drops by 15.88% and 17.21%, and on MATH500, the drops are 6.80% and 7.80%. This phenomenon indicates that LLMs are highly sensitive to slight changes in the input text

in mathematical reasoning tasks, highlighting their vulnerability and lack of robustness in mathematical reasoning. Further analysis also find that after being attacked, the length of the answers generated by the model increasing significantly compare to when it is not attacked. This not only makes its reasoning process more lengthy and complex, but also increases the consumption of computing resources to a certain extent, causing a significant increase in reasoning cost and reducing overall reasoning efficiency. Our contributions can be summarized as follows:

- We propose MSCR, an automated adversarial attack method based on multi-source candidate replacement. This method can generate adversarial samples by perturbing one word, thereby systematically revealing the vulnerability of LLMs in mathematical task.
- We conduct extensive attack experiments on twelve LLMs across two different math benchmarks to comprehensively evaluate the effectiveness of our proposed method, and we apply the generated adversarial samples to two commercial LLMs. The results show that even a slight perturbation of a single word in the question text can significantly reduce the reasoning accuracy of nearly all models across various math benchmarks.
- We further analyze the experimental results and the responses generated by LLMs, finding that LLMs not only exhibit significant robustness issues when facing minor input perturbations, but also show a significant increase in the length of the generated responses. This phenomenon indicates that when LLMs are under attack, their reasoning path tend to be lengthy and complex, resulting in additional consumption of computing resources, increased reasoning costs, and ultimately reduce overall reasoning efficiency.

2 Related work

Text adversarial attacks have recently garnered widespread attention as a key method for evaluating the robustness of natural language processing models. Ebrahimi et al. (2017) achieves misclassification through gradient-guided character replacement. Jin et al. (2020) generates candidate replacement words based on word embeddings and greedy search, significantly improving readability. With the rise of pre-trained models, Li et al. (2020) significantly improves attack efficiency by leveraging MLMs and word importance ranking, but its effectiveness in context-dependent tasks is limited. Garg & Ramakrishnan (2020) makes some improvements in semantic preservation but still faces the challenge of a large search space. To further overcome the contradiction between discrete optimization and semantic constraints, Guo et al. (2021) introduces Gumbel-softmax for continuous optimization and combines model perplexity and BERTScore as differentiable constraints, achieving stronger cross-model transfer while maintaining semantic similarity. Wang et al. (2023) introduces gradient optimization to directly search the entire vocabulary space, significantly improving attack success rate and demonstrating the cross-model transferability of adversarial examples. Zhao et al. (2022) adopts a two-stage optimization method of greedy keyword screening and genetic algorithm refinement to achieve extremely low perturbation rate in long text tasks.

In recent years, with the widespread adoption of LLMs, researchers have proposed jailbreak attacks, which use adversarial prompts to bypass security alignment restrictions and generate violating content. Yong et al. (2023) relied on manually constructed semantically coherent templates. While the approach offer some stealth, it suffer from limited generalization and are easily offset by iterative model updates. To improve automation, Zou et al. (2023) first formalized the jailbreak problem as a discrete optimization task, achieved impressive success rates. However, the generated garbled suffixes are easily intercepted by perplexity checks, resulting in insufficient robustness. Subsequently, Liu et al. (2023) utilizes a hierarchical genetic algorithm to generate semantically coherent adversarial prompts, effectively circumventing perplexity checks and significantly improving the attack success rate of aligned models. Yao et al. (2024) achieves cross-model vulnerability detection by




Figure 2: Performance changes of various LLMs under attacks by the MSCR algorithm.

constructing constraint templates and problem sets. Jia et al. (2024) integrates multi-coordinate updates with harmful target templates, achieving a significant attack success rate. Meanwhile, Chao et al. (2025) automated jailbreak attacks using a dual-model adversarial framework, efficiently jailbreak models like GPT-4 and Gemini-Pro within 20 queries. This method leverages black-box interaction to generate semantic jailbreak prompts, significantly reducing labor costs, but the prompt diversity is limited by predefined strategies.

The performance of LLMs in mathematical reasoning tasks has attracted widespread attention, but doubts about their essential capabilities have gradually emerged. Shi et al. (2023) perturbs the GSM8K benchmark by adding redundant information that is semantically related to the original problem but does not affect the solution, indicating that LLMs are easily disturbed by irrelevant context in arithmetic reasoning. Even with the CoT, self-consistent decoding or prompt enhancement, it is still difficult to cure logical confusion; Li et al. (2024) generates eight types of variants from five dimensions, including numerical modification and deletion, sentence insertion, and problem reconstruction, revealing the shortcomings of LLMs in critical thinking and operational generalization; Mirzadeh et al. (2024) uses parsable symbolic templates to modify numerical values and proper nouns, generates large-scale samples, and evaluates model performance through accuracy distribution,

further revealing the nature of LLMs that is highly dependent on surface pattern matching. Its performance fluctuates greatly under numerical changes or redundant sentence interference; Shrestha et al. (2025) constructed a new benchmark using six different levels of numerical perturbations, verifying that LLMs have a numerical generalization bottleneck; Huang et al. (2025) systematically revealed, by constructing MATH-P-Simple and MATH-P-Hard, that LLMs’ accuracy dropped significantly when faced with problems with similar semantics but fundamentally different logical structures, suggesting that they rely on pattern memory rather than deep reasoning.

3 Method

MSCR is an automated adversarial attack method for multi-source candidate replacements. This method uses three sources of information: the cosine similarity of word vectors in the model embedding space, the WordNet dictionary, and a MLM. This method provides synonyms or antonyms as the initial set of candidate replacement words and then filters to form the final set of candidate replacement words. The method also supports replacement operations, replacing the original word with a candidate word of similar meaning, to achieve the desired perturbation effect.

Next, we will discuss the three sources of information for providing candidate replacement words and the specific algorithmic framework.

3.1 Perturbation word information sources

Cosine similarity of word vectors. Because text data is essentially a discrete sequence of symbols, it’s difficult to directly apply continuous, subtle perturbations to the input space. Furthermore, since word vector embedding can capture more nuanced semantic and grammatical features in a high-dimensional continuous space. Therefore, one of the information sources for our set of candidate replacement words is the cosine similarity of word vectors in the embedding space. Specifically, we first use the embedding matrix of the model, which maps each word in the vocabulary to a dense vector representation. To avoid bias caused by differences in word vector modulus length, we normalize all word vectors before calculation, ensuring that they lie on the unit hypersphere. Next, for the target word w in the input math problem, the algorithm obtains its corresponding normalized vector v_w and sequentially calculates the cosine similarity with the other candidate vectors v_i in the embedding space: $\text{sim}(v_w, v_i) = \frac{v_w \cdot v_i}{\|v_w\| \|v_i\|}$. This metric effectively measures the directional proximity of words in the semantic space. In this way, we obtain a preliminary set of candidate replacement words.

Building on this foundation, we add a sophisticated filtering mechanism to further screen candidates with high similarity rankings. Specifically, we first retrieved the top 10 semantically closest candidates from the aforementioned candidate words based on cosine similarity. These candidates then underwent multi-layered filtering: First, we implemented character-level filtering to remove control characters, non-Latin characters, and pure punctuation to ensure text structural integrity and readability. Second, we introduced substring matching and morphological constraints. For candidates that were perturbed solely on the original word, we retained words containing the original word stem or minor orthographic changes, thereby maintaining semantic consistency within local phrases or specialized terminology.

The WordNet dictionary. In order to make up for the lack of semantic generalization of the cosine similarity perturbation method based on word vectors, while taking into account the limitations of the model’s own knowledge, the WordNet dictionary is introduced as a second source of candidate replacement words. This provides high-quality synonym candidates based on a semantic network constructed by human experts. For the target word w , the algorithm first queries all its synonym sets in WordNet and traverses all word units under each synonym set to extract words that are clearly defined as having a synonym relationship with w at the lexicographic level. After obtaining a preliminary set of candidate replacement words from WordNet, we calculated the cosine similarity between the candidate words and




Figure 3: Overview of the MSCR attack flow.

the original word, sorted the candidate set in descending order by similarity, and retained the top five most relevant candidates.

Masked language model. Given the outstanding performance of MLMs in natural language processing tasks, especially their powerful capabilities in various downstream tasks as well as their small parameter size and storage occupancy, this algorithm introduces MLMs prediction as the third source of information to generate replacement vocabulary that is highly compatible with the specific context. Specifically, the algorithm first replaces the target word w with a special mask token to form a structured input sequence, and then inputs it into a MLM. The model then computes the conditional probability distribution for each word in the vocabulary at that mask position. Based on this probability distribution, the algorithm adopts a probability-based filtering mechanism: only high-confidence candidates with a predicted probability greater than 0.1 are retained, and low-probability noise candidates that are grammatically correct but semantically weak or illogical are significantly eliminated, ensuring that the candidate set contains the semantically adapted items with the highest model confidence, thereby improving the quality of the candidate set.

The final candidate replacement word set consists of the union of the three aforementioned sub-candidate sets. Through this multi-source candidate word integration strategy, the algorithm effectively combines the model’s internal distributed knowledge with external structured knowledge, covering diverse domains and semantic scenarios. This generates a more natural, semantically similar, and high-quality candidate replacement word set for adversarial attacks.

3.2 Attack algorithm process

By summarizing the research of (Shi et al., 2023; Li et al., 2024; Mirzadeh et al., 2024; Shrestha et al., 2025), and others, we find that when current LLMs face mathematical reasoning tasks, any slight perturbation may cause the LLMs to produce incorrect results. To better explore the vulnerabilities of LLMs in mathematical reasoning tasks, for each target word in the original input sentence, we first generate a corresponding set of candidate replacement words using the three aforementioned information sources, then perturb the original problem. When replacing perturbed words, the cosine similarity candidate words of the word vectors are primarily used for local replacement of the current word, WordNet and MLM candidate words use a global replacement strategy, replacing all words in the sentence that match the original word, ensuring the readability and grammatical integrity of the generated text. After each candidate word replacement, the algorithm constructs a new adversarial input and calls the target model for reasoning and solving. Finally, when determining whether the attack is successful, the algorithm employs a two step evaluation mechanism. First, the output of the target model is compared to the original answer

to check whether the replacement has successfully altered the model’s prediction. If the current candidate word replacement causes the model’s output to differ from the original answer, it is considered a preliminary attack success, and the algorithm immediately returns the generated adversarial example along with the corresponding model response. If the output is consistent, it is considered an attack failure, and the algorithm proceeds to try the next candidate word in priority order, continuing until all candidates are exhausted or a model error is triggered. Next, for the questions and answers judged as preliminary attack successes, the algorithm further invokes an advanced commercial LLM for secondary evaluation to enhance the robustness of the assessment. If this commercial model confirms that the target model has indeed produced an incorrect answer for the perturbed question, it is regarded as a final attack success. The detailed algorithm flow is shown in Figure 3.

4 Experiment

In this section, we introduce our experimental setup (Sec. 4.1) and detailed experimental results of our algorithm (Sec. 4.2).

4.1 Experimental setup

Models: In our experiments, we evaluate a wide variety of LLMs, including reasoning models, math-specific models, and general-purpose models:

- Qwen: Qwen2.5-Math-1.5B-Instruct, Qwen2.5-Math-7B-Instruct (Yang et al., 2024);
- DeepSeek: DeepSeek-R1-Distill-Qwen-1.5B, DeepSeek-R1-Distill-Qwen-7B, DeepSeek-R1-Distill-Qwen-32B (Guo et al., 2025);
- Meta Llama: Meta-Llama-3-8B-Instruct, Meta-Llama-3-70B-Instruct (AI@Meta, 2024);
- Gemma: gemma-2-9b-it, gemma-2-27b-it (Team, 2024), gemma-3-4b-it, gemma-3-12b-it, gemma-3-27b-it (Team, 2025).

Additionally, we run transferability experiments of the adversarial samples on OpenAI-o3 (OpenAI, 2025) and GPT-4o (Hurst et al., 2024) to further evaluate the attack effectiveness of MSCR. The MLM used is bert-large-uncased (Devlin et al., 2018). For secondary evaluation, the commercial LLM used is qwen3-max-preview.

Dataset: To evaluate the performance of MSCR, our experiments use the GSM8K and MATH500 datasets. The GSM8K dataset assesses basic mathematical reasoning, and the MATH500 dataset evaluates more advanced reasoning capabilities.

Parameter Settings: All LLMs set the temperature parameter to 0.6 to balance exploration and exploitation during generation. For each model and each dataset, we run the experiment three times independently and average the results to ensure the credibility of the experimental results.

4.2 Experimental results

In Sec. 4.2.1, we present the main performance of our algorithm, revealing the vulnerability of LLMs’ mathematical reasoning capabilities when facing minor perturbations. In Sec. 4.2.2, we demonstrate the changes in the length of responses of LLMs when facing minor perturbations under our attack framework, revealing that minor perturbations increase the inference cost of LLMs. The full results with standard deviations are presented in Appendix B.

4.2.1 Main performance

We conduct large-scale adversarial attack experiments using twelve LLMs on two mathematical reasoning benchmarks of different difficulty levels, GSM8K and MATH500, with

Table 1: The performance of various LLMs under the MSCR algorithm attack (accuracy, %).

Model	GSM8K			MATH500		
	Original	Attack	Decrease	Original	Attack	Decrease
DeepSeek-R1-Distill-Qwen-32B	94.39	75.59	18.80↓	93.00	79.20	13.80↓
DeepSeek-R1-Distill-Qwen-7B	93.71	68.16	25.55↓	92.80	78.60	14.20↓
DeepSeek-R1-Distill-Qwen-1.5B	83.09	57.54	25.55↓	81.20	63.20	18.00↓
Qwen2.5-Math-7B-Instruct	94.09	54.81	39.28↓	83.20	61.60	21.60↓
Qwen2.5-Math-1.5B-Instruct	85.90	36.01	49.89↓	76.20	49.80	26.40↓
Meta-Llama-3-70B-Instruct	93.86	65.43	28.43↓	57.60	45.00	12.60↓
Meta-Llama-3-8B-Instruct	83.02	47.69	35.33↓	30.60	6.60	24.00↓
gemma-3-27b-it	98.10	77.79	20.31↓	92.00	75.00	17.00↓
gemma-3-12b-it	97.73	75.51	22.22↓	89.80	70.40	19.40↓
gemma-3-4b-it	94.54	63.84	30.70↓	81.40	56.20	25.20↓
gemma-2-27b-it	92.65	59.74	32.91↓	68.30	39.00	29.30↓
gemma-2-9b-it	90.75	53.07	37.68↓	47.00	11.60	35.40↓

Table 2: Performance of commercial LLMs under other adversarial samples (accuracy, %).

Model	GSM8K			MATH500		
	Original	Attack	Decrease	Original	Attack	Decrease
OpenAI-o3	97.41	81.53	15.88↓	98.20	91.40	6.80↓
GPT-4o	96.28	79.07	17.21↓	76.00	68.20	7.80↓

results summarized in Table 1. Our findings show that even when applying only a single-word perturbation to the input question, the attack framework can significantly weaken the problem-solving ability of LLMs, leading to substantial increases in error rates. On the perturbed GSM8K benchmark, all models except DeepSeek-R1-Distill-Qwen-32B experience an accuracy drop of more than 20%, with half of them dropping by over 30%. Qwen2.5-Math-1.5B-Instruct suffers the most severe degradation, with accuracy falling by 49.89%. On the more challenging MATH500 benchmark, all models show an accuracy drop of over 10% after perturbation, and half of them decline by more than 20%. For example, gemma-2-9b-it’s accuracy decreases by 35.40%, corresponding to a 75.32% relative reduction, while Meta-Llama-3-8B-Instruct drops from 30.60% to only 6.60%, corresponding to a 78.43% relative reduction. Moreover, as shown in Table 1, within the same model family, models with larger parameter sizes and better baseline performance generally suffer smaller accuracy drops, demonstrating relatively higher robustness. However, even these stronger models still experience reductions of over 20% on GSM8K and over 10% on MATH500. These results collectively reveal the significant vulnerability of current LLMs in mathematical reasoning tasks, indicating that even larger and stronger models are highly susceptible to being misled by low-cost input perturbations, which ultimately leads to incorrect answers.

To further evaluate the attack effectiveness of MSCR, we examine the transferability of its generated adversarial samples across different models. Specifically, we transfer the perturbed questions generated by MSCR targeting gemma-3-27b-it to OpenAI-o3 and GPT-4o for testing. As shown in Table 2, both models exhibit significant performance degradation on the perturbed GSM8K and MATH500 benchmarks. On GSM8K, the accuracy of OpenAI-o3 and GPT-4o drops by 15.88% and 17.21%, respectively, while on MATH500 the drops are 6.80% and 7.80%, respectively, with error rates of model-generated answers increasing sharply. These results provide strong evidence that MSCR achieves effective attacks, not only significantly reducing model accuracy but also demonstrating strong cross-model transferability.

Table 3: The response length of various LLMs under the MSCR algorithm attack. We use the number of tokens in the response of LLMs to measure the response length.

Model	GSM8K			MATH500		
	Original	Attack	Ratio	Original	Attack	Ratio
DeepSeek-R1-Distill-Qwen-32B	727.0	1294.3	$1.78 \times$	1668.9	1931.3	$1.16 \times$
DeepSeek-R1-Distill-Qwen-7B	1204.5	2190.2	$1.82 \times$	2204.7	3081.9	$1.40 \times$
DeepSeek-R1-Distill-Qwen-1.5B	1754.3	2698.0	$1.54 \times$	2521.6	3303.8	$1.31 \times$
Qwen2.5-Math-7B-Instruct	305.4	331.3	$1.08 \times$	692.5	995.1	$1.44 \times$
Qwen2.5-Math-1.5B-Instruct	306.8	330.5	$1.08 \times$	579.0	745.9	$1.29 \times$
Meta-Llama-3-70B-Instruct	170.4	207.2	$1.22 \times$	283.5	304.9	$1.08 \times$
Meta-Llama-3-8B-Instruct	174.3	203.1	$1.17 \times$	426.1	515.3	$1.21 \times$
gemma-3-27b-it	273.9	523.8	$1.91 \times$	698.3	1033.8	$1.48 \times$
gemma-3-12b-it	310.8	608.8	$1.96 \times$	755.4	1165.2	$1.54 \times$
gemma-3-4b-it	411.8	881.7	$2.14 \times$	926.3	1308.3	$1.41 \times$
gemma-2-27b-it	173.7	188.0	$1.08 \times$	359.1	378.2	$1.05 \times$
gemma-2-9b-it	174.6	190.0	$1.09 \times$	396.8	413.8	$1.04 \times$

4.2.2 Response length collapse

Next, we further analyze the changes in response length generated by LLMs when facing perturbed questions compared to their original responses. As shown in Table 3, under our perturbation framework, models often generate longer responses when handling perturbed questions, and in some cases, the average response length even approaches or exceeds twice the original average response length. This phenomenon is more pronounced in more powerful models, but less pronounced in relatively weaker models. For example, the gemma2 series models exhibit almost no noticeable change in response length. We speculate that this may be because weaker models do not have true mathematical reasoning capabilities. Their problem-solving approach is closer to memory-based or template-based formal solutions, often classifying questions by identifying keywords and directly applying existing patterns to generate answers. As a result, when facing with perturbed questions, they do not generate additional reasoning processes, leading to minimal changes in response length.

Furthermore, we visualize and conduct more detailed statistical analyses of the response lengths of different models when handling various questions. The results are shown in Figure 4. Whether on the GSM8K or MATH500, models generate responses that are longer than the original response lengths for most perturbed questions, and in some cases, the response length even reaches 4 to 10 times the original length. In some reasoning models, the response length even exceeds 10 times the original length, resulting in a very severe response length collapse. These phenomena indicate that current LLMs are highly susceptible to minor perturbations, which makes the reasoning process longer and more complex, further increasing the reasoning cost and computational resource consumption, and ultimately reducing overall reasoning efficiency.

5 Conclusion

This paper proposes MSCR, an automated adversarial attack method based on multi-source candidate replacements. By combining cosine similarity in the embedding space of LLMs, the WordNet dictionary, and contextual predictions from a MLM, it generates high-quality candidate replacements and enables a systematic evaluation of the robustness of LLMs in mathematical reasoning. Extensive experiments on 12 LLMs using the GSM8K and MATH500 benchmarks show that this method, by applying only a slight single-word perturbation to the input question, can significantly reduce the mathematical reasoning accuracy of nearly all models, with a maximum drop of 49.89%. Meanwhile, we conduct adversarial sample transferability experiments on OpenAI-o3 and GPT-4o, and the performance of




Figure 4: The distribution of the ratio between the response length of LLMs for perturbed questions and the original length. Here, we only present the visualization results for DeepSeek-R1-Distill-Qwen-32B and gemma-3-27b-it; visualizations for additional models can be found in the Appendix C.

these two commercial LLMs is also significantly degraded. Notably, the attack not only exposes the models’ general vulnerability in reasoning but also causes a substantial decline in efficiency. The average length of the generated responses increases dramatically, reaching up to 2.14 times the original length, some reasoning models even exhibit a response length collapse exceeding 10 times the original length, resulting in longer reasoning paths and a sharp rise in computational resource consumption. These findings strongly reveal the significant robustness challenges and efficiency bottlenecks faced by current LLMs in mathematical reasoning tasks. The automated and scalable attack method proposed in this work provides a crucial basis and important direction for systematically evaluating model vulnerabilities, understanding their reasoning mechanisms, and designing more robust and efficient defense strategies in the future.

References

- AI@Meta. Llama 3 model card. 2024. URL https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. In 2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML), pp. 23–42. IEEE, 2025.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. arXiv preprint arXiv:2110.14168, 2021.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva,INDERJIT Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. arXiv preprint arXiv:2507.06261, 2025.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. CoRR, abs/1810.04805, 2018. URL <http://arxiv.org/abs/1810.04805>.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. Hotflip: White-box adversarial examples for text classification. arXiv preprint arXiv:1712.06751, 2017.
- Siddhant Garg and Goutham Ramakrishnan. Bae: Bert-based adversarial examples for text classification. arXiv preprint arXiv:2004.01970, 2020.
- Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. Gradient-based adversarial attacks against text transformers. arXiv preprint arXiv:2104.13733, 2021.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948, 2025.
- 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. arXiv preprint arXiv:2103.03874, 2021.
- Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle Cai, Hui Yuan, Runzhe Wang, et al. Math-perturb: Benchmarking llms’ math reasoning abilities against hard perturbations. arXiv preprint arXiv:2502.06453, 2025.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. arXiv preprint arXiv:2410.21276, 2024.
- Xiaojun Jia, Tianyu Pang, Chao Du, Yihao Huang, Jindong Gu, Yang Liu, Xiaochun Cao, and Min Lin. Improved techniques for optimization-based jailbreaking on large language models. arXiv preprint arXiv:2405.21018, 2024.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In Proceedings of the AAAI conference on artificial intelligence, volume 34, pp. 8018–8025, 2020.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. Bert-attack: Adversarial attack against bert using bert. arXiv preprint arXiv:2004.09984, 2020.
- Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. Gsm-plus: A comprehensive benchmark for evaluating the robustness of llms as mathematical problem solvers. arXiv preprint arXiv:2402.19255, 2024.
- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. arXiv preprint arXiv:2310.04451, 2023.
- Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncl Tuzel, Samy Bengio, and Mehrdad Farajtabar. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. arXiv preprint arXiv:2410.05229, 2024.
- OpenAI. O3 and o4 mini system card, 2025. URL <https://openai.com/index/o3-o4-mini-system-card/>.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. arXiv preprint arXiv:2402.03300, 2024.

- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pp. 31210–31227. PMLR, 2023.
- Safal Shrestha, Minwu Kim, and Keith Ross. Mathematical reasoning in large language models: Assessing logical and arithmetic errors across wide numerical ranges. *arXiv preprint arXiv:2502.08680*, 2025.
- Gemma Team. Gemma. 2024. doi: 10.34740/KAGGLE/M/3301. URL <https://www.kaggle.com/m/3301>.
- Gemma Team. Gemma 3. 2025. URL <https://goo.gle/Gemma3Report>.
- Yimu Wang, Peng Shi, and Hongyang Zhang. Gradient-based word substitution for obstinate adversarial examples generation in language models. *arXiv preprint arXiv:2307.12507*, 2023.
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems*, 36:80079–80110, 2023.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Dongyu Yao, Jianshu Zhang, Ian G Harris, and Marcel Carlsson. Fuzzllm: A novel and universal fuzzing framework for proactively discovering jailbreak vulnerabilities in large language models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4485–4489. IEEE, 2024.
- Zheng-Xin Yong, Cristina Menghini, and Stephen H Bach. Low-resource languages jailbreak gpt-4. *arXiv preprint arXiv:2310.02446*, 2023.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.
- Xingyi Zhao, Lu Zhang, Depeng Xu, and Shuhan Yuan. Generating textual adversaries with minimal perturbation. *arXiv preprint arXiv:2211.06571*, 2022.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, et al. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*, 2025.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022.
- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.

A Usage Statement for Large Language Models

In this paper, we use LLMs for text polishing. Specifically, for certain paragraph expressions, we use LLMs to improve the fluency of the text. In addition, we also employ LLMs for literature retrieval, and some of the references in Sec. 2 are obtained through LLMs assisted search. At the same time, LLMs are used to help check grammar.

B Full Results

In Table 4 and Table 5, we present the full experimental results with standard deviations of various LLMs’ performance under perturbations from the MSCR algorithm. “Original” indicates the original accuracy of LLMs, “Attack” indicates the accuracy of LLMs after being attacked, and “Decrease” indicates the decrease in the accuracy of LLMs.

Table 4: Full performance of various LLMs on the GSM8K benchmark under the MSCR framework attack (accuracy, %).

Model	Original	Attack	Decrease
DeepSeek-R1-Distill-Qwen-32B	94.39 (± 0.25)	75.59 (± 0.22)	18.80 (± 0.01) $\downarrow$
DeepSeek-R1-Distill-Qwen-7B	93.71 (± 0.32)	68.16 (± 0.35)	25.55 (± 0.65) $\downarrow$
DeepSeek-R1-Distill-Qwen-1.5B	83.09 (± 0.40)	57.54 (± 0.06)	25.55 (± 0.63) $\downarrow$
Qwen2.5-Math-7B-Instruct	94.09 (± 0.17)	54.81 (± 0.67)	39.28 (± 0.82) $\downarrow$
Qwen2.5-Math-1.5B-Instruct	85.90 (± 0.25)	36.01 (± 0.62)	49.89 (± 0.59) $\downarrow$
Meta-Llama-3-70B-Instruct	93.86 (± 0.22)	65.43 (± 0.16)	28.43 (± 0.12) $\downarrow$
Meta-Llama-3-8B-Instruct	83.02 (± 0.77)	47.69 (± 0.19)	35.33 (± 0.90) $\downarrow$
gemma-3-27b-it	98.10 (± 0.13)	77.79 (± 0.17)	20.31 (± 0.04) $\downarrow$
gemma-3-12b-it	97.73 (± 0.32)	75.51 (± 0.25)	22.22 (± 0.11) $\downarrow$
gemma-3-4b-it	94.54 (± 0.32)	63.84 (± 0.31)	30.70 (± 0.04) $\downarrow$
gemma-2-27b-it	92.65 (± 0.28)	59.74 (± 0.25)	32.91 (± 0.19) $\downarrow$
gemma-2-9b-it	90.75 (± 0.10)	53.07 (± 0.35)	37.68 (± 0.42) $\downarrow$

Table 5: Full performance of various LLMs on the MATH500 benchmark under the MSCR algorithm attack (accuracy, %).

Model	Original	Attack	Decrease
DeepSeek-R1-Distill-Qwen-32B	93.00 (± 0.25)	79.20 (± 0.16)	13.80 (± 0.10) $\downarrow$
DeepSeek-R1-Distill-Qwen-7B	92.80 (± 0.25)	78.60 (± 0.19)	14.20 (± 0.10) $\downarrow$
DeepSeek-R1-Distill-Qwen-1.5B	81.20 (± 0.25)	63.20 (± 0.19)	18.00 (± 0.10) $\downarrow$
Qwen2.5-Math-7B-Instruct	83.20 (± 0.10)	61.60 (± 0.10)	21.60 (± 0.17) $\downarrow$
Qwen2.5-Math-1.5B-Instruct	76.20 (± 0.25)	49.80 (± 0.25)	26.40 (± 0.25) $\downarrow$
Meta-Llama-3-70B-Instruct	57.60 (± 0.25)	45.00 (± 0.25)	12.60 (± 0.33) $\downarrow$
Meta-Llama-3-8B-Instruct	30.60 (± 0.10)	6.60 (± 0.10)	24.00 (± 0.16) $\downarrow$
gemma-3-27b-it	92.00 (± 0.25)	75.00 (± 0.19)	17.00 (± 0.16) $\downarrow$
gemma-3-12b-it	89.80 (± 0.25)	70.40 (± 0.10)	19.40 (± 0.16) $\downarrow$
gemma-3-4b-it	81.40 (± 0.49)	56.20 (± 0.34)	25.20 (± 0.53) $\downarrow$
gemma-2-27b-it	68.30 (± 0.34)	39.00 (± 0.25)	29.30 (± 0.42) $\downarrow$
gemma-2-9b-it	47.00 (± 0.38)	11.60 (± 0.18)	35.40 (± 0.20) $\downarrow$

In Table 6 and Table 7, we present the full results of response length variation with standard deviations for various LLMs under perturbations from the MSCR algorithm. “Original” indicates the original response length of LLMs, “Attack” indicates the response length of LLMs after being attacked, and “Ratio” indicates how many times the response length of LLMs after being attacked is the original response length.

Table 6: Full response length variation of various LLMs on the GSM8K benchmark under the MSCR algorithm attack.

Model	Original	Attack	Decrease
DeepSeek-R1-Distill-Qwen-32B	727.0 (± 1.8)	1294.3 (± 1.1)	$1.78 \times$
DeepSeek-R1-Distill-Qwen-7B	1204.5 (± 10.1)	2190.2 (± 8.6)	$1.82 \times$
DeepSeek-R1-Distill-Qwen-1.5B	1754.3 (± 5.6)	2698.0 (± 52.6)	$1.54 \times$
Qwen2.5-Math-7B-Instruct	305.4 (± 3.5)	331.3 (± 15.5)	$1.08 \times$
Qwen2.5-Math-1.5B-Instruct	306.8 (± 10.0)	330.5 (± 35.1)	$1.08 \times$
Meta-Llama-3-70B-Instruct	170.4 (± 0.9)	207.2 (± 0.8)	$1.22 \times$
Meta-Llama-3-8B-Instruct	174.3 (± 0.6)	203.1 (± 3.9)	$1.17 \times$
gemma-3-27b-it	273.9 (± 2.1)	523.8 (± 3.3)	$1.91 \times$
gemma-3-12b-it	310.8 (± 8.8)	608.8 (± 9.7)	$1.96 \times$
gemma-3-4b-it	411.8 (± 0.2)	881.7 (± 3.9)	$2.14 \times$
gemma-2-27b-it	173.7 (± 0.6)	188.0 (± 1.0)	$1.08 \times$
gemma-2-9b-it	174.6 (± 1.0)	190.0 (± 0.7)	$1.09 \times$

Table 7: Full response length variation of various LLMs on the MATH500 benchmark under the MSCR algorithm attack.

Model	Original	Attack	Decrease
DeepSeek-R1-Distill-Qwen-32B	1668.9 (± 27.8)	1931.3 (± 17.5)	$1.16 \times$
DeepSeek-R1-Distill-Qwen-7B	2204.7 (± 41.4)	3081.9 (± 16.5)	$1.40 \times$
DeepSeek-R1-Distill-Qwen-1.5B	2521.6 (± 71.7)	3303.8 (± 118.7)	$1.31 \times$
Qwen2.5-Math-7B-Instruct	692.5 (± 6.3)	995.1 (± 11.1)	$1.44 \times$
Qwen2.5-Math-1.5B-Instruct	579.0 (± 4.8)	745.9 (± 7.3)	$1.29 \times$
Meta-Llama-3-70B-Instruct	283.5 (± 4.7)	304.9 (± 5.6)	$1.08 \times$
Meta-Llama-3-8B-Instruct	426.1 (± 4.5)	515.3 (± 6.6)	$1.21 \times$
gemma-3-27b-it	698.3 (± 3.3)	1033.8 (± 1.6)	$1.48 \times$
gemma-3-12b-it	755.4 (± 5.6)	1165.2 (± 4.9)	$1.54 \times$
gemma-3-4b-it	926.3 (± 17.0)	1308.3 (± 12.0)	$1.41 \times$
gemma-2-27b-it	359.1 (± 5.0)	378.2 (± 1.4)	$1.05 \times$
gemma-2-9b-it	396.8 (± 20.0)	413.8 (± 16.3)	$1.04 \times$

C More Results of Response Length Ratio Distribution

To better illustrate the phenomenon of response length collapse when LLMs face slight perturbations, we present the visualization results of the response length ratio distribution for other models, as shown in Figure 5 and Figure 6.




Figure 5: Other LLMs on the GSM8K benchmark showing the distribution of the ratio between response length for perturbed questions and the original response length.




Figure 6: Other LLMs on the MATH500 benchmark showing the distribution of the ratio between response length for perturbed questions and the original response length.

D More Perturbation Examples

As shown below, we present more examples of the performance of LLMs under the MSCR framework when facing perturbations. These examples further demonstrate the vulnerability of LLMs' mathematical reasoning ability when facing slight perturbations.

Perturbation Example of gemma-2-27b-it on GSM8K

Original question: Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?

Original answer: 70000

Perturbation question: **he** decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?

Model Response: Here's how to solve this problem step-by-step:

1. Step 1: Calculate the increase in value. The house's value increased by 150% of its initial value, which is $0.150 * \$80,000 = \$12,000$;
2. Step 2: Calculate the new value of the house. Add the increase in value to the initial value: $\$80,000 + \$12,000 = \$92,000$;
3. Step 3: Calculate the total cost. Add the purchase price and the repair costs: $\$80,000 + \$50,000 = \$130,000$;
4. Step 4: Calculate the profit. Subtract the total cost from the new value of the house: $\$92,000 - \$130,000 = -\$38,000$;
5. Answer: He made a loss of \$38,000.

Is the answer correct? No.

Perturbation Example of DeepSeek-R1-Distill-Qwen-32B on GSM8K

Original question: Kylar went to the store to buy glasses for his new apartment. One glass costs \$5, but every second glass costs only 60% of the price. Kylar wants to buy 16 glasses. How much does he need to pay for them?

Original answer: 64

Perturbation question: Kylar went **_to** the store to buy glasses for his new apartment. One glass costs \$5, but every second glass costs only 60% of the price. Kylar wants to buy 16 glasses. How much does he need to pay for them?

Model Response: To determine how much Kylar needs to pay for 16 glasses, we can break down the cost based on the given pricing structure:

1. Cost of the first glass: \$5.00;
2. Cost of each subsequent glass: 60% of \$5.00 = \$3.00;
3. Number of glasses beyond the first one: $16 - 1 = 15$;
4. Total cost for the additional 15 glasses: $15 \times \$3.00 = \45.00 ;
5. Total cost: $\$5.00 + \$45.00 = \$50.00$;
6. Therefore, Kylar needs to pay 50 dollars for the 16 glasses.

Is the answer correct? No.

Perturbation Example of Qwen2.5-Math-7B-Instruct on GSM8K

Original question: Carla is downloading a 200 GB file. Normally she can download 2 GB/minute, but 40 of the way through the download, Windows forces a restart to install updates, which takes 20 minutes. Then Carla has to restart the download from the beginning. How long does it take to download the file?

Original answer: 160

Perturbation question: **carla** is downloading a 200 GB file. Normally she can download 2 GB/minute, but 40 of the way through the download, Windows forces a restart to install updates, which takes 20 minutes. Then Carla has to restart the download from the beginning. How long does it take to download the file?

Model Response: To determine the total time it takes Carla to download the 200 GB file, we need to break down the process into segments and calculate the time for each segment:

1. Calculate the time to download 40 of the file: The file size is 200 GB. 40 of 200 GB is 0.40 times 200 = 80 GB. Carla downloads at a rate of 2 GB per minute. Time to download 80 GB is $\frac{80}{2} = 40$ minutes;
2. Add the time for the restart: The restart takes 20 minutes;
3. Calculate the time to download the remaining 60 of the file: The remaining file size is 200 - 80 = 120 GB. Carla downloads at a rate of 2 GB per minute. Time to download 120 GB is $\frac{120}{2} = 60$ minutes;
4. Add the time to download the file from the beginning: The file size is 200 GB. Carla downloads at a rate of 2 GB per minute. Time to download 200 GB is $\frac{200}{2} = 100$ minutes;
5. Sum up all the times: Time to download 40 of the file 40 minutes. Time for the restart 20 minutes. Time to download the remaining 60 of the file 60 minutes. Time to download the file from the beginning 100 minutes. Total time = 40 + 20 + 60 + 100 = 220 minutes;
6. Therefore, the total time it takes Carla to download the file is 220 minutes.

Is the answer correct? No.

Perturbation Example of gemma-3-27b-it on GSM8K

Original question: Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?

Original answer: 70000

Perturbation question: Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house **to** 150%. How much profit did he make?

Model Response: Let's break it down step by step:

1. Calculate the total cost: Josh spent \$80,000 + \$50,000 = \$130,000;
2. Calculate the increased value: The value increased to 150% of the original purchase price, so the new value is \$80,000 * 1.50 = \$120,000. However, the problem states the repairs *increased* the value to 150%. This means the final value is 150% of the original cost *plus* repairs. So, the base for the 150% calculation is \$80,000 + \$50,000 = \$130,000. The new value is \$130,000 * 1.50 = \$195,000;

3. Calculate the profit: Profit is the new value minus the total cost: $\$195,000 - \$130,000 = \$65,000$;

4. Final Answer: The final answer is 65000;

Is the answer correct? No.

Perturbation Example of gemma-2-27b-it on MATH500

Original question: What is the smallest positive perfect cube that can be written as the sum of three consecutive integers?

Original answer: 27

Perturbation question: What **Is** the smallest positive perfect cube that can be written as the sum of three consecutive integers?

Model Response: Here's how to solve this problem step-by-step:

1. Represent the consecutive integers: Let the three consecutive integers be n , $n + 1$, and $n + 2$;
2. Set up the equation: The sum of these integers is $n + (n + 1) + (n + 2) = 3n + 3$;
3. Express the sum as a perfect cube: We want to find the smallest n such that $3n + 3$ is a perfect cube;
4. Test values: Let's start testing small values of n : If $n = 1$, then $3n + 3 = 6$ (not a perfect cube) ; If $n = 2$, then $3n + 3 = 9$ (3 cubed);
5. We found that when $n = 2$, the sum of the consecutive integers is 9, which is 3 cubed.
6. Answer: 9.

Is the answer correct? No.

Perturbation Example of gemma-3-27b-it on MATH500

Original question: Below is the graph of $y = a \sin(bx + c) + d$ for some positive constants a , b , c , and d . Find the smallest possible value of c .

```
import TrigMacros;
```

```
size(400);
```

```
real f(real x)
{
    return 2*sin(3*x + pi) + 1;
}
```

```
draw(graph(f,-3*pi,3*pi,n=700,join=operator ..),red);
trig_axes(-3*pi,3*pi,-4,4,pi/2,1);
layer();
rm_trig_labels(-5,5, 2);
```

```
label("$1$", (0,1), E);
label("$2$", (0,2), E);
label("$3$", (0,3), E);
```

```
label("$-1$", (0,-1), E);
label("$-2$", (0,-2), E);
label("$-3$", (0,-3), E);
```

Original answer: π

Perturbation question: **under** is the graph of $y = a \sin(bx + c) + d$ for some positive constants a , b , c , and d . Find the smallest possible value of c .

```
import TrigMacros;

size(400);

real f(real x)
{
    return 2*sin(3*x + pi) + 1;
}

draw(graph(f,-3*pi,3*pi,n=700,join=operator ..),red);
trig_axes(-3*pi,3*pi,-4,4,pi/2,1);
layer();
rm_trig_labels(-5,5, 2);

label("$1$", (0,1), E);
label("$2$", (0,2), E);
label("$3$", (0,3), E);
label("$-1$", (0,-1), E);
label("$-2$", (0,-2), E);
label("$-3$", (0,-3), E);
```

Model Response: Here's how to solve this problem step-by-step:

1. Let the given function be $y = a \sin(bx + c) + d$.
2. From the graph, we can see that the maximum value of y is 3 and the minimum value is -1 .
3. Thus, $a + d = 3$ and $-a + d = -1$. Adding these equations gives $2d = 2$, so $d = 1$. Then $a + 1 = 3$, so $a = 2$.
4. The period of the function is $\frac{2\pi}{b}$. From the graph, we can see that the function completes one period from $-\frac{\pi}{3}$ to $\frac{\pi}{3}$, so the period is $\frac{\pi}{3} - (-\frac{\pi}{3}) = \frac{2\pi}{3}$. Thus, $\frac{2\pi}{b} = \frac{2\pi}{3}$, so $b = 3$.
5. Now we have $y = 2 \sin(3x + c) + 1$.
6. The graph is $y = 2 \sin(3x + c) + 1$.
7. When $x = 0$, we have $y = 2 \sin c + 1$. From the graph, when $x = 0$, $y = 2 \sin(3(0) + c) + 1 = 2 \sin c + 1$.
8. From the graph, we see that when $x = 0$, $y = 2 \sin(c) + 1 = 0$.
9. Then $2 \sin c = -1$, so $\sin c = -\frac{1}{2}$.
10. The general solution for $\sin c = -\frac{1}{2}$ is $c = \frac{7\pi}{6} + 2n\pi$ or $c = \frac{11\pi}{6} + 2n\pi$ for some integer n .
11. We want to find the smallest positive value of c .
12. If $c = \frac{7\pi}{6} + 2n\pi$, then for $n = -1$, $c = \frac{7\pi}{6} - 2\pi = \frac{7\pi-12\pi}{6} = -\frac{5\pi}{6}$.
13. For $n = 0$, $c = \frac{7\pi}{6}$.
14. If $c = \frac{11\pi}{6} + 2n\pi$, then for $n = -1$, $c = \frac{11\pi}{6} - 2\pi = \frac{11\pi-12\pi}{6} = -\frac{\pi}{6}$.

15. For $n = 0$, $c = \frac{11\pi}{6}$.
16. The smallest positive value of c is $\frac{7\pi}{6}$. However, we can write $y = 2\sin(3x + c) + 1 = 2\sin(3x + c + 2\pi) + 1$.
17. We can also write $y = 2\sin(3x + c) + 1 = 2\sin(3x + c - 2\pi) + 1$.
18. We can also write $y = 2\sin(3x + c) + 1 = 2\sin(3x + c + k(2\pi)) + 1$ for any integer k .
19. We have $y = 2\sin(3x + c) + 1$. We want the graph to be $y = 2\sin(3x + \pi) + 1$.
20. We have $3x + c = 3x + \pi$, so $c = \pi$.
21. We have $\sin c = -\frac{1}{2}$, so $c = \frac{7\pi}{6}$ or $c = \frac{11\pi}{6}$.
22. We can write $y = 2\sin(3x + \frac{7\pi}{6}) + 1$.
23. We can also write $y = 2\sin(3x + \frac{11\pi}{6}) + 1$.
24. We can write $y = 2\sin(3x + \frac{7\pi}{6} - 2\pi) + 1 = 2\sin(3x - \frac{5\pi}{6}) + 1$.
25. We can write $y = 2\sin(3x + \frac{11\pi}{6} - 2\pi) + 1 = 2\sin(3x - \frac{\pi}{6}) + 1$.
26. The graph is $y = 2\sin(3x + \pi) + 1$. Thus, we want $c = \pi$.
27. We have $\sin c = -\frac{1}{2}$.
28. If $c = \pi$, then $\sin c = \sin \pi = 0$.
29. If $c = \frac{7\pi}{6}$, then $\sin c = \sin \frac{7\pi}{6} = -\frac{1}{2}$.
30. If $c = \frac{11\pi}{6}$, then $\sin c = \sin \frac{11\pi}{6} = -\frac{1}{2}$.
31. We want to find the smallest positive value of c .
32. We have $y = 2\sin(3x + c) + 1$.
33. At $x = 0$, we have $y = 2\sin c + 1 = 0$, so $\sin c = -\frac{1}{2}$.
34. Then $c = \frac{7\pi}{6} + 2n\pi$ or $c = \frac{11\pi}{6} + 2n\pi$.
35. We want the smallest positive value of c .
36. If $n = -1$, $c = \frac{7\pi}{6} - 2\pi = -\frac{5\pi}{6}$, which is negative.
37. If $n = -1$, $c = \frac{11\pi}{6} - 2\pi = -\frac{\pi}{6}$, which is negative.
38. If $n = 0$, $c = \frac{7\pi}{6}$ or $c = \frac{11\pi}{6}$. The smallest is $\frac{7\pi}{6}$.
39. However, we can also write $2\sin(3x + c) = 2\sin(3x + c + 2\pi)$.
40. Also, $2\sin(3x + c) = 2\sin(3x + c - 2\pi)$.
41. The graph is $y = 2\sin(3x + \pi) + 1$, so $c = \pi$.
42. We have $\sin c = -\frac{1}{2}$, so $c = \frac{7\pi}{6} + 2n\pi$ or $c = \frac{11\pi}{6} + 2n\pi$.
43. If we take $c = \frac{7\pi}{6}$, then $y = 2\sin(3x + \frac{7\pi}{6}) + 1$.
44. If we take $c = \frac{11\pi}{6}$, then $y = 2\sin(3x + \frac{11\pi}{6}) + 1$.
45. If we want the graph to be $y = 2\sin(3x + \pi) + 1$, then $c = \pi$. But $\sin \pi = 0$, which is not $-\frac{1}{2}$. Let $3x + c = 3x + \pi$. Then $c = \pi$.
46. Since $c = \frac{7\pi}{6}$ and $c = \frac{11\pi}{6}$ are solutions to $\sin c = -\frac{1}{2}$, we can subtract 2π from these to get negative values.
47. We want the smallest positive value of c , which is $\frac{7\pi}{6}$.
48. Final Answer: The final answer is $\boxed{\frac{7\pi}{6}}$

Is the answer correct? No.

Perturbation Example of Meta-Llama-3-70B-Instruct on MATH500

Original question: If $f(x) = \frac{3x-2}{x-2}$, what is the value of $f(-2) + f(-1) + f(0)$? Express your answer as a common fraction.

Original answer: $\frac{14}{3}$

Perturbation question: **.If** $f(x) = \frac{3x-2}{x-2}$, what is the value of $f(-2) + f(-1) + f(0)$? Express your answer as a common fraction.

Model Response: Here's how to solve this problem step-by-step:

1. We have that $f(-2)$ is undefined, $f(-1) = \frac{3(-1)-2}{-1-2} = \frac{-5}{-3} = \frac{5}{3}$, and $f(0) = \frac{3(0)-2}{0-2} = \frac{-2}{-2} = 1$. Therefore, $f(-2) + f(-1) + f(0)$ is undefined.

Is the answer correct? No.

Perturbation Example of DeepSeek-R1-Distill-Qwen-32B on MATH500

Original question: The proper divisors of 12 are 1, 2, 3, 4 and 6. A proper divisor of an integer N is a positive divisor of N that is less than N . What is the sum of the proper divisors of the sum of the proper divisors of 284?

Original answer: 284

Perturbation question: **the** proper divisors of 12 are 1, 2, 3, 4 and 6. A proper divisor of an integer N is a positive divisor of N that is less than N . What is the sum of the proper divisors of the sum of the proper divisors of 284?

Model Response: To find the sum of the proper divisors of the sum of the proper divisors of 284, we proceed as follows:

1. Find the proper divisors of 284: 1, 2, 4, 71, 142;
2. Calculate the sum of these proper divisors: $1 + 2 + 4 + 71 + 142 = 220$;
3. Find the proper divisors of 220: 1, 2, 4, 5, 10, 11, 20, 22, 44, 55, 110;
4. Calculate the sum of these proper divisors: $1 + 2 + 4 + 5 + 10 + 11 + 20 + 22 + 44 + 55 + 110 = 220$;
5. Final Answer: 220.

Is the answer correct? No.

Perturbation Example of Qwen2.5-Math-7B-Instruct on MATH500

Original question: Find the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$. (Do not use a calculator!)

Original answer: 13535

Perturbation question: Find the greatest integer less **Ethan** $(\sqrt{7} + \sqrt{5})^6$. (Do not use a calculator!)

Model Response: To find the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$, we start by letting $x = \sqrt{7} + \sqrt{5}$. We also consider its conjugate $y = \sqrt{7} - \sqrt{5}$. Note that $0 < y < 1$ because $\sqrt{7} > \sqrt{5}$:

1. First, we calculate x^2 and y^2 :

$$x^2 = (\sqrt{7} + \sqrt{5})^2 = 7 + 2\sqrt{35} + 5 = 12 + 2\sqrt{35}$$

$$y^2 = (\sqrt{7} - \sqrt{5})^2 = 7 - 2\sqrt{35} + 5 = 12 - 2\sqrt{35}$$

2. Next, we find x^4 and y^4 :

$$x^4 = (x^2)^2 = (12 + 2\sqrt{35})^2$$

$$= 144 + 2 \cdot 12 \cdot 2\sqrt{35} + (2\sqrt{35})^2$$

$$= 144 + 48\sqrt{35} + 4 \cdot 35$$

$$= 144 + 48\sqrt{35} + 140$$

$$= 284 + 48\sqrt{35}$$

$$y^4 = (y^2)^2 = (12 - 2\sqrt{35})^2$$

$$= 144 - 2 \cdot 12 \cdot 2\sqrt{35} + (2\sqrt{35})^2$$

$$= 144 - 48\sqrt{35} + 4 \cdot 35$$

$$= 144 - 48\sqrt{35} + 140$$

$$= 284 - 48\sqrt{35}$$

3. Now, we calculate x^6 and y^6 :

$$x^6 = x^2 \cdot x^4 = (12 + 2\sqrt{35})(284 + 48\sqrt{35})$$

$$= 12 \cdot 284 + 12 \cdot 48\sqrt{35} + 2\sqrt{35} \cdot 284 + 2\sqrt{35} \cdot 48\sqrt{35}$$

$$= 3408 + 576\sqrt{35} + 568\sqrt{35} + 2 \cdot 48 \cdot 35$$

$$= 3408 + 1144\sqrt{35} + 3360 = 6768 + 1144\sqrt{35}$$

$$y^6 = y^2 \cdot y^4 = (12 - 2\sqrt{35})(284 - 48\sqrt{35})$$

$$= 12 \cdot 284 - 12 \cdot 48\sqrt{35} - 2\sqrt{35} \cdot 284 + 2\sqrt{35} \cdot 48\sqrt{35}$$

$$= 3408 - 576\sqrt{35} - 568\sqrt{35} + 2 \cdot 48 \cdot 35$$

$$= 3408 - 1144\sqrt{35} + 3360 = 6768 - 1144\sqrt{35}$$

4. Since $0 < y < 1$, it follows that $0 < y^6 < 1$. Therefore, $6768 - 1144\sqrt{35}$ is a number between 0 and 1 less than 6768. This means that $6768 + 1144\sqrt{35}$ is a number slightly less than 6769. Thus, the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$ is:

6768

Is the answer correct? No.