

Evaluating LLMs’ Mathematical and Coding Competency through Ontology-guided Interventions

Pengfei Hong², Navonil Majumder¹, Deepanway Ghosal¹ *
Somak Aditya³, Rada Mihalcea⁴, Soujanya Poria⁵

¹ Singapore University of Technology and Design, ²HiThink Research,
³ IIT Kharagpur, ⁴ University of Michigan, ⁵Nanyang Technological Research

Abstract

Recent advancements in Large Language Models (LLMs) have showcased striking results on existing logical reasoning benchmarks, with some models even surpassing human performance. However, the true depth of their competencies and robustness in reasoning tasks remains an open question. To this end, in this paper, we focus on two popular reasoning tasks: arithmetic reasoning and code generation. Particularly, we introduce (i) a general ontology of perturbations for math and coding questions, (ii) a semi-automatic method to apply these perturbations, and (iii) two datasets, GSMORE and HUMANEVAL-CORE, respectively, of perturbed math and coding problems to probe LLM capabilities in numeric reasoning and coding tasks. Through comprehensive evaluations of both closed-source and open-source LLMs, we show a significant performance drop across all the models against the perturbed questions, suggesting that the current LLMs lack robust problem solving skills and structured reasoning abilities in many areas, as defined by our ontology. We open-source the datasets and source code at: <https://github.com/declare-lab/LLM-ReasoningTest>.

1 Introduction

Logical reasoning, which includes mathematics and programming, in a structured and well-defined domain, becomes increasingly complex with the increasing presence of diverse situations, events, and contexts represented in natural language. However, the process of solving a mathematical expression should remain consistent and a logical reasoner should be robust to *certain* changes in the context. Consider, problems I) *I have two apples and Sam has three more. How many apples are there?*, and II) *I see two wombats. And my friend sees another three. How many wombats are there?*. Variations I and II should yield the

same expression ($2 + 3$), and a similar process to solve resulting in the same answer (5). Current state-of-the-art Large Language Models (LLM) have shown impressive performance on a *wide* variety of mathematical problems (Cobbe et al., 2021a) and coding problems (Chen et al., 2021a) expressed in natural language. However, these evaluations barely test the depth of LLMs’ expertise, and thus we do not currently have fine insights into the LLM capabilities in these domains. For example, in mathematics, GPT-4’s performance monotonically decreases from GSM-8k (Cobbe et al., 2021a) (92%; 5-shot CoT) on grade school mathematical problems requiring rigorous arithmetic and logical reasoning to solve; to MMLU-Math (87.5%) (Hendrycks et al., 2020) on a collection of mathematical problems, ranging in difficulty from elementary to advanced levels; and to MATH (50.36%) (Hendrycks et al., 2021) on challenging competitive mathematics problems. Similar variance in LLM performance can also be observed for coding challenges (Chen et al., 2021a). Such high-level evaluations show the *breadth* of LLM capabilities but do not provide any insights into the finer and more fundamental math- and coding-relevant capabilities and shortcomings of LLMs as (i) many LLMs like GPT-4 (OpenAI, 2023) are exposed to publicly available math and coding datasets during pre-training, and (ii) many datasets focus on advanced branches of mathematics and problems, skipping the fundamentals. Hence, we instead focus on depth, encapsulated by one fundamental question:

Research Question: How robust are the capabilities of large language models (LLMs) in reasoning and understanding the problem-solving process?

Our Approach: Introducing small modifications to the selected seed math problems from GSM8K and HumanEval, where leading LLMs,

*Now at Google Deepmind.

such as o1-preview, have shown excellent performance. Our goal is to examine whether these models can effectively handle slight variations to the seed questions, which they previously solved successfully. An example is presented in Figure 2. If the models struggle with these modified problems, it would suggest that there is still room for improvement in LLMs’ core reasoning abilities.

Evaluating LLMs on these modified questions could provide fine-grained insights into the robustness of the reasoning abilities in the context of math and coding. Specifically, our evaluation benchmark is based on our novel ontology of perturbation operations that lists various changes across a diverse set of factors. We apply these perturbations to existing arithmetic and coding problems. These perturbations allow us to assess whether the model truly comprehends underlying concepts and structures or regurgitates. For instance, while a model may correctly answer a math question in a dataset like GSM8k, it may struggle with a simple perturbation to the question, such as replacing numerical values (Srivastava et al., 2024) in math questions with variables, which challenges the model to establish relationships among the variables, revealing its deeper understanding (or lack thereof). By introducing these ontological perturbations, (1) we gain insights into the models’ reasoning abilities and (2) uncover strategies for future data augmentation that can then be utilized to enhance LLMs through weakly supervised fine-tuning methodologies.

Our perturbations extend beyond mere numerical alterations to seed questions; they also challenge with the underlying concepts and principles required to solve a problem. For instance, we introduce a category *Critical Thinking*, which includes scenarios where multiple correct answers may be possible, depending on the constraints on various variables, as demonstrated in the example below:

Original Question: A merchant must choose between two purchasing plans: jewelry valued at \$5,000 or electronic gadgets valued at \$8,000. His financial advisor predicts a **2.5%** increase in the jewelry market and a **1.2%** increase in the electronic gadgets market over the next month. If the merchant aims to maximize profit by the end of this month, what would be his profit?

Perturbed Question: A merchant must choose between two purchasing plans: jewelry valued at \$5,000 or electronic gadgets valued at \$8,000. His financial advisor predicts that the jewelry market

will rise by **$x\%$** and the electronic gadgets market by **1.2%** over the next month. If the merchant aims to maximize profit by the end of this month, what would his profit be? Such instances are unique to our datasets making the evaluation more rigorous.

We sample five questions from GSM8K, subjecting them to 44 types of perturbations, followed by extensive manual filtering and validation, resulting in 216 perturbed questions in GSMORE. As our ontology is generic, we also apply these perturbations to randomly sampled five code-generation questions from HumanEval producing 219 questions in HUMANEVAL-CORE.

This work makes the following contributions:

1. We propose a manually-curated, novel, extensive, and extensible ontology of perturbation operations for math and coding problems.
2. We provide a strategy to semi-automatically apply such perturbations first through GPT-4, followed by manual filtering. We refer to this dataset as GSMORE— Mathematics-oriented Robustness Evaluation, consisting of 216 examples starting from just five questions in GSM8K.
3. We show that the ontology and the semi-automatic pipeline can be easily extended to domains like code-generation, by creating the dataset HUMANEVAL-CORE — Coding-oriented Robustness Evaluation, having 219 samples from five seed questions in HumanEval.
4. We further evaluate the math and code-based reasoning abilities of the latest LLMs, including GPT-4o and o1, using our datasets, revealing their vulnerability to these perturbations.

The integration of our ontology with the newly introduced dataset, referred to as GSMORE and HUMANEVAL-CORE, paves the way for a fresh perspective in assessing the math and code reasoning capabilities of LLMs. This evaluation approach transcends the often unidimensional leaderboard performance metrics by delving into LLM performance on finely defined classes of problems.

2 The Ontology of Perturbations

2.1 Need for Ontology of Perturbations

We first identify a set of factors relevant to the solution of a structured reasoning problem expressed in natural language (similar to Kaushik et al. (2021)); and perturb a seed question under these set of factors semi-automatically in a model-agnostic way (i.e., not necessarily adversarial to a target model).

In the NLI context, [Kaushik et al. \(2021\)](#) utilized human workers to directly perturb a hypothesis, keeping the premise constant; and in a post-hoc way, identifies the categories (or factors) which such revisions pertain to. Previous works ([Xu et al., 2023](#); [Li et al., 2024a](#); [Wang et al., 2024](#)) discuss ways of perturbation, by identifying a set of factors that is specifically designed to increase the complexity of a seed questions in limited ways. The categories are broad and do not exploit the *logical* nature of the underlying domain (along with the *linguistic* dimensions of the instruction). This is where, we believed, an ontological approach may help, where broader categories can help us generalize, while fine-grained sub-categories exploit the domain-specific characteristics.

Let’s take mathematics for example. The solution to a reasoning problem can depend on the number and complexity of operations, variables, functions, and possible existing theorems (external knowledge). Similarly, code generation problems can depend on the data structures, variables, functions, and libraries it needs access to. On top of this well-defined set of factors existing in structured reasoning problems, the list of factors expands as the problem is expressed in natural language. Entities and relations expressed in the text need to be mapped to variables and constants (in both). Physical actions (giving and taking apples) may need to be mapped to mathematical operations (or code). It is clear that the set of *logical* and *linguistic* factors co-exist in these reasoning problems, detailed in Appendix I.1. Therefore we come up with an extensible ontology, capturing the above nuances. We believe it will capture and categorize the factors where LLMs fail over multiple domains. As others have shown, the same process can be enabled to perform data augmentations.

2.2 The Ontology

Extending SVAMP ([Patel et al., 2021](#))-like perturbations, we propose a set of high-level categories that are applicable to a broad class of reasoning tasks, expressed in natural language. We primarily identified the following hierarchy (see Table 7):

Level I: Aspect. There are two aspects to these perturbations: (i) *structural perturbation* and (ii) *representational perturbation*. *Structural perturbation* covers all perturbations that probe the underlying reasoning path (or structure) in different ways, by slightly varying the logic behind the question or probing intermediate steps, seeking explanations.

Representational perturbations involves modification of the encoding of the question or solution while preserving the underlying logic of the original question.

Level II: Target. The subject of change in each *aspect* is gradually refined into multiple *Targets*. For example, the target of *logic*, under *structural perturbations*, deals with perturbations that alter the reasoning path in different controlled ways.

Level III: Dimension. This is a further refinement that defines the exact target dimensions (the WHAT) in the reasoning process (question, reasoning, computation, answer expression etc.) to which the perturbations are applied.

Level IV: Category. This level captures the method (the HOW) through which the higher-level *Dimension* perturbation is achieved. These methods are domain dependent and, thus, their implementations vary from math to coding problems.

3 Curation of GSMORE and HUMANEVAL-CORE

Our objective is to assess the resilience of LLMs to perturbations of math and coding questions along various dimensions. Thus, as seed datasets, we use GSM8K ([Cobbe et al., 2021b](#))—a collection of mathematical problems demanding rigorous arithmetic and logical reasoning—and HumanEval ([Chen et al., 2021a](#)) for coding. Five questions¹ from GSM8K are perturbed using our ontological framework (see Appendix I) to generate GSMORE. On the other hand, we sampled five coding problems from HumanEval dataset ([Chen et al., 2021a](#)) that were perturbed using the ontology explained in Appendix I. These perturbations are aimed at modifying the problems in terms of complexity and representation to assess the robustness of the LLMs to these ontological categories of perturbations. Fig. 1 shows examples of three perturbed questions and answers from GSMORE and HUMANEVAL-CORE. Examples and definitions of all the remaining perturbations are present in Appendix I. We use a three-staged combination of automatic generation from GPT-4 ([OpenAI, 2023](#)) with human verification and annotation to create GSMORE and HUMANEVAL-CORE: (i) perturbed question generation (§3.1), (ii) filtering and

¹Math questions in the GSM8K dataset take between two and eight steps to solve. We randomly chose five questions that take three to seven steps to solve. We cover various topics involving algebraic questions, physical application questions, and decision-based application questions.

validation of generated questions (§3.2), and (iii) annotating final answers (§3.3).

3.1 Perturbed Question Generation

In the first stage, our objective is to create perturbed questions from the source GSM8K/HumanEval questions for each perturbation type. We write prompt templates for each perturbation type and fill them with a source question to create the input prompt to GPT-4. Each template captures the essence of the respective perturbation type (Appendix I.2, Appendix I.3, Appendix I.4) to instruct GPT-4 on how to perturb the source question.

For example, the prompt for *Remove Constraint (G1.)* for our running example is as follows:

```
Instruction: Rewrite the original mathematical context below based on the #Rewrite Requirement#.
Your output should only be #Rewritten Context#.
#Original Context#: John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick.
#Original Query#: What is the total inner volume of all 3 boxes?
#Rewrite Requirement#: 1. Remove some constraints or information from the original context.
2. Make sure the rewritten question can still be solved, but the answer is simpler.
#Rewritten Context#:
```

This prompt to GPT-4 generated: *John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. What is the total volume of all 3 boxes?*. The black text in this prompt marks the static template components to enforce the intended perturbation, while the blue text indicates the source question. These templates are used iteratively to generate perturbed questions for GPT-4.

3.2 Filtering and Validation

Unfortunately, GPT-4-generated perturbed questions sometimes lack meaning and suitability for robustness testing due to complex and open-ended perturbation types, leading to errors in generation. As noted in Li et al. (2024a), GPT-4 may i) fail to incorporate perturbations, such as missing values in *Data Restructuring*, ii) introduce unintended changes. We aim to maintain Human Understandability, Logical Coherence, and Instruction Adherence, as detailed in Appendix D.2.

To ensure these qualities and relevance, we use a semi-automatic filtering process. Initially, GPT-

4 performs an automated check against the three criteria, discarding any questions that do not meet them. Questions that repeatedly fail are then reviewed and handled by a human annotator.

Human Verification. Despite automatic verification, perturbed questions still have limitations, so we conduct a final human verification to refine them. Our findings show that 36% of the filtered questions needed minor rewording, 31% contained significant inaccuracies or failed the filtering, and 33% were correct as is. Thus, the final questions in GSMORE and HUMAN-EVAL-CORE are high-quality, understandable, logically coherent, and aligned with the intended perturbation method. Human verification is performed by five PhD computer science students, with each question revised by two annotators and verified by two others.

3.3 Obtaining Final Answers of the Perturbed Questions

Finally, we also annotate the gold answer for the perturbed questions. We engaged the same five annotators for this process. Each gold answer was initially annotated by one annotator. Subsequently, the annotated responses underwent verification by the other two annotators.

3.4 Statistics of GSMORE and HUMAN-EVAL-CORE

We sampled five questions from GSM8K and HumanEval and perturbed them using GPT-4 in 44 distinct perturbation categories. Following a rigorous process of filtering and validation, we retained a total of 216 and 219 perturbed questions in GSMORE and HUMAN-EVAL-CORE, respectively. We specify the detailed statistics in Appendix C and the details of the five selected questions from each dataset in Appendix F and Appendix G respectively.

4 Experiments

4.1 Evaluation Protocol

Owing to the loosely controlled format of the LLM responses to the majority of the questions, calculating accuracy through direct string matching with the annotated answer may not always be reliable. Additionally, in the context of *concept analysis*, curating an exhaustive list of correct answers could be intractable. For instance, the category *optimize solution (G27.)* asks to further optimize the provided solution. There could be numerous distinct

Original	Logic Alteration	Concept Analysis	Format Constraint
GSM8K		GSMoer	
<p>Question: John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the inner volume of all 3 boxes?</p> <p>Answer: Walls are 1 inch thick, reducing each dimension by 2 inches. Thus, the internal dimensions become 3x4x2=24 cubic inches, making the total volume for all 3 boxes 3x24=72 cubic inches.</p>	<p>Question - Variable Relationship: John has X boxes. Each box is Y inches by 6 inches by 4 inches. ... If the total inner volume of all the boxes is 72 cubic inches, find the equation that relates X and Y?</p> <p>Answer: Walls are 1 inch thick, reducing each dimension by 2 inches. Thus, the internal dimensions become (Y-2)x4x2 cubic inches for one box. Given the total volume for all X boxes is 72 cubic inches, the equation relating X and Y simplifies to $X*(Y-2)=9$.</p>	<p>Question - Step Necessity: ... Suppose we want to find out the total inner volume of all 3 boxes. To solve this math question, is there a way to determine the total inner volume of all 3 boxes without calculating the inner volume of one box?</p> <p>One possible answer: Calculate by using: total inner volume = total outer volume - volume of the wall. Yes.</p>	<p>Question - Reasoning Format: Answer the following question with only base-2 coded values. ... What is the total inner volume of all 3 boxes?</p> <p>Answer (in binary): Walls are 1(bin) inch thick, reducing each dimension by 10(bin) inches. Thus, the internal dimensions become: 11x100x10=11000(bin) cubic inches. making the total volume for all 11(bin) boxes: 11x11000=1001000 (bin) cubic inches.</p>
HumanEval		HumanEval-Core	
<p>Question: <pre>def gcd(a:int, b:int): """Return the greatest common divisor of two integers a and b >>>gcd(25, 15) 5"""</pre></p> <p>Answer: <pre>while b: a, b = b, a % b return abs(a)</pre></p>	<p>Question - Variable Relationship: If the following program output integer 7, what can we find about the relationship between its input parameters: a, b? <pre>def function(a: int, b: int): while b: a, b = b, a % b return abs(a)</pre></p> <p>Answer: Both a and b share 7 as their greatest common divisor.</p>	<p>Question - Step Necessity: Justify the necessity and functionality of operation 'abs(a)' in fulfilling the requirement. {##Original Question##} {##Original Answer##}</p> <p>One possible answer: Ensures GCD result is non-negative, aligning with mathematical expectations.</p>	<p>Question - Reasoning Format: Fulfill the coding requirement below without using python keyword "while" inside the answer. {##Original Question##} {##Original Answer##}</p> <p>Answer: <pre>if b == 0: return abs(a) else: return gcd(b, a % b)</pre></p>

Figure 1: Examples of the original questions and perturbed questions with *Logic*, *Concept* and *Format* as Targets. The targeted change for each question is highlighted in yellow background.

valid ways to optimize the given solution. To address these challenges, we manually evaluated the generated outputs. For HUMAN-EVAL-CORE, we designed test cases where possible. The evaluation is fully automated for those instances. For the remaining instances without test cases, we manually verified the answers returned by the LLMs.

Automated Evaluation. With the recent advancements in the reasoning abilities of large language models (LLMs), we also utilize GPT-4o as an evaluator to assess the correctness of answers against the manually annotated ground truth. GPT-4o proved to be an effective evaluator, showing a high level of agreement with human judgments across various models. We report the automated evaluation results and their agreement with manual evaluation in Table 2.

4.2 Experimental Setup

We evaluated five prominent closed- and open-sourced LLMs on our benchmark. The closed-sourced LLMs are o1-preview, GPT-4o, GPT-4, GPT-3.5, and Gemini-1.5-pro. The remaining open-sourced LLMs include one general-purpose LLM and one LLM finetuned on task-specific datasets. The general-purpose LLM is Llama3-8B-Instruct and task-specific LLMs are MetaMath-70B-V1.0 and CodeLlama-70B-Instruct for cod-

ing and math, respectively. MetaMath-70B-V1.0 is finetuned on a mixture of datasets from Metamath (Yu et al., 2023b) and Mistral (Jiang et al., 2023) and CodeLlama-70B-Instruct is finetuned on publicly available coding and coding-related instructions (Rozière et al., 2023). Model Details are specified in Appendix D.1. We listed the prompts used for these models in Appendix J. Each question is evaluated with pass@1 metric under zero-shot setting. More details are in Appendix J on the evaluation settings.

4.3 Experimental Results and Analyses

Large Drop in Reasoning Performance of LLMs.

The results in Table 1 show that perturbed questions significantly challenge all models in both math and coding contexts. GPT-4’s accuracy decreased notably, as did other LLMs, with all showing a performance decline of over 30 points. Notably, closed-source models outperformed open-source ones in every tested aspect. Models like CodeLlama and Metamath fine-tuned on specific tasks, performed better in logic alteration and representational perturbations but worse in concept analysis. This suggests fine-tuning may restrict broader reasoning capacities. In general, LLMs handled logic alteration better than concept analysis, indicating their robustness in abstract reasoning yet limitations in understanding deeper mathematical con-

	Models	Original		Structural								Representational				Weighted Avg.	
		Seed Samp.	All Samp.	Logic Alteration				Concept Analysis				Question Format			A. Format		
				Gran. Adjust.	Reason Adjust.	Compute. Adjust.	Formul Adjust.	Avg. Perf.	Quest. Under.	Sol. Under.	Crit. Think.	Avg. Perf.	Form. Change.	Form. Comp.	Avg. Perf.		Form. Constraint
MATH	o1-preview	100.0	95.2	100.0	92.5	92.3	77.1	88.9	95.0	95.0	56.0	80.0	95.0	80.0	92.0	90.0	86.7
	GPT-4o	100.0	95.1	100.0	92.5	81.8	71.4	85.8	95.0	90.0	48.0	75.4	100.0	80.0	96.0	85.0	83.8
	GPT-4	100.0	94.6	100.0	80.0	90.9	60.0	78.3	85.0	65.0	48.0	64.6	90.0	60.0	84.0	65.0	74.2
	GPT-3.5	80.0	78.9	75.0	27.5	54.6	25.7	38.7	55.0	45.0	12.0	35.4	35.0	40.0	36.0	5.0	35.8
	Gemini	80.0	91.7	90.0	50.0	81.8	37.1	56.6	60.0	20.0	16.0	30.8	55.0	20.0	48.0	30.0	46.2
	Llama3	60.0	80.7	50.0	12.5	18.2	5.7	17.9	35.0	60.0	4.0	30.8	5.0	60.0	16.0	5.0	26.2
	Metamath	80.0	82.3	70.0	15.0	27.3	11.4	25.5	30.0	25.0	4.0	18.5	35.0	80.0	44.0	20.0	21.3
Average	85.7	88.4	83.6	52.9	63.8	41.2	56.0	65.0	57.1	26.9	47.9	59.3	60.0	59.4	42.9	53.5	
CODE	o1-preview	100.0	92.4	95.0	62.5	46.7	80.0	71.8	95.0	85.0	64.0	80.0	75.0	80.0	76.0	95.0	76.8
	GPT-4o	80.0	90.2	90.0	57.5	53.3	74.3	68.1	80.0	95.0	56.0	75.4	65.0	60.0	64.0	70.0	70.0
	GPT-4	80.0	76.5	90.0	37.5	46.7	50.0	52.3	65.0	80.0	44.0	61.5	65.0	40.0	60.0	55.0	56.7
	GPT-3.5	80.0	64.9	90.0	67.5	40.0	29.4	40.7	60.0	75.0	40.0	56.9	50.0	40.0	48.0	45.0	47.1
	Gemini	80.0	71.9	80.0	32.5	53.3	23.5	41.3	65.0	75.0	44.0	60.0	45.0	40.0	44.0	35.0	47.3
	Llama3	60.0	62.2	45.0	12.5	33.3	11.8	21.1	50.0	50.0	8.0	33.9	25.0	40.0	28.0	20.0	36.6
	CodeLlama	60.0	67.8	80.0	40.0	40.0	11.8	38.5	35.0	35.0	28.0	32.3	40.0	0.0	32.0	40.0	26.3
Average	77.1	75.1	79.1	39.6	44.7	40.1	47.7	64.3	70.7	40.6	57.1	52.1	42.9	50.3	51.4	51.5	

Table 1: Model performance on math and coding across various *Dimensions* (Level III of ontology). All the average reported is weighted average.

	Models	Original		Structural								Representational				Weighted Avg.	Auto-eval Agreement	
		Seed Samp.	All Samp.	Logic Alteration				Concept Analysis				Question Format			A. Format			
				Gran. Adjust.	Reason Adjust.	Compute. Adjust.	Formul Adjust.	Avg. Perf.	Quest. Under.	Sol. Under.	Crit. Think.	Avg. Perf.	Form. Change.	Form. Comp.	Avg. Perf.			Form. Constraint
MATH	o1-preview	100.0	95.2	95.0	90.0	91.7	62.9	82.2	75.0	50.0	20.0	46.2	85.0	80.0	84.0	85.0	71.9	84.3
	GPT-4o	100.0	95.1	100.0	92.5	81.8	71.4	85.8	95.0	90.0	48.0	75.4	100.0	80.0	96.0	85.0	71.3	85.6
	GPT-4	100.0	94.6	90.0	67.5	90.9	51.4	68.9	45.0	15.0	24.0	27.7	85.0	40.0	76.0	80.0	58.3	81.0
	GPT-3.5	80.0	78.9	75.0	25.0	54.5	14.3	34.0	50.0	30.0	8.0	27.7	40.0	60.0	44.0	25.0	32.4	85.2
	Gemini	80.0	91.7	80.0	60.0	81.8	31.4	56.6	45.0	10.0	12.0	21.5	45.0	40.0	44.0	70.0	45.8	86.6
	Llama3	60.0	80.7	65.0	37.5	45.5	14.3	35.8	15.0	15.0	8.0	12.3	30.0	40.0	32.0	30.0	27.8	88.9
	Metamath	80.0	82.3	75.0	22.5	45.5	14.3	32.1	30.0	10.0	8.0	15.4	40.0	100.0	52.0	45.0	30.6	89.8
Average	85.7	88.4	82.9	56.4	70.2	37.1	56.5	50.7	31.4	18.3	32.3	60.7	62.9	61.1	60.0	48.3	85.9	
CODE	o1-preview	100.0	92.4	90.0	55.0	46.7	71.4	58.5	70.0	60.0	48.0	65.5	70.0	80.0	72.0	85.0	65.9	86.4
	GPT-4o	80.0	90.2	90.0	47.5	46.7	65.7	60.9	65.0	60.0	32.0	50.7	60.0	40.0	56.0	70.0	58.2	85.5
	GPT-4	80.0	76.5	100	37.5	46.7	50.0	52.3	75.0	65.0	48.0	61.5	65.0	40.0	60.0	55.0	74.5	71.7
	GPT-3.5	80.0	64.9	73.7	35.0	40.0	29.4	40.7	55.0	70.0	40.0	56.9	50.0	40.0	48.0	45.0	46.3	87.2
	Gemini	80.0	71.9	80.0	32.5	53.3	23.5	41.3	65.0	55.0	24.0	46.2	45.0	40.0	44.0	35.0	43.2	85.4
	Llama3	60.0	62.2	45.0	12.5	33.3	11.8	21.1	50.0	40.0	0.0	46.1	25.0	40.0	28.0	25.0	36.6	91.8
	CodeLlama	60.0	67.8	80.0	40.0	40.0	11.8	38.5	25.0	15.0	16.0	18.0	40.0	0.0	32.0	40.0	26.4	79.0
Average	77.1	75.1	79.8	37.1	43.8	37.7	44.8	57.9	52.1	29.7	49.3	50.7	40.0	48.6	50.7	50.2	83.9	

Table 2: Model performance on math using automated evaluation across various *Dimensions* (Level III of ontology). All the average reported is weighted average.

	Models	G. Adj.	R. Adj.	C. Adj.	F. Adj.	Avg.
MATH	GPT-4	100	77.5	90.91	71.43	81.13
	GPT-3.5	90	50	90.91	40	58.49
	Gemini	95	57.50	63.64	45.71	61.32
CODE	GPT-4	100	50	46.67	55.88	59.43
	GPT-3.5	82.35	42.50	40	26.47	43.40
	Gemini	70.59	25	53.33	26.47	36.79

Table 3: The impact of incorporating the original question and answer into the prompt on the performance of *logic* Target within the GSMORE and HUMAN-EVAL-CORE. The reported average is a weighted average.

	Models	G. Adj.	R. Adj.	C. Adj.	F. Adj.	Avg.
Self-C	GPT-4	95	87.5	90.91	65.71	82.08
	GPT-3.5	60	45	45.45	25.71	41.51
	Gemini	75	45	81.82	40	52.83
POT	GPT-4	95	90	81.82	68.57	83.02
	GPT-3.5	75	57.5	54.55	25.71	50
	Gemini	90	60	63.64	45.71	61.32

Table 4: The impact of using prompting techniques on the performance of *Logic* Target within the GSMORE and HUMAN-EVAL-CORE. Self-C stands for Self-Consistency prompting (Wang et al., 2022) and POT stands for Program of Thought (Chen et al., 2022).

cepts. GPT-4 demonstrated resilience across various question types, outshining others, especially in handling different problem-solving frameworks, although it still struggled more in math than in coding in concept analysis. We include Target-wise (Level II) performance analysis in ???. The latest LLMs, such as o1-preview and GPT-4o, have outperformed their predecessor, GPT-4. However,

there is still a performance decline of 9% for o1-preview and 12% for GPT-4o. Notably, o1-preview demonstrates the highest resilience to perturbation tests in both Math and Coding questions, highlighting the significance of learning to reason.

4.4 Analyses

Incorporation of Original Answer in Prompt.

In Table 3, providing models with the correct answer to the original question along with the prompt significantly improves their ability to solve perturbed questions, particularly in the *Computational Adjustment* dimension. However, performance remains weak in *Formula Adjustment*, highlighting challenges in abstract reasoning despite access to solutions. Notably, even equipped with correct answers, some models like Gemini and GPT-3.5 still fail on simpler question variants, underscoring their low sensitivity to semantic perturbations.

Prompting Techniques. In Table 4, different prompting techniques greatly influence model performance in *Logic Alteration* tasks. The Program-of-Thoughts technique notably boosts reasoning capabilities in closed-source models by reducing logical errors, leading to better performance in symbolic manipulation for GPT-4. Conversely, the Self-Consistency method shows only minor improvements and even a performance decline in the Gemini model, suggesting difficulties in effective in-context learning for new unseen tasks.

Identified Vulnerabilities in Reasoning. The *Formulation Adjustment* dimension presents a significant challenge to both closed-source and open-source models, largely due to the demands of abstract reasoning. Instead of reasoning a number or code as the answer, this involves the manipulation of abstract math and coding concepts in the logical space behind the surface of the problem. For example, in the *WhatIf* category, models must hypothesize outcomes by changing certain events under consistent conditions, which requires a nuanced grasp of the problem-solving framework. The *Critical Thinking* dimension tests a model’s ability to scrutinize relationships between pieces of information, demanding a comprehensive analysis to identify inconsistencies without a predefined solution path. This emphasizes the necessity for models to thoroughly understand and navigate through all possible avenues to effectively resolve conflicts or discrepancies. Furthermore, the *Format Change* dimension poses difficulties to models like GPT-3.5 attempt to follow these constraints but often fail to maintain the integrity of their reasoning when adapting to new formats, highlighting a lack of flexibility in handling varied task demands.

5 Discussions

5.1 Difficulty Change by Perturbations

Performance drop may result from an increased scale of reasoning or a higher level of abstract reasoning required. To explore this, the experiments measured changes in the scale and depth of reasoning by comparing the number of reasoning steps and the depth required for each perturbed question against its original version. The difficulty was also evaluated through A/B testing and by recording human performance and response times in various perturbation categories, as detailed in Appendix E (summarize in Table 5). Table 9 conducted human evaluation on 44 perturbation types, 11 increased the number of reasoning steps needed, and 10 required deeper reasoning compared to the original questions. Although more complex questions increased the time humans needed to respond, human performance remained almost the same. However, LLMs showed a notable decrease in performance - 12.4% for increased reasoning steps and 3.7% for deeper reasoning compared to when the number of reasoning steps or difficulty remained the same. Further, there was also more than a thirty percent change in model performance for perturbed questions of equal difficulty, indicating that models have troubling reasoning robustness issues.

Category	Human Acc(Δ)	Model Acc(Δ)	Time Consumption
Number of Reasoning Steps			
⊕	95.2(-4.8)	32.0(-44.0)	178%
⊖	98.4(-1.6)	44.4(-31.6)	39%
Reasoning Depth			
⊕	97.3(-2.7)	38.4(-37.6)	113%
⊖	97.7(-2.3)	42.1(-33.9)	62%

Table 5: Summary of Human and Model Accuracy, and Time Consumption by Number of Steps and Conceptual Depth of Questions. ⊕ indicates an increase; ⊖ indicates no change in reasoning steps or depth; Δ stands for the performance change relative to the original question.

5.2 Design Choices behind Ontology

An effective perturbation type maintains control over most variables while introducing only unidirectional changes to the original questions. Ideally, these perturbations should be noticeable to humans yet subtle enough that the required changes in skills for solving these variant questions do not significantly alter human reasoning, due to inherent human cognitive priors. Any data perturbation ontology necessitates predefined assumptions about

which aspects of the data are mutable and how these changes might influence the outcomes. Therefore, recognizing and understanding these assumptions is crucial for enhancing future data augmentation efforts.

5.3 Scaling to More Instances

Our human-in-the-loop approach may restrict scaling to more instances; however, our primary focus is on evenly evaluating performance across various perturbation categories, rather than on scaling. Nonetheless, it is feasible to expand the dataset through a multi-agent approach (Wang et al., 2024), which selectively filters out the more challenging samples. Our initial experiments, indicate that GPT-4 can achieve a perturbation success rate of over 90% for perturbation categories with same reasoning depth in Table 9.

6 Related Works

A variety of datasets are developed to assess reasoning capabilities across multiple domains. In causal reasoning, the key datasets are by Huang et al. (2023); Bondarenko et al. (2022). For coding, notable contributions have been made by Chen et al. (2021b); Austin et al. (2021). Additionally, mathematical reasoning has been addressed at different educational levels: grade-school (Cobbe et al., 2021b,a), high school (Hendrycks et al., 2021), and college level (Sawada et al., 2023; Zheng et al., 2021). Despite the advancements in large language models (Ahn et al., 2024), recent studies (Mondorf and Plank, 2024) contend that these models more closely resemble stochastic parrots (Bender et al., 2021) than true systematic reasoners, exhibiting significant limitations particularly in scenarios not covered by their training data (Bender et al., 2021; Wan et al., 2024).

Therefore, recent works have focused on the robustness of reasoning under various perturbations that alter reasoning questions. Different domain-specific methods have been proposed for generating test cases for reasoning tasks (Yu et al., 2023a; Wu et al., 2023), as summarized in Table 6. For mathematics, contemporary works have employed numerical or symbolic substitutions (Li et al., 2024a; Zhou et al., 2023; Meadows et al., 2023; Wang et al., 2024; Patel et al., 2021), insertion of irrelevant distractors (Shi et al., 2023; Li et al., 2023), functional equivalence (Srivastava et al., 2024), and reverse prediction (Yu et al., 2023b; Berglund et al.,

2023; Deb et al., 2023) to uncover conceptual errors (Sanyal et al., 2022), cognitive biases (Dasgupta et al., 2022), and sensitivity to reasoning context (Wu et al., 2023). Similarly, Abstract Syntax Tree guided perturbations are performed on the domain of code (Naik et al., 2023; Allamanis et al., 2021). However, to our knowledge, perturbation methods across the domain of reasoning tasks are lacking. In this work, we consolidate and develop a broader underlying ontology that connects and expands upon previous methods for perturbing reasoning datasets. This new framework is both systematic and hierarchical, and it is readily adaptable to various domains, including mathematics and coding.

A distinct line of research focuses on evaluating reasoning through non-conclusion-based assessments, which provide deeper insights into models' reasoning behaviors. For example, ReasonEval (Xia et al., 2024) analyzes the *reliability* and *redundancy* of generated reasoning steps, highlighting the qualitative aspects of reasoning. Similarly, Li et al. (2024b) perform error identification within the reasoning path rather than simply identifying the correct answer. Furthermore, Zeng et al. (2023) explore the robustness of models across varied potential reasoning paths, reinforcing the idea that higher accuracy does not necessarily improve reasoning quality. Our ontology extends these approaches by including perturbations on various concepts related to reasoning path and question understanding, thereby enriching the framework for assessing reasoning capabilities.

Beyond perturbation-based question generation, educational measurement research has long sought assessment instruments that can **reliably discriminate** learners according to specific mathematical skills. For example, Homjan et al. (2022) develop a diagnostic test targeting grade-4 addition and subtraction, while Rani and Anisha (2017) design and standardise a geometry achievement test for grade-9 students. These studies emphasise psychometric qualities such as item difficulty, discrimination index, and overall test reliability rather than the micro-level composition of individual questions. The resulting insights into difficulty calibration, diversity of item types, and longitudinal tracking of learner progress can inform how we evaluate the evolving capabilities of large language models (LLMs) on mathematically grounded tasks.

7 Conclusion

Our study evaluated the robustness of several prominent Large Language Models (LLMs) in handling mathematical and coding problems. By employing an ontology for random perturbations on questions from the GSM8K and HumanEval datasets, we crafted two specialized datasets, GSMORE and HUMAN-EVAL-CORE, containing 216 and 219 questions, respectively. These datasets target a broad variations of mathematical and coding problem-solving and analytical skills, resulting in notable performance drops in LLMs upon evaluation. The introduction of GSMORE and HUMAN-EVAL-CORE provides a new framework for assessing LLMs’ abilities in mathematics and coding, while also revealing their vulnerabilities in consistent reasoning across different formats. This research highlights the complex challenges that LLMs face, stressing the importance of continued exploration into their strengths and weaknesses in logical reasoning tasks.

Limitations

Despite our attempt to construct a novel systematic ontology to evaluate an LM’s “real” robustness and reasoning capabilities in structured reasoning tasks, it may not precisely reflect LLM’s true ability due to several factors.

Incompleteness. In our endeavor to develop a comprehensive ontology for evaluating Language Models’ (LMs) responses to perturbed questions across various reasoning scenarios, we recognize significant limitations. Firstly, despite our efforts, the ontology may not fully capture all essential aspects of reasoning abilities, lacking in breadth and depth. Secondly, the complexity within each reasoning category can vary significantly. For instance, within the *Computation Demand* category, adjusting the number of digits in mathematical operations allows us to modulate the reasoning challenge. However, creating a benchmark that exhaustively encompasses all facets of reasoning behavior is an unattainable goal. Such an exhaustive compilation is beyond the scope of any single study and necessitates collective efforts from the broader research community.

Scalability. The size of our dataset is constrained due to the human in the loop required for its preparation. Each question generated by GPT-4 needs to be meticulously reviewed to ensure it is solvable

and accurately reflects the intended perturbation specific to its category, without introducing unintended modifications. Furthermore, confirming the accuracy of answers is a critical step, as many questions do not yield answers that exactly match a predefined format. This verification process limits our ability to expand the dataset on a large scale, as it relies on manual effort.

8 Potential Risks

Not applicable.

9 Ethical Considerations

Not applicable.

Acknowledgements

The work is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG3-GV2023-010). Any opinions, findings, and conclusions, or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore.

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*, abs/2402.00157.
- Miltiadis Allamanis, Henry Jackson-Flux, and Marc Brockschmidt. 2021. [Self-supervised bug detection and repair](#). In *Neural Information Processing Systems*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and Charles Sutton. 2021. [Program synthesis with large language models](#). *ArXiv*, abs/2108.07732.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.
- Lukas Berglund, Meg Tong, Max Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. 2023. [The reversal curse: LLMs trained on "a is b" fail to learn "b is a"](#). *ArXiv*, abs/2309.12288.
- Alexander Bondarenko, Magdalena Wolska, Stefan Heindorf, Lukas Blübaum, Axel-Cyrille Ngonga

- Ngomo, Benno Stein, Pavel Braslavski, Matthias Hagen, and Martin Potthast. 2022. [Causalqa: A benchmark for causal question answering](#). In *International Conference on Computational Linguistics*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021a. [Evaluating large language models trained on code](#).
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021b. [Evaluating large language models trained on code](#).
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2022. [Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks](#). *ArXiv*, abs/2211.12588.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021a. [Training verifiers to solve math word problems](#). *ArXiv*, abs/2110.14168.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021b. [Training verifiers to solve math word problems](#). *ArXiv*, abs/2110.14168.
- Ishita Dasgupta, Andrew Kyle Lampinen, Stephanie C. Y. Chan, Antonia Creswell, Dharshan Kumaran, James L. McClelland, and Felix Hill. 2022. [Language models show human-like content effects on reasoning](#). *ArXiv*, abs/2207.07051.
- Aniruddha Deb, Neeva Oza, Sarthak Singla, Dinesh Khandelwal, Dinesh Garg, and Parag Singla. 2023. [Fill in the blank: Exploring and enhancing llm capabilities for backward reasoning in math word problems](#). *ArXiv*, abs/2310.01991.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2022. [Pal: Program-aided language models](#). *ArXiv*, abs/2211.10435.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Xiaodong Song, and Jacob Steinhardt. 2020. [Measuring massive multitask language understanding](#). *ArXiv*, abs/2009.03300.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the math dataset](#).
- Suchart Homjan, Krapan Sri-ngan, and Wanida Homjan. 2022. [Construction of diagnostic test in mathematics on addition and subtraction basic for primary students](#). *Journal of Education and Learning*, 11(3):88–94.
- Yinya Huang, Ruixin Hong, Hongming Zhang, Wei Shao, Zhicheng YANG, Dong Yu, Changshui Zhang, Xiaodan Liang, and Linqi Song. 2023. [Clomo: Counterfactual logical modification with large language models](#). *ArXiv*, abs/2311.17438.
- Albert Qiaochu Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Giana Lengyel, Guillaume Lample, Lucile Saulnier, L'elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *ArXiv*, abs/2310.06825.
- Divyansh Kaushik, Amrith Setlur, Eduard Hovy, and Zachary C Lipton. 2021. [Explaining the efficacy of counterfactually augmented data](#). *International Conference on Learning Representations (ICLR)*.
- Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. 2024a. [Gsm-plus: A comprehensive benchmark for evaluating the robustness of llms as mathematical problem solvers](#). *ArXiv*, abs/2402.19255.
- Xiaoyuan Li, Wenjie Wang, Moxin Li, Junrong Guo, Yang Zhang, and Fuli Feng. 2024b. [Evaluating mathematical reasoning of large language models: A focus on error identification and correction](#).

- Zekun Li, Baolin Peng, Pengcheng He, and Xifeng Yan. 2023. [Do you really follow me? adversarial instructions for evaluating the robustness of large language models.](#) *ArXiv*, abs/2308.10819.
- Jordan Meadows, Marco Valentino, Damien Teney, and André Freitas. 2023. [A symbolic framework for evaluating mathematical reasoning and generalisation with transformers.](#)
- Philipp Mondorf and Barbara Plank. 2024. [Beyond accuracy: Evaluating the reasoning behavior of large language models - a survey.](#) *ArXiv*, abs/2404.01869.
- Atharva Naik, Soumitra Das, Jyothi Vedurada, and Somak Aditya. 2023. [Sync: A structurally guided hard negative curricula for generalizable neural code search.](#) In *International Joint Conference on Natural Language Processing*.
- OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Arkil Patel, S. Bhattamishra, and Navin Goyal. 2021. [Are nlp models really able to solve simple math word problems?](#) In *North American Chapter of the Association for Computational Linguistics*.
- Reena Rani and Anisha. 2017. [Construction and standardization of mathematics achievement test for ixth grade students.](#) *Educational Quest: An International Journal of Education and Applied Social Science*, 8(2):651–655.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, I. Evtimov, Joanna Bitton, Manish P Bhatt, Cristian Cantón Ferrer, Aaron Grattafori, Wenhan Xiong, Alexandre D’efossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. [Code llama: Open foundation models for code.](#) *ArXiv*, abs/2308.12950.
- Soumya Sanyal, Zeyi Liao, and Xiang Ren. 2022. [Robustlr: A diagnostic benchmark for evaluating logical robustness of deductive reasoners.](#)
- Tomohiro Sawada, Daniel Paleka, Alexander Havrilla, Pranav Tadepalli, Paula Vidas, Alexander Kranias, John J. Nay, Kshitij Gupta, and Aran Komatsuzaki. 2023. [Arb: Advanced reasoning benchmark for large language models.](#) *ArXiv*, abs/2307.13692.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Huai hsin Chi, Nathanael Scharli, and Denny Zhou. 2023. [Large language models can be easily distracted by irrelevant context.](#) In *International Conference on Machine Learning*.
- Saurabh Srivastava, B AnnaroseM, V AntoP, Shashank Menon, Ajay Sukumar, T AdwaithSamod, Alan Philipose, Stevin Prince, and Sooraj Thomas. 2024. [Functional benchmarks for robust evaluation of reasoning performance, and the reasoning gap.](#) *ArXiv*, abs/2402.19450.
- Yuxuan Wan, Wenxuan Wang, Yiliu Yang, Youliang Yuan, Jen tse Huang, Pinjia He, Wenxiang Jiao, and Michael R. Lyu. 2024. [A & b == b & a: Triggering logical reasoning failures in large language models.](#) *ArXiv*, abs/2401.00757.
- Siyuan Wang, Zhuohan Long, Zhihao Fan, Zhongyu Wei, and Xuanjing Huang. 2024. [Benchmark self-evolving: A multi-agent framework for dynamic llm evaluation.](#) *ArXiv*, abs/2402.11443.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Huai hsin Chi, and Denny Zhou. 2022. [Self-consistency improves chain of thought reasoning in language models.](#) *ArXiv*, abs/2203.11171.
- Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim, Jacob Andreas, and Yoon Kim. 2023. [Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks.](#)
- Shijie Xia, Xuefeng Li, Yixin Liu, Tongshuang Wu, and Pengfei Liu. 2024. [Evaluating mathematical reasoning beyond accuracy.](#) *arXiv preprint arXiv:2404.05692*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. [Wizardlm: Empowering large language models to follow complex instructions.](#) *arXiv preprint arXiv:2304.12244*.
- Long Long Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zheng Li, Adrian Weller, and Weiyang Liu. 2023a. [Metamath: Bootstrap your own mathematical questions for large language models.](#) *ArXiv*, abs/2309.12284.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2023b. [Metamath: Bootstrap your own mathematical questions for large language models.](#)
- Zhongshen Zeng, Pengguang Chen, Shu Liu, Haiyun Jiang, and Jiaya Jia. 2023. [Mr-gsm8k: A meta-reasoning benchmark for large language model evaluation.](#)
- Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. 2021. [Minif2f: a cross-system benchmark for formal olympiad-level mathematics.](#) *ArXiv*, abs/2109.00110.
- Zihao Zhou, Qiufeng Wang, Mingyu Jin, Jie Yao, Jianan Ye, Wei Liu, Wei Wang, Xiaowei Huang, and Kaizhu Huang. 2023. [Mathattack: Attacking large language models towards math solving ability.](#) In *AAAI Conference on Artificial Intelligence*.

Variant Name	Parent Domain (Dataset)	Type	Annotation	Dimension	# Categories
SVAMP (Patel et al., 2021) *	math (ASDiv-A)	Equation-formed list	Human (Q,A)	V L	3
MetaMathQA (Yu et al., 2023a)	math (GSM8K, MATH)	Dynamic CheckList	GPT-3.5-Turbo	V R	4
GSM-HARD (Gao et al., 2022)	math (GSM8K)	Program-formed CheckList	Codex (Q,A), Human (A)	V	1
GSM-IC (Shi et al., 2023) *	math (GSM8K)	Static Checklist	Human (Q)	L	3
GSM-PLUS (Li et al., 2024a) *	math (GSM8K)	Dynamic CheckList	GPT-4, Human (Q,A)	R L T C	8
GSMORE-HUMANEVAL-CORE (Our) *	math (GSM8K), code (HumanEval)	Dynamic Ontology	GPT-4, Human (Q,A)	R L C T F S V	44

Table 6: Overview of variants in reasoning datasets arising from perturbation types. * refers to datasets specifically designed to evaluate the robustness of model performance. Different letters represent different perturbation types: [R]epresentational Change, [L]ogic Alteration, [C]oncept Analysis, Critical [T]hinking, [F]ormulation Adjustment, [S]caling, [V]alue Replacement.

Aspect (Level I)	Target (Level II)	Dimension (Level III)	Category (Level IV)	Math	Code
Structural Perturbation <i>Def: Modification on specific aspects of logic or concepts to alter the reasoning process required to reach the answer</i>	Logic Alteration <i>Def: Modifications to the reasoning framework or logic underpinning a problem.</i>	Granularity Adjustment <i>Def: Fine-grained sub-tasks of the original question</i>	G1. Remove Constraint G2. Partial Solution G3. Solution Plan G4. Detail Expansion	Remove Constraint Median Inquiry Solution Plan Detail Elaboration	Remove Constraint Helper Function Solution Plan Example Detail
		Reasoning Adjustment <i>Def: Target at logical structure of the original</i>	G5. Add Restriction G6. Subsequent Question G7. Concurrent Question G8. Change Question G9. Info Recombination G10. Domain Knowledge G11. Complex Reality G12. General Solution	Restrict Question Further Question Parallel Question Change Query Info Recombination Theoretical Challenge Value Probability Code Implementation	Restrict Requirement Further Requirement Parallel Requirement Change Docstring Info Recombination Code Import Example Boundary Higher Order
		Computation Adjustment <i>Def: Target at values or entities</i>	G13. Computation Demand G14. Change Value G15. Change Operation	Value Big Change Subject Change Calculation	Generalize Parameter Parameter Content Variable Type
		Formulation Adjustment <i>Def: Reformulate question for solution form to be an abstract expression.</i>	G16. Symbolic Response G17. Value Relationship G18. Variable Group G19. Backward Reasoning G20. Counterfactual G21. Solve Value G22. Identify Range	Variable Response Variable Relation Variable Scaling Variable Adaptation WhatIf Question Solve X Variable Range	Code Execution Parameter Relationship Variable Substitution Reverse Engineering WhatIf Code Solve Input Variable Range
	Concept Analysis <i>Def: Examination and Analysis of the underlying concepts and principles of a problem</i>	Question Understanding <i>Def: Interpretation of the information inside the question</i>	G23. Inherent Premise G24. Complete Missing G25. Question Formulation G26. Add Misinformation	Identify Assumption Info Sufficiency Question Formulation Introduce Distraction	Test Case Incomplete Answer Question Formulation Introduce Bias
		Solution Understanding <i>Def: Assessment of the problem-solving processes</i>	G27. Optimize Solution G28. Step Functionality G29. Theoretical Basis G30. Cost Analysis	Info Necessity Step Necessity Theoretical Basis Solution Efficiency	Reduce Complexity Step Necessity Theoretical Basis Code Complexity
		Critical Thinking <i>Def: Identification of noise, inaccuracies and inconsistencies</i>	G31. Seek Clarification G32. Conditional Analysis G33. Conflicting Information G34. Surface Error G35. Hidden Error	Introduce Ambiguity Discuss Separately Introduce Contradiction Value Uncommon Value Error	Example Requirement Incomplete Requirement Wrong Example Runtime Error Logical Error
	Representational Perturbation <i>Def: Preservation of the underlying logic and conceptual framework, but modification of the encoding or representation</i>	Question Format <i>Def: Direct modification on the encoding of the question while keeping the logical structure intact</i>	G36. Setting Rephrase G37. Change Sequence G38. Close Format G39. Data Restructuring	Change Setting Change Sequence True False Value Structuring	Realworld Usecase Parameter Sequence True False Complex Docstring
		Format Comparison <i>Def: Comparing two problem of different forms</i>	G40. Identical Problem	Identical Question	Identical Code
		Answer Format <i>Def: Indirect modification on the output form</i>	Format Constraint <i>Def: Add constraint on the solution</i>	G41. Reasoning Format G42. Reasoning Style G43. Alternative Answer G44. New Rule	Binary Coded X Language Alternative Answer Define Rules

Table 7: Our proposed ontology framework with domain, dimension, mathematical and code realization categories.

A Recommendations

Based on our findings, we make the following recommendations as strategies to address the weaknesses we identified in the logical reasoning competencies of LLMs.

Diversify the Datasets and Formats Used in Fine-tuning. If a model is trained exclusively on a single problem-solving method, its capability to adapt to questions presented in different formats and solve a diverse array of problems diminishes. To counter this, we suggest boosting the model’s resilience

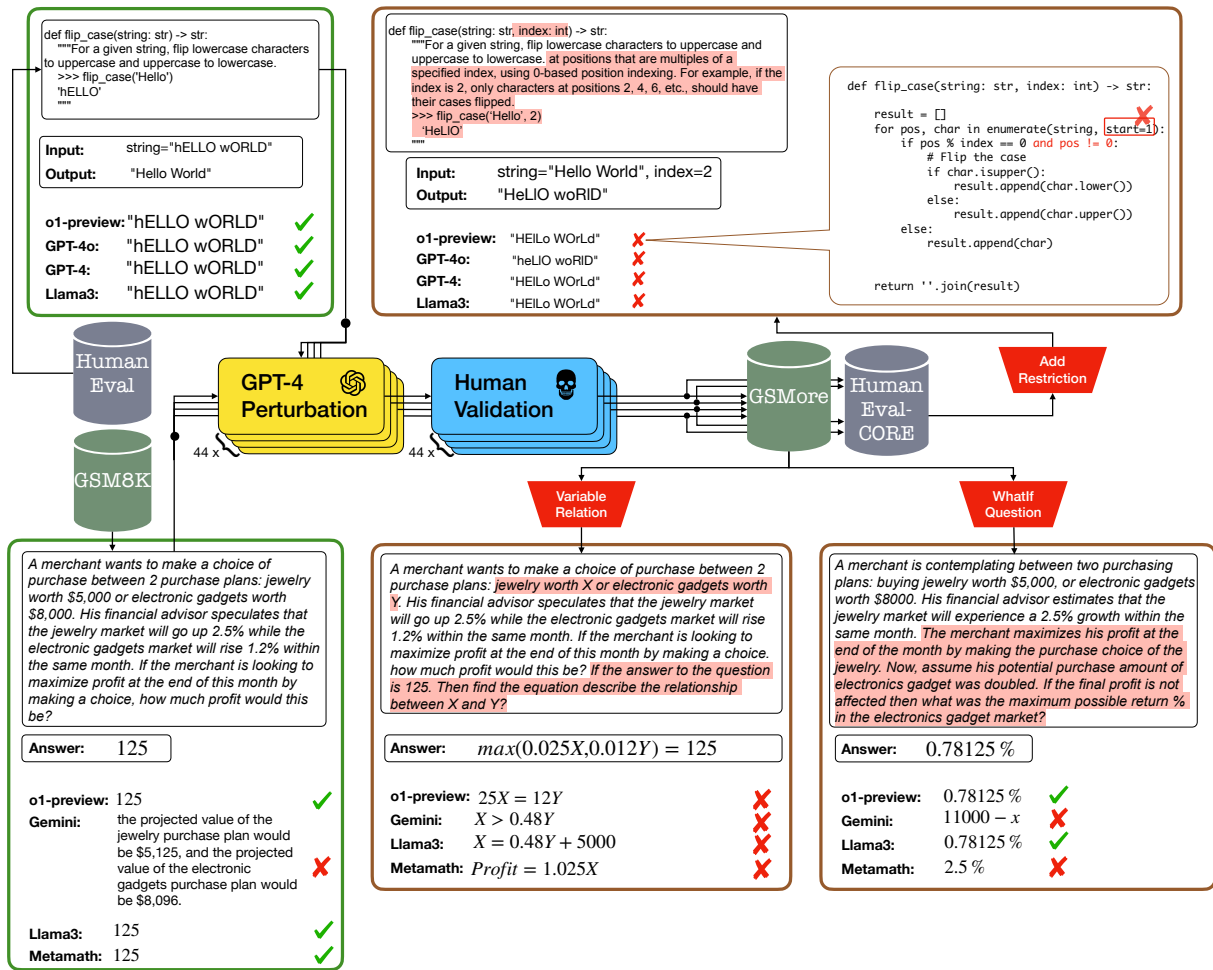


Figure 2: A semi-automated pipeline of creating GSMORE, from five simple questions from GSM8k. An analogous pipeline is used to create the perturbations of the coding questions from HumanEval, named HUMAN EVAL-CORE.

to perturbations by fine-tuning it with datasets in a variety of formats and adding augmented instructions.

Include More Complex Open-Ended Questions. It is also crucial to move beyond simple multiple-choice questions, and include open-ended questions that test the model’s comprehension of mathematical concepts in the fine-tuning dataset, as this enhances its overall understanding and interpretation of questions.

B Experimental Results

We have also extend our evaluation on smaller sized open sourced models:

C Dataset Details

In particular, there are a total of 5 math questions for each category except *Change Subject* and *Reverse Engineering*, which have 3 and 4 questions, respectively, in GSMORE. Likewise, all but *Reverse Engineering* perturbation—with 4 questions—have 5 coding questions in HUMAN EVAL-CORE.

D Experiment Details

D.1 Model Details

- We use o1-preview with model version dated "2024-08-06".
- We use GPT-4o with model version dated "2024-08-06".
- We use GPT-4 with model version named "1106-Preview".

	Logic Alteration (%)				Format (%)	
	Comp. Adj.	Formul. Adj.	Gran. Adj.	Reason Adj.	Format Change	Format Constraint
Math deepseekr1¹	90.9	54.3	90.0	67.5	80.0	60.0
Math gemma3²	72.7	28.6	85.0	55.0	70.0	40.0
Math qwen2.5³	81.8	54.3	100.0	70.0	90.0	65.0
Code deepseekr1¹	53.3	60.0	92.3	69.7	93.3	100.0
Code gemma3²	33.3	40.0	69.2	60.6	86.7	0.0
Code qwen2.5³	53.3	40.0	69.2	27.3	60.0	100.0

Table 8: Percentage scores for math and code models on logic–alteration and format dimensions.

- For GPT-3.5, we use version "GPT-3.5-turbo-0613".
- We use Gemini-1.5-pro for Gemini model.
- Llama3-Instruct <https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>.
- MetaMath-70B-V1.0 <https://huggingface.co/meta-math/MetaMath-70B-V1.0>.
- CodeLlama-70B-Instruct <https://huggingface.co/codellama/CodeLlama-70b-Instruct-hf>.

D.2 Filtering Criteria

- (i) **Human Understandability:** The generated questions should be comprehensible to humans. The language, structure, and presentation of the questions should be clear and easy to understand. Vague or confusing questions should be rejected.
- (ii) **Logical Coherence:** The questions must make logical sense. They should not contain contradictions², nonsensical premises, or incoherent elements.
- (iii) **Instruction Adherence:** The generated questions should closely adhere to the instructions in the prompt for the specific perturbation type. The question should not deviate from the intended method of perturbation.

E Benchmark Difficulty Evaluation

The evaluation of difficulty was conducted by three undergraduate students. Each participant was presented with questions to solve on paper, without access to calculators or computers. Their task completion time for each question was recorded. The students also noted any changes in the number of steps required to solve the perturbed questions compared to the original ones, indicating whether the steps increased or remained approximately the same. Furthermore, they assessed whether the perturbed variants required the application of more advanced mathematical concepts or skills.

F Original Questions from GSM8K

The following selected questions are from the GSM8K dataset, specifically chosen for their variations in complexity. Each of the five questions requires between 3 to 7 steps to solve, illustrating the range of reasoning complexity present in the GSM8K dataset. These questions span a wide array of everyday topics that involve the application of mathematical principles, including physical dimensions, profit maximization, purchasing decisions, time management, and solving multi-variable equations. Those 5 questions demands diversity of mathematical problem-solving skills in different situations.


F.1 Question 1

²Except for the *conflicting information* (G33.) type, where we intentionally introduce contradictions.

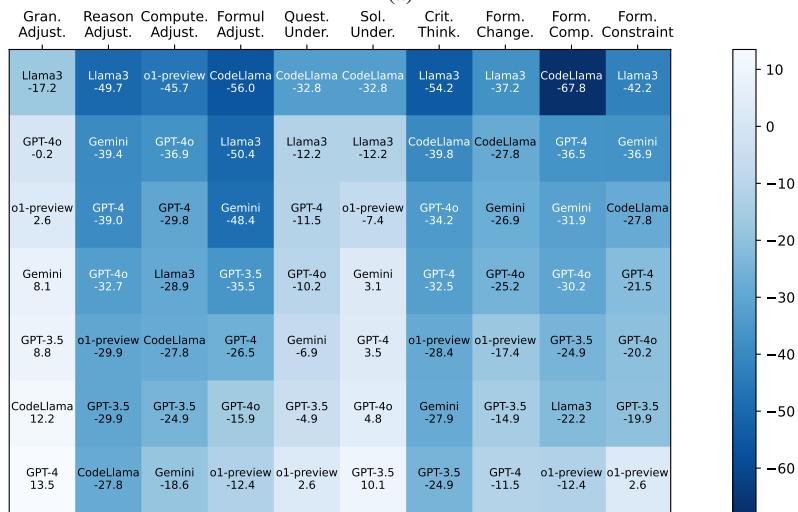
Dimension	Category	Human Acc	Model Acc	Time Consump	Steps	Reasoning Depth
Granularity Adjustment	Remove Constraint	100	84	-70	⊖	⊖
	Partial Solution	100	70	-40	⊖	⊖
	Solution Plan	100	76	-50	⊖	⊖
	Detail Expansion	100	70	-50	⊖	⊖
Reasoning Adjustment	Add Restriction	100	22	+100	⊕	⊖
	Subsequent Question	100	34	+50	⊖	⊖
	Concurrent Question	100	36	+150	⊖	⊖
	Change Question	100	42	-70	⊖	⊕
	Info Recombination	87	28	+40	⊕	⊖
	Domain Knowledge	80	56	+450	⊕	⊕
	Complex Reality	100	32	+100	⊕	⊖
	General Solution	100	24	+0	⊖	⊕
Computation Adjustment	Computation Demand	100	36	+20	⊖	⊖
	Change Value	100	56	-10	⊖	⊖
	Change Operation	100	66	+0	⊖	⊖
Formulation Adjustment	Symbolic Response	100	42	+100	⊖	⊕
	Value Relationship	93	20	+100	⊖	⊕
	Variable Group	100	24	+140	⊕	⊕
	Backward Reasoning	100	26	+100	⊖	⊕
	Counterfactual	100	18	+160	⊕	⊖
	Solve Value	100	28	+140	⊖	⊖
	Identify Range	93	26	-40	⊖	⊖
Question Understanding	Inherent Premise	100	38	+160	⊖	⊖
	Complete Missing	100	60	-50	⊖	⊖
	Question Formulation	93	50	+200	⊖	⊖
	Add Misinformation	100	68	+50	⊖	⊖
Solution Understanding	Optimize Solution	100	50	+160	⊕	⊕
	Step Functionality	100	42	+100	⊖	⊕
	Theoretical Basis	100	62	-50	⊖	⊖
	Cost Analysis	100	58	+50	⊖	⊕
Critical Thinking	Seek Clarification	80	26	-50	⊖	⊖
	Conditional Analysis	93	16	+200	⊕	⊖
	Conflicting Information	100	8	+50	⊖	⊖
	Surface Error	100	44	+50	⊖	⊖
	Hidden Error	93	30	+200	⊖	⊖
Format Change	Setting Rephrase	100	50	+0	⊖	⊖
	Change Sequence	100	52	+0	⊖	⊖
	Close Format	93	36	+20	⊖	⊖
	Data Restructuring	100	40	+160	⊕	⊖
Format Comparison	Identical Problem	87	42	+20	⊖	⊖
Format Constraint	Reasoning Format	100	30	+200	⊕	⊖
	Reasoning Style	100	34	+170	⊖	⊖
	Alternative Answer	100	28	+60	⊖	⊖
	New Rule	87	36	+250	⊕	⊖
Average		97.7	41.2	+74.3	N/A	N/A

Table 9: Comparison of Average Baselines: Human vs. Models. Displays accuracy rates for participants and models, and time change percentage for solving perturbed vs. original questions. ⊕ indicates an increase; ⊖ signifies equal reasoning depth.

A merchant wants to make a choice of purchase between 2 purchase plans: jewelry worth \$5,000 or electronic gadgets worth \$8,000. His financial advisor speculates that the jewelry market will go up 2.5% while the electronic gadgets market will rise 1.2% within the same month. If the merchant is looking to maximize profit at the end of this month by making a choice, how much profit would this be?



(a)



(b)

Figure 3: Performance drop of LLMs for various dimensions (Level III) of perturbations in (a) GSMORE and (b) HUMAN-EVAL-CORE as compared to the performance on the original GSM8K and HumanEval test datasets.

Answer: If he purchases jewelry, he will make a profit of 2.5% which is $5000 \cdot (2.5/100) = 125$. If he purchases electronic gadgets, he will make a profit of 1.2% which is $8000 \cdot (1.2/100) = 96$. If he wants to maximize profit, since $125 > 96$, he will choose to purchase jewelry, thereby making a profit of 125

F.2 Question 2

Question 2: John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes?

Answer: The walls subtract $2 \cdot 1 = 2$ inches from each dimension. So each box has $5 - 2 = 3$ inch width It also has a $6 - 2 = 4$ inch height. Finally, it has a $4 - 2 = 2$ inch depth. So the inner volume of one box is $4 \cdot 3 \cdot 2 = 24$ cubic inches. So in total the inner volume of the 3 boxes is $3 \cdot 24 = 72$ cubic inches

F.3 Question 3

Question 3: 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?

Answer: The discount price of one glass is $60/100 * 5=3$. If every second glass is cheaper, that means Kylar is going to buy $16 / 2 = 8$ cheaper glasses. So for the cheaper glasses, Kylar is going to pay $8 * 3 = 24$. And for the regular-priced glasses, Kylar will pay $8 * 5 = 40$. So in total Kylar needs to pay $24 + 40 = 64$ for the glasses he wants to buy.

F.4 Question 4

Question 4: Vicki is planning a pop concert at her high school. The show will be 2 hours. She is allowing each group 2 minutes to get on stage, 6 minutes to perform, and then 2 minutes to exit the stage. If she allows a 10-minute intermission, how many groups can perform in the concert?

Answer: First, we should convert the 2 hours of showtime into minutes for our calculations. Since there are 60 minutes in 1 hour, the show will be $2 * 60 = 120$ minutes. Of those 120 minutes, 10 will be used for intermission, so $120 - 10 = 110$ minutes for performances. Each group will use 2 minutes to get on stage + 6 minutes to perform + 2 minutes to exit the stage = 10 minutes of show time. Of the 110 minutes of performances, 10 are used per group, so $110 \text{ minutes} / 10 = 11$ groups can perform.

F.5 Question 5

Question 5: Together Lily, David, and Bodhi collected 43 insects. Lily found 7 more than David. David found half of what Bodhi found. How many insects did Lily find?

Answer: Let B = the number of insects Bodhi collected. David = $B/2$, Lily = $B/2 + 7$. $B + B/2 + 7 = 43$. Simplify: $2B = 36$. Simplify B = 18 insects. David = $18/2 = 9$ insects. Lily = $9 + 7 = 16$ insects. Lily found 16 insects.

G Original Questions from HumanEval

The following selected questions are from the HumanEval dataset, specifically chosen for their variations in complexity. Each of the five questions requires different number of lines code to solve, illustrating the range of reasoning complexity present in the HumanEval dataset. These questions includes basic programming concepts such as string manipulation, list indexing, classic algorithm, math problem and state conditions. Those 5 questions demands diversity of programming skills and concepts in different situations.

G.1 Question 1

```
1 def flip_case(string: str) -> str:
2
3     """For a given string, flip lowercase characters to uppercase and uppercase to
4     lowercase.
5
6     >>> flip_case('Hello')
7     'hELLO'
8     """
9     return string.swapcase()
```

G.2 Question 2

```
1 def greatest_common_divisor(a: int, b: int) -> int:
2
3     """ Return a greatest common divisor of two integers a and b
```

```

4
5 >>> greatest_common_divisor(3, 5)
6 1
7 >>> greatest_common_divisor(25, 15)
8 5
9 """
10
11 while b:
12     a, b = b, a \% b
13 return abs(a)

```

G.3 Question 3

```

1 def derivative(xs: list):
2
3     """ xs represent coefficients of a polynomial.
4     xs[0] + xs[1] * x + xs[2] * x^2 + ....
5     Return derivative of this polynomial in the same form.
6
7     >>> derivative([3, 1, 2, 4, 5])
8     [1, 4, 12, 20]
9     >>> derivative([1, 2, 3])
10    [2, 6]
11    """
12    if len(xs) == 1: return [0]
13    if len(xs) == 0: return []
14    return [(i * x) for i, x in enumerate(xs)][1:]

```

G.4 Question 4

```

1 def sum_squares(lst):
2
3     """
4     This function will take a list of integers. For all entries in the list, the
5     function shall square the integer entry if its index is a
6     multiple of 3 and will cube the integer entry if its index is a multiple of 4
7     and not a multiple of 3. The function will not
8     change the entries in the list whose indexes are not a multiple of 3 or 4. The
9     function shall then return the sum of all entries.
10
11     Examples:
12     For lst = [1,2,3] the output should be 6
13     For lst = [] the output should be 0
14     For lst = [-1,-5,2,-1,-5] the output should be -126
15     """
16
17    result =[]
18    for i in range(len(lst)):
19        if i%3 == 0:
20            result.append(lst[i]**2)
21        elif i% 4 == 0 and i%3 != 0:
22            result.append(lst[i]**3)
23        else:
24            result.append(lst[i])
25    return sum(result)

```

G.5 Question 5

```

1 def is_nested(string):
2
3     """
4     Create a function that takes a string as input which contains only square
5     brackets.
6     The function should return True if and only if there is a valid subsequence of
7     brackets
8     where at least one bracket in the subsequence is nested.
9     Examples:
10    [[]] output: True

```

```

9      [[]] output: False
10     [] output: False
11     [[]] output: True
12     [[]][[] output: True
13     """
14
15     stack = []
16     depth = 0
17     for i, char in enumerate(string):
18         if char == '[':
19             stack.append('[')
20             if depth > 0:
21                 depth -= 1
22         elif char == ']':
23             if len(stack) > 0:
24                 stack.pop()
25                 depth += 1
26             if depth >= 2:
27                 return True
28             if len(stack) == 0:
29                 depth = 0
30     return False

```

H Ontology

The summary of our proposed ontological categories is shown in Table 7.

I Ontology of Perturbations

I.1 Principles behind Ontology

Consider a maths question:

Question: John has 3 boxes, each of which is externally measured as 5 inches by 6 inches by 4 inches. The boxes have walls that are 1 inch thick. What is the total inner volume of all the boxes?

We consider the following eight aspects of such questions:

- (i) **Information:** Each sentence clause that is mentioned inside the question. For example: Each box is 5 inches by 6 inches by 4 inches.
- (ii) **Query:** What is being asked by the Question that can be calculated with the given Information? For example: What is the total inner volume of all the boxes?
- (iii) **Values:** Values inside the Information. For example, 3 boxes in this particular instance.
- (iv) **ToolBox:** Mathematical concepts, formulas, and operations that are relevant to solving a specific problem. For example: Multiplication is used to calculate the volume of a rectangular prism (box) as length \times width \times height and Subtraction is used to adjust the external dimensions to account for the wall thickness.
- (v) **Mathematical Structure:** Chain of thought and problem-solving strategies that outline how the 'Tools' in Toolbox are organized to transition from the given data to the final answer. For example, to solve the question above: first, Subtract the thickness of the walls; second, calculate the volume of one box; third, multiply the volume of one box by the number of boxes.
- (vi) **Query Representation:** The Format of how Information and Values are presented. For example: the sequence of Information presented.
- (vii) **Final Answer:** Final answer to the Query. For example: 72
- (viii) **Answer Representation:** The Format of the answer presented.

In a similar vein, consider a coding question:

```

1 def greatest_common_divisor(a: int, b: int) -> int:
2     """ Return the greatest common divisor of two integers a and b
3     Example:
4     >>> greatest_common_divisor(3, 5)
5         1
6     >>> greatest_common_divisor(25, 15)
7         5
8     """

```

We can decompose the coding question into the following aspects:

- (i) **Question Header:** The name of the function, in the case above, `greatest_common_divisor`
- (ii) **Docstring:** Defines the requirement for the final output. For example, Return the greatest common divisor of two integers
- (iii) **Values:** The type and structure of input arguments. In the above example, `a` (integer type) and `b` (integer type)
- (iv) **Examples:** Demonstrations of how the function is used. In the case above,

```

1 """
2 >>> greatest\_common\_divisor(3, 5)
3     1
4 >>> greatest\_common\_divisor(25, 15)
5     5
6 """

```

- (v) **Toolbox:** Libraries and operations that can be used to achieve a function.
- (vi) **Code Structure:** Sequence of steps of code to fulfill the requirement specified in **Docstring**
- (vii) **Question Representation:** Format of how the **Question header** and **Docstring** is presented
- (viii) **Answer Representation:** Format of how the **Code Structure** is presented.

The perturbations in the ontology we introduce operate on these eight aspects of a maths or coding question. Each perturbation changes only one or two aspects of the original question.

We broadly group these perturbations into two main categories: *Structural Perturbation* and *Representational Perturbation*. *Structural Perturbations* generate new questions by modifying the specific targeted aspects of inherent logic, framework, or concepts in the original question. *Structural Perturbation* is further categorized into *Logic Alteration* and *Concept Analysis*. *Logic-Alteration* perturbations changes the logic underpinning a problem through addition or removal of information, or it changes the reasoning framework of the original problem. The *Concept Analysis* questions, however, examines the underlying concepts and principles of the problem. Rather than solving a specific problem, these questions focus on analyzing the process of problem solving, and how it get the solutions, which may require a deeper understanding of the question and problem solving framework. Details and examples for each of these perturbation types are presented below.

Unlike *Structural Perturbations*, *Representational Perturbations* retain the logical structure of the original solution, only to exclusively change the representation or encoding of the information present in the question or in the answer. In our ontology, *Representational Perturbation* has only two manifestations, *Format Change*, which directly alters the representation of the questions and answers. *Format Constraint*, which add constraint that indirectly alters the format of the answers. More details and examples are below.

For each of the above broad perturbation types, we further define many dimensions of perturbations. We apply specific methods to introduce variations or *perturbations* to the questions along these dimensions. Each dimension can further manifest in various ways that correspond to some method of perturbation. For example, a dimension such as “simplify question” can be realized in different ways for the “logic alteration” perturbation type. These perturbations can affect the difficulty level of the questions, making them either more challenging or simpler. Additionally, some perturbations may result in questions that do not have a definitive answer.

I.2 Logic Alteration

This category groups all the perturbations that have a definitive final answer. The final answer can be in the format of a value (Math) or code (HumanEval) (for dimension “Question Simplification”, “Reasoning Adjustment”, “Computation Adjustment”) or a mathematical expression (Math) or Natural Language (Code) (for dimension “Symbolic Reasoning”). For logic alteration questions, if the final answer is normalized to the most simplified form. The generated answer can be deemed correct only if it can also normalize to the same form.

(i) **Granularity Adjustment:** This dimension aims to make the question easier to solve. It can achieve this by using four ways:

G1. Remove Constraint: Remove one piece of constraint that make the question easier to solve

Remove Constraint (Math): Delete one piece of **information** from the original question that does not make the question unsolvable. The aim is to simplify the question. Example:

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. What is the total volume of all 3 boxes?

Remove Constraint (Code): Simplify the coding requirement by removing one constraint or transformation in the **Docstring**

Generate a python function that fulfills the requirement in docstring and examples usages below.

Changed from G.1:

```
1 def change_case(string: str) -> str:
2
3     """For a given string, convert all uppercase characters to lowercase.
4
5     >>> change_case('Hello')
6     'hello'
7     """
```

G2. Partial Solution: The answer only need to solve parts of the original question

Median Inquiry: Change the original **query** to ask one of the intermediate values that is used to solve for the final answer of the original query. The aim is to simplify the question. Example:

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the inner volume of one box?

Helper Function: Provide a helper function alongside the coding question that achieves partial function in **Code Structure**

Changed from Appendix G.1: Generate a python function that fulfills the requirement in docstring and examples usages below. You should complete the function using helper function.

```
1 def helper_function(char: str) -> str:
2     """Checks if a given character is uppercase or lowercase, and flips its case
3     ."""
4
5     if char.isupper():
6         return char.lower()
7     elif char.islower():
8         return char.upper()
9     else:
10        return char
11
12 def flip_case(string: str) -> str:
13
14     """For a given string, flip lowercase characters to uppercase and uppercase
15     to lowercase by using the helper function above to achieve the requirement
16     >>> flip_case('Hello')
17     'hELLO'
18     """
```

- G3. Solution Plan:** Besides the original question, provide a high level plan of how the question should be answered, the solution will only need to execute the abstract plan.

Solution Plan (Math): Provide the original question along with its **mathematical structure** (problem strategy) to the question, ask the model to solve the question by following the strategy.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes? Follow this plan to solve the question: [#Solution Plan#] Subtract the thickness of the walls from each dimension of the box to get the inner dimensions. Determine the width, height, and depth of the inner box. Calculate the inner volume of one box by multiplying the width, height, and depth. Calculate the total inner volume by multiplying the inner volume of one box by the number of boxes.

Solution Plan (Code): Provide the high level plan of how the code need to be written along with the question. **Changed from G.1:** Generate a python function that fulfills the requirement in docstring and examples usages below. You should follow the solution plan when solving the problem.

```
1 def flip_case(string: str) -> str:
2     """
3     Inverts the case of each character in the provided string.
4
5     This function takes a string as an argument and returns a new string with
6     each character's case inverted.
7     Uppercase letters are converted to lowercase, and lowercase letters are
8     converted to uppercase.
9
10    Solution Plan:
11    1. Create a result variable to hold the updated string.
12    2. Iterate through each character in the string.
13    3. Check if the character is uppercase; if so, convert it to lowercase and
14    add it to the result.
15    4. If the character is lowercase, convert it to uppercase and add it to the
16    result.
17    5. After iterating through all characters, return the result.
18    """
```

- G4. Detail Expansion:** Besides the original question, provide a few key important details or explanations without which is hard to solve the question.

Detail Elaboration: Provide original question along with the **toolbox** (commonsense knowledge) to solve the question.

Changed from F.2:

John has 3 boxes. Each box has outside dimensions of 5 inches by 6 inches by 4 inches. The walls of each box are 1 inch thick, uniformly throughout each face of the box, thereby reducing the inner dimensions of each box. The material of the boxes is uniformly distributed and does not bulge or cave in thereby affecting the inner volume. There are no internal structures or partitions inside the boxes that could further reduce the inner volume. What is the total inner volume of all 3 boxes?

Example Detail: Besides providing the input and output of each **example**, it also provide a step by step explanation of how the input is transformed to the output. **Changed from G.3:** Generate a python function that fulfills the requirement in docstring and examples usages below.

```
1 def derivative(xs: list):
2     """ xs represent coefficients of a polynomial.
3     xs[0] + xs[1] * x + xs[2] * x^2 + ....
4     Return derivative of this polynomial in the same form.
5
6     >>> derivative([3, 1, 2, 4, 5]) calculates the derivative as [1*1, 2*2, 3*4,
7     4*5] resulting in [1, 4, 12, 20].
8
9     >>> derivative([1, 2, 3]) calculates the derivative as [1*2, 2*3] resulting
10    in [2, 6].
11    """
```

(ii) **Reasoning Adjustment:** This dimension targets to partially change the logical structure of the original problem. It can be achieved through eight ways:

G5. Add Restriction: Add a new piece of condition or requirement to the answer of the question.

Restrict Question: Adding a new piece of **information** that serves as a constraint or modifier on the query. Example:

Changed from F.2:

John has 3 boxes, each of which is externally measured as 5 inches by 6 inches by 4 inches. The boxes have walls that are 1 inch thick. There is also an added wooden board divider in the middle across the smallest dimension which is also 1 inch thick. What is the total inner volume of all the boxes?

Restrict Requirement: Add a piece of information that serves as a constraint or modifier on the function.

Changed from G.1

```
1 def flip_case(string: str, index: int) -> str:
2
3     """For a given string, flip lowercase characters to uppercase and uppercase
4     to lowercase. Only flip the case for characters at indices which are
5     multiples of the provided index.
6     Note: If the index provided is 2, only the characters at the 2nd, 4th, 6th
7     positions and so on will have their cases flipped.
8
9     >>> flip_case('Hello', 2)
10    'HeLlO'
11    """
```

G6. Subsequent Question: Adding an additional query or requirement based on the answer of the original question.

Further Question: Adding an additional **query** that will need extra steps of calculation based on the final answer of the original query.

Changed from F.2: John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes? **If John wants to entirely fill these boxes with small cubes each measuring 0.5 inches on all sides, then how many cubes will he need?**

Further Requirement: Adding an additional requirement of transformation based on the output of the original function.

```
1 def flip_case_count(string: str) -> Tuple[str, int]:
2
3     """
4     For a given string, flip lowercase characters to uppercase and uppercase to
5     lowercase. Additionally, return the number of case flips performed.
6
7     >>> flip_case_count('Hello')
8     ('hELLO', 5)
9     """
```

G7. Concurrent Question: Adding an additional query or requirement that is independent from the original question.

Parallel Question: Adding an additional **query** along with the original query based on the information given in the question, the added **query** should inquiry a value that is irrelevant of the original answer. Example:

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes? **What is the total volume of the material used to build the boxes?**

Parallel Requirement: Adding an additional requirement in **Docstring** that does not rely on the output of the original question.

Changed from G.1:

```
1 def flip_case_and_count(string: str) -> Tuple[str, int]:
2
3     """For a given string, not only should you flip lowercase characters to
4     uppercase and uppercase to lowercase. You should also output another Title
5     case where only the first letter of each word is capitalized"""
6
7     """>>> flip_case_and_count('Hello')
8     ('hELLO', 'Hello')
9     """
```

G8. Change Question: Change the current query or requirement to a different but similar one based on the existing information provided inside the question.

Change Query: Change the **query** to ask for another value that requires more computation based on the information given in the question.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total **outer volume** of all 3 boxes?

Change Docstring: Change the **Docstring** to another requirement based on the input given in the question header.

Changed from G.3:

```
1 def calc_derivative(xs: list):
2
3     """ xs represent coefficients of a polynomial.
4     xs[0] * (exp(x))^0 + xs[1] * (exp(x))^1 + xs[2] * (exp(x))^2 + ....
5     Return derivative of this polynomial in the same form.
6     >>> derivative([3, 1, 2, 4, 5])
7     [1, 4, 12, 20]
8     >>> derivative([1, 2, 3])
9     [2, 6]
10    """
```

G9. Info Recombination: Combine the fundamental concepts or frameworks from another question with the original question.

Info Recombination (Math): Graft **mathematical structure** from another question and combine with the original question.

Changed from F.4:

Vicki and James are planning an event at their high school that combines a pop singing concert and dance events. The whole event will be 2 hours long. Vicki is allowing each musical group 2 minutes to get on stage, 6 minutes to perform, and then 2 minutes to exit the stage. James will also perform two solo dance routines, each lasting five minutes. Considering a 10-minute intermission during the show, how many musical groups can perform at the concert?

Info Recombination (Code): Merge the requirement from another coding question with existing question. **Changed from G.1:**

```
1 def flip_case_and_odd_sum(string: str) -> tuple:
2     """
3     Given a string, flip lowercase characters to uppercase and uppercase to
4     lowercase.
5     Also return the odd letters that are in even positions of the original
6     string.
7     string Index starts from 0, alphabet index start from 1. Aa is 1, Bb is 2..
8     Examples:
9     >>> flip_case_and_odd_sum('Hello')
10    ('hELLO', 'o')
11    """
```

- G10. Domain Knowledge:** Introduce a specific knowledge in math or code and merge it with the question.
Theoretical Challenge: Incorporate a specific theorem into the question so that perturbed question requires a new **toolbox** to solve.

Changed from F.2:

John has an infinite number of boxes numbered as first, second, third, and so on. The first box is 5 inches by 6 inches by 7 inches in size. Starting from the second box each box is half the size of the previous box in each dimension. What is the total volume all the boxes combined?

Code Import: The requirement requires to use a specific python library to solve the problem.

Changed from G.1: Rewrite the function below to take in batch input parameters and use the multicore cpu for efficiency.

- G11. Complex Reality:** Add an aspect of complexity in the real world scenario.

Value Probability: Introduce concept of uncertainty to deterministic values and calculate the estimation. The perturbed question will require **toolbox** (knowledge) of probability.

Changed from F.1:

A merchant wants to make a choice of purchase between 2 purchase plans: jewelry worth \$5,000 or electronic gadgets worth \$8,000. His financial advisor speculates that the jewelry market has a 70% chance to go up 2.5% and a 30% chance to remain the same, while the electronic gadgets market will rise 1.2% within the same month. If the merchant is looking to maximize profit at the end of this month by making a choice, how much estimated profit would this be?

Example Boundary: Add boundary examples along with the existing **examples**. The boundary examples contains input that does not met requirement specified in the docstring. **Changed from G.3:** Write a function to fulfill the requirement and all the examples inside the docstring

```

1 def derivative(xs: list):
2
3     """ xs represent coefficients of a polynomial.
4     xs[0] + xs[1] * x + xs[2] * x^2 + ....
5     Return derivative of this polynomial in the same form. The solution should
6     pass all the test cases specified below
7
8     # Regular case
9     >>> derivative([3, 1, 2, 4, 5])
10    [1, 4, 12, 20]
11    # Smaller case
12    >>> derivative([1, 2, 3])
13    [2, 6]
14    # Special case with empty list
15    >>> derivative([])
16    []
17    # Boundary case, the shortest polynomial
18    >>> derivative([1])
19    [0]
20    # Boundary case, all-zero polynomial
21    >>> derivative([0.0, 0.0, 0.0])
22    [0, 0]
23    """

```

- G12. General Solution:** Provide the solution in a more general scenario.

Code Implementation: Develop a code function to solve the question in general.

Changed from F.2:

Original Examples # Can you write a Python code to find out what is the total inner volume of all 3 boxes?

Higher Order: Write a higher order function that can solve the coding problem in general.

Changed from G.2

```

1 def greatest_common_divisor(numbers: list[int]) -> int:
2     """

```



```

3     Calculates the greatest common divisor (GCD) of a list of integers.
4     Returns the GCD as an integer.
5
6     Examples:
7     - For numbers = [20, 40, 60], the function returns 20.
8     - For numbers = [35, 14], the function returns 7.
9     """

```

(iii) **Computation Adjustment:** While retaining the **Logical Structure**, this type aims to change one single reasoning step of the original question.

- G13. Computation Demand:** Change the value to complex values that put a high demand on computation.
Value Big: Significantly increasing the magnitude of values that pose a challenge for calculations.
Changed from F.2:

John has 3000 boxes. Each box is 500 inches by 600 inches by 400 inches. The walls are 100 inches thick. What is the total inner volume of all the boxes?

Generalize Parameter: Extend the current parameter into different python object types

Changed from G.2:

```

1 def find_common_divisor(value1: Union[int, float, str], value2: Union[int, float
2     , str]) -> float:
3     """
4     Takes two values (int, float, or float in string format) and finds the
5     largest float that divides both into integers.
6     Inputs can be a mix of types. Returns the divisor as a float.
7
8     Examples:
9     print(find_common_divisor("0.5", 1)) # 0.5
10    print(find_common_divisor(0.25, "1.25")) # 0.25
11    """

```

- G14. Change Value:** Change the content of the value to a different one.
Change Subject: If there are multiple mentions in the question, Exchange **values** of names or references in the question.
Changed from F.5:

Together David, Bodhi, and Lily collected 43 insects. **David** found 7 more than **Bodhi**. **Bodhi** found half of what **Lily** found. How many insects did Lily find?

Parameter Content: Change the format or meaning of the input parameter.

Changed from G.3:

```

1 def derivative(polynomial: str):
2
3     """ 'polynomial' is a string that stands for polynomial for form
4     coefficients_0 + coefficients_1 * x + coefficients_2 * x^2 + ....
5     This function will return the derivative of the aforementioned polynomial in
6     the same format.
7
8     >>> derivative('3 +1x + 2x^2 + 4x^3 + 5x^4')
9     '1 + 4x + 12x^2 + 20x^3'
10    >>> derivative('1 - 2x + 3x^2')
11    '-2 + 6x'
12    """

```

- G15. Change Operation:** Change one operation regarding how the **Values** are processed.
Change Calculation: Change no more than 3 words in original question so that the **toolbox** (mathematical operations) involved in the calculation are changed.
Changed from F.2:

John has 3 boxes. The inner dimension of each box is 3 inches by 4 inches by 2 inches. The walls are 0.5 inches thick. What is the total outer volume of all 3 boxes?

Variable Type: Change the python object type of the original parameter while keep its content the same, also specify the return variable to be in a specific type.

Changed from G.3:

```
1 def derivative(xs: list[str]) -> list[str]:
2
3     """ xs represent coefficients of a polynomial.
4     xs[0] + xs[1] * x + xs[2] * x^2 + ...
5     Return derivative of this polynomial in the same form.
6     """
```

(iv) **Formulation Adjustment:** This dimension tests the abstract reasoning ability of LLMs under the same logical structure of the original question. This dimension focuses on solving the general version of the original reasoning problem, rather than focusing on to get a standard solution. For math, We change the context to include one or more symbolic variables to replace its original **values**.

G16. Symbolic Response: Use logic to infer the final output after a sequence of steps.

Variable Response: Replace one **value** inside the question with a variable and answer with the variable included.

Changed from F.2:

John has X boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all the boxes as a function of X?

Code Execution: Given Docstring requirement, and specific input parameter, find the output for the function without writing any code. **Changed from G.1:** Find the output of the following function description, if the input is:string = "Hello World&7"

```
1 def flip_case(string: str) -> str:
2     """For a given string, flip lowercase characters to uppercase and uppercase
   to lowercase."""
```

G17. Value Relationship: Identify the relationship between input values or parameters if the output or the final answer is given.

Variable Relationship: Replace a pair of **values** inside the question with variables. After answering the original question, the variable forms a relationship. Query that relationship.

Changed from F.2:

John has X boxes. Each box is Y inches by 6 inches by 4 inches. The walls are 1 inch thick. If the total inner volume of all the boxes is 72 cubic inches, then find the equation that relates X and Y?

Parameter Relationship: Given the output of the function, categorize the possible groups of inputs parameters into the question. **Changed from G.2:** If the below program output integer 7. What is the relationship between a and b

```
1 def function(a: int, b: int) -> int:
2     while b:
3         a, b = b, a % b
4     return a
```

G18. Variable Group: Change a group of several input values or parameters to variables.

Variable Scaling: After answering the question, change the **query** to: if certain factual numbers in the question is scaled up by x, how will the final answer change?

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. Now, the number of boxes, the box outer dimensions, and the wall thickness are all scaled up by a

factor of X. What is the total inner volume of all the boxes as a function of X?

Variable Substitution: Change one or more variables inside the docstring to input parameters.

Changed from G.1:

```
1 def flip_case(string: str, specific_value: str) -> str:
2
3     """ """ For a given string and specific value, flip the specific value from
4     lowercase to uppercase or uppercase to lowercase. The function will only
5     flip the case of the specific value in the string.
6     >>> flip_case('Hello', 'h')
7     'hello'
8     """
```

G19. Backward Reasoning: Reverse the reasoning process, reason from how to reach input from output.

Variable Adaptation: If the answer to the question add or subtract by a certain amount x, pick one value inside the **Information** and ask how it should change if other values are kept the same.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. If the total inner volume of all 3 boxes increases by a certain variable X, how should the thickness of the walls adjust correspondingly if the number of boxes and the external dimensions of each box stay the same? Write the answer as a function of X.

Reverse Engineering: Change the **Docstring, Function Header, and Examples** to find the function that can reverse engineer the original function. Specifically, mapping the output back to its original input. **Changed from G.1:** Create a function that reverses the following function's process, effectively transforming its output back into the original input

```
1 def function(string: str) -> str:
2     return string.swapcase()
```

G20. Counterfactual: What would the outcome be if X had happened instead of Y, given the same initial conditions and context.

WhatIf Question: First mask some number of values inside the question and answer the original question. What if we change one value inside the question, how will the final answer change? (The final answer should not have variables included as the masked value could be solved given the final answer.)

Changed from F.2:

John has 3 boxes. Each box is 5 inches in width by 6 inches in length and a few inches in height. The walls are 1 inch thick. The total inner volume of all the boxes combined is 72 cubic inches. Now, if the thickness of the walls is half of its original thickness, then what will be the total inner volume?

WhatIf Code: WhatIf the **code structure** or **input value** is changed, and some condition is masked.

Changed from G.1: Find the output of the 'changed_function', if the input is the same.

```
1 We know that if we input masked_input to the `original_function`, the output is
2 following:
3 >>> original_function(masked_input)
4 'HELLO'
5 Here is the `original_function`:
6 def original_function(string: str) -> str:
7     return string.swapcase()
8
9 Here is the `changed_function`:
10 def changed_function(string: str) -> str:
11     return string.swapcase()[::-1]
12
13 What will be the output for `changed_function(masked_input)`"
```

G21. Solve Value: Mask one variable's value inside question, given answer, infer the masked value.

Solve X: Replace one value inside the question with X and solve for X.

Changed from F.2:

John has X boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. If the total inner volume of all 3 boxes is 72 cubic inches what is the value for X?

Solve Input: Determine the input value of the function, based on the known output value.

Changed from G.1 What is input to the following function, if the output is: "hELLO 9"

```
1 def function(string: str) -> str:  
2     return string.swapcase()
```

G22. Identify Range: Find what are possible constraint on the values.

Variable Range: Replace one **value** with variable, and change the **query** to find the possible range of values based on the question.

Changed from F.2:

John has 3 boxes. Each box is X inches by 6 inches by 4 inches. The walls are 1 inch thick. Suppose we want to find out the total inner volume of all the boxes. What are the possible ranges of values of variable X based on the given information?

Parameter Range: Identify what are the constraint on the input parameter, or what is the range of output parameter if input parameter is constraint to take certain value.

Changed from G.3: If all the item inside the input list is smaller than 1, what are the constraints on the output from this function below?

```
1 def function(xs: list):  
2     return [(i * x) for i, x in enumerate(xs)][1:]
```

I.3 Concept Analysis

This perturbation type encompasses questions that concentrate on the model's capabilities beyond mere problem-solving accuracy. The responses to these questions should be in natural language format. Instead of just assessing whether the model can correctly predict answers to new questions, we aim to examine the depth of knowledge the models possess and understanding of important concepts and rationales in the process of solving the original questions. Essentially, we are asking: *Does the model predict correctly because it truly understands the question?* To test this, we observe how the model behaves in different or unusual scenarios that are not typically presented in standard questions.

(i) **Question Understanding:** This dimension examines how model decompose, interpret and analyze the information inside the question.

G23. Inherent Premise: Identify the underlying premise of the question.

Identify Assumption: Identify one hidden commonsense assumption in the question that requires the answer to be answerable.

Changed from F.3:

You do not need to solve the question below, just identify one important hidden assumption that is required for the question to be answerable. 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?

Test Case: List different boundary test cases that is valid for the input of the question. **Changed from G.1:** Provide input parameters for the test cases of the specified coding problem. These parameters should encompass boundary conditions within the scope defined by the function's requirements specification, and avoid scenarios that fall outside of these requirements.

```
1 def flip_case(string: str) -> str:  
2     """For a given string, flip lowercase characters to uppercase and uppercase  
   to lowercase.  
3     """
```

G24. Complete Missing: Fulfill the missing information in the question by analyze how the information is structured and presented inside the question.

Info Sufficiency: Mask or delete an important piece of **information** and ask what additional information is needed to make the question answerable.

Changed from F.2:

John owns 3 boxes, each measuring 5 inches by 6 inches by 4 inches. Each box also had inner walls with certain non-zero thicknesses. Suppose you want to find out the total inner volume of all the boxes. What information is missing to calculate that?

Incomplete Answer: Given the question, mask partial answer of the original, the model need to infer the missing lines based on the context. **Changed from G.1:** Complete the function below by predicting what is inside the masked code paragraph

```
1 def flip_case(string: str) -> str:
2     """For a given string, flip lowercase characters to uppercase and uppercase
3     to lowercase.
4     >>> flip_case('Hello')
5     'hELLO'
6     """
7     [masked code paragraph]
8     if char.isupper():
9         result += char.lower()
10    else:
11        result += char.upper()
12    return result
```

G25. Question Formulation: Formulate the question based on its answer.

Question Formulation - (Math): Formulate a **question** to the chain of thought **gold answer**.

Changed from F.1:

Formulate a math application **question** that requires the following **mathematical structure** (calculations): $5000 \cdot (2.5/100) = \$125$ $8000 \cdot (1.2/100) = \$96$ $\$125 > \96 $\$125$ Math Question: Ask potential structures of math application.

Question Formulation - (Code): Formulate a concise coding requirement by looking at the function code.

Changed from G.2: Write a concise code description for the following code of its functionality no more than 1 sentence.

```
1 def function(a,b):
2     while b:
3         a, b = b, a % b
4     return a
```

G26. Add Misinformation: Add a piece of distracting information that can mislead the answer.

Introduce Distraction: Add a Potentially Distracting **information** that will not affect the answer to the question.

Changed from F.1:

A merchant is considering a decision between the following purchase plans: jewelry with a value of \$5,000, a trip to Europe costing \$7,000, or electronic gadgets worth \$8,000. His financial advisor predicts that the jewelry market will increase by 2.5%, the travel market will stay relatively stable with little to no change, and the electronic gadgets market will rise by 1.2% within the same month. He recently also came into an inheritance of \$20,000 that he doesn't need to use right away. If the merchant's goal is to maximize profit at the end of this month by making a purchase choice, how much profit would this be?

Introduce Bias: Change the python header to describe another function requirement, and change all the examples demonstrations bias towards a specific output **Changed from G.1**

```
1 def uppercase(string: str) -> str:
```



```

2 """For a given string, flip lowercase characters to uppercase and uppercase
3 to lowercase.
4 >>> flip_case('hello')
5 'HELLO'
6 """

```

(ii) **Solution Understanding:** This dimension focuses on the problem-solving process to get to the final answer and how to optimize it.

G27. Optimize Solution: Assess whether the current state is optimal or if improvements are necessary.

Info Necessity: Check If there is redundant **information** given in the question, if yes, identify the redundant information.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. Suppose we want to find out the total inner volume of all 3 boxes. To solve this math question, is there a way to determine the total inner volume of all 3 boxes without calculating the inner volume of one box?

Reduce Complexity: Assess whether the complexity of the current code be further reduced.

Changed from G.3: Optimize the code below to more efficiently achieve the same requirement specified in the docstring

```

1 def derivative_polynomial(coefficients, derivative=None, index=0):
2     """
3     This function calculates the derivative of a polynomial using recursion.
4     coefficients: List of coefficients of the polynomial.
5     derivative: List to store the coefficients of the derivative. Initially None
6     .
7     index: Current index in the coefficients list.
8
9     The base case of the recursion is when the index is equal to the length of
10    the coefficients list.
11    """
12    if index > 0:
13        derivative_coefficient = index * coefficients[index]
14        derivative.append(derivative_coefficient)
15    return derivative_polynomial(coefficients, derivative, index + 1)

```

G28. Step Functionality: Whether there are alternative answers that follow the constraint.

Step Necessity: Whether there are any alternative solutions **reasoning steps** without calculating an specific intermediate value.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. Suppose we want to find out the total inner volume of all 3 boxes. To solve this math question, is there a way to determine the total inner volume of all 3 boxes without calculating the inner volume of one box?

Step Necessity: Provide one line of code inside the Python function, and explain the functionality of that line of code in the context of the whole solution.

Changed from G.2: Explain what is the the line below the comment functionality?

```

1 def greatest_common_divisor(a: int, b: int) -> int:
2
3     """ Return a greatest common divisor of two integers a and b
4     >>> greatest_common_divisor(3, 5)
5     1
6     >>> greatest_common_divisor(25, 15)
7     5
8     """
9     while b:
10        a, b = b, a % b
11    # What is the functionality of `abs()`
12    return abs(a)

```

- G29. Theoretical Basis:** Identify the theory or principles in solving the question in general.
Theoretical Basis (Math): Identify the underlying arithmetic or algebraic rules (**toolbox**) that govern the solution to the question.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. Assume you want to find out the total inner volume of all 3 boxes. Can you identify one underlying mathematical theory which is required to do that?

Theoretical Basis (Code): Request explanation on essential python concepts required to solve the question, for example, related to python objects and programming skills.

Changed from G.1: Please describe to me in simple terms, assuming I have no knowledge of programming. Your task isn't to solve the coding problem itself, but rather to identify the programming concepts in Python that would be necessary to address the problem presented below.

```
1 def flip_case(string: str) -> str:
2     """For a given string, flip lowercase characters to uppercase and uppercase
3     to lowercase.
4     >>> flip_case('Hello')
5     'hELLO'
6     """
```

- G30. Cost Analysis:** Analyze the computational cost regarding the solution.
Solution Efficiency: Compare two solution plans on solving the question and evaluate which one uses less computation.

Changed from F.2:

Evaluate which solution plan is more efficient in solving the question?
Plan 1: Calculate the volume of the outer dimensions for one box, calculate the volume of the material used for the walls for one box, subtract the latter from the former to find the inner volume of one box, and then multiply this by 3 for all boxes.
Plan 2: Calculate the inner dimensions of a single box by subtracting twice the thickness of the walls from each outer dimension, then find the volume of this inner space and multiply by 3 for all boxes.

Code Complexity: Analyze the time complexity and space complexity of the provided code solution.

Changed from G.1 Analyze the time and space complexity regarding to input parameter string of the following function.

```
1 def flip_case(string: str) -> str:
2     """For a given string, flip lowercase characters to uppercase and uppercase
3     to lowercase.
4     >>> flip_case('Hello')
5     'hELLO'
6     """
```

(iii) **Critical Thinking:** In this dimension, deliberate errors are introduced into the question or in a provided example answer. The purpose is to see if the LLM can identify and rectify these errors. This tests the LLM's error detection capabilities, which is crucial for reliability in practical applications.

- G31. Seek Clarification:** The question requires to be clarified first before answering.
Introduce Ambiguity: Introduce Ambiguity to the question implicitly by changing the original **information**, so that the question cannot be solved without clarification.

Changed from F.2:

John has three 5x6x4 inch boxes. A particular side of each box have 1 inch thick walls. What does the total inner capacity of these boxes amount to?

Example Requirement: Remove the coding requirement in the docstring, instead only provide examples as a coding requirement. The provided examples will define and demonstrate the expected behavior in various scenarios. **Changed from G.1:** Begin by analyzing the function's behavior specified in the docstring to understand its pattern, and then proceed to code the function accordingly.

```

1 def flip_case(string: str) -> str:
2     """
3     function('Hello') == 'hELLO'
4     function('Python 3.8') == 'pYTHON 3.8'
5     function('123abcXYZ') == '123ABCxyz'
6     function('MixedCASE123') == 'mIXEDcase123'
7     function('ALLUPPERCASE') == 'alluppercase'
8     """

```

G32. Conditional Analysis: Based on different possible situations of the question, the answer should separately presented.

Discuss Separately: Introduce new **information** containing variables or conditions that require the answer to be discussed separately based on conditions or variables.

Changed from F.1:

A merchant wants to make a choice of purchase between 2 purchase plans: jewelry worth \$5,000 or electronic gadgets worth \$8,000. His financial advisor speculates that the jewelry market will go up $x\%$ while the electronic gadgets market will rise 1.2% within the same month. If the merchant is looking to maximize profit at the end of this month by making a choice, how much profit would this be?

Incomplete requirement: Left some condition unspecified in the docstring. **Changed from G.1:**

```

1 def flip_case(ch: str) -> str:
2
3     """For a given string, all the letters inside the string should be changed.
4     flip lowercase characters to uppercase.
5     >>> flip_case('h')
6     'H'
7     """

```

G33. Conflicting Information: Introduce a new piece of information that is conflicting with existing information. This will make the question unanswerable, so the if the LLM can spot the error without mentioning.

Introduce Contradiction: Add a piece of contradicting **information** to the question and check if LLM can spot the problem.

Changed from F.2:

John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. Each box is also 8 inches in width. What is the total inner volume of all 3 boxes?

Wrong Example: Include an example that is conflicting with the requirement specified in the docstring.

Changed from G.1:

```

1 def flip_case(string: str) -> str:
2     """ """For a given string, flip lowercase characters to uppercase and
3     uppercase to lowercase.
4     >>> flip_case('Hello')
5     'hello'
6     """

```

G34. Surface Error: Introduce an obvious error that can be spotted without reasoning.

Value Uncommon: Change the **values** so that it seems weird or unusual by commonsense knowledge standards.

Changed from F.2:

Can you spot anything unusual in the following question? John has 3 boxes. Each box measures 50000 miles by 60000 miles by 40000 miles. The walls of the boxes are 100 miles thick. What is the total inner volume of all 3 boxes?

Runtime Error: Introduce a piece of error that will cause a runtime error or syntax error in python.

Changed from G.1: Debug the error in the following code

```

1 def flip_case(string, str) -> str:

```

```

2     """For a given string, flip lowercase characters to uppercase and uppercase
3     to lowercase.
4     >>> flip_case('Hello')
5     'hELLO'
6     """
7     return string.swapcase()

```

G35. Hidden Error: Introduce a hidden error that need logical reasoning to spot.

Value Error: Change the **values** so that the question does not make sense.

Changed from F.4:

Vicki is planning a pop concert at her high school. The show will be 2 minutes. She is allowing each group 2 hours to get on stage, 6 hours to perform, and then 2 hours to exit the stage. If she allows a 10-hour intermission, how many groups can perform in the concert?

Logical Error: Introduce the change in the code that will cause a Logical Error in Python. **Changed from G.1**

```

1 def flip_case(string: str) -> str:
2     """For a given string, flip lowercase characters to uppercase and uppercase
3     to lowercase.
4     >>> flip_case('Hello')
5     'hELLO'
6     """
7     string = list(string.swapcase())
8     return string

```

I.4 Question Format

This dimension is inspired by metamath (Yu et al., 2023b). It involves changing the **question representation** by modifying the question encoding or specify the representation of the answer in different ways while keeping the underlying logical structure and conceptual framework of the original question intact. The objective is to verify whether the LLM can still provide correct answers even when the format or presentation of the question changes. This tests the model’s ability to reason irrespective of how it’s presented. It also tests models’ instruction following ability where the **answer representation** must follow a certain format.

(i) **Format Change:**

G36. Setting Rephrase: Rephrase the question in another setting.

Change Setting: Rephrase by changing the application setting and values inside the information, while keeping the core mathematical structure intact.

Changed from F.2:

Maria has 4 cuboids. Each cuboid is 7 feet by 9 feet by 6 feet. The walls are 2 feet thick. What is the total volume of all the cuboids?

Realworld Usecase: Frame the requirement in docstring into a problem that will happen in a realworld scenario. **Changed from G.1:**

```

1 def switch_text_case(text: str) -> str:
2     """
3     Imagine you're working on a document and you've mistaken the case in the
4     text you write. You wrote all the lower case letters in uppercase and vice
5     versa, suppose you want to correct all of them using python.
6     """

```

G37. Change Sequence: Change the order of the information and names of the variables that is originally presented in the question.

Change Sequence: Change the sequence of information given in the question without affecting the solvability of the question.

Changed from F.2:

The walls of John's boxes are 1 inch thick. Each of these boxes measures 5 inches by 6 inches by 4 inches. John has 3 boxes. What is the total inner volume of all 3 boxes?

Parameter Sequence: Change the sequence of the input parameter and change the input parameter names.

Changed from G.2

```
1 def munchee_bunchee(xray: int, yoyo: int) -> int:
2
3     """ Return a common divisor that is the largest of two integers xray and
4     yoyo
5     >>> munchee_bunchee(3, 5)
6     1
7     >>> munchee_bunchee(25, 15)
8     5
9     """
```

G38. Close Format: Rewrite the sentence as a closed-format question that evaluates the correctness of possible answers.

True False: Evaluate a potentially misleading answer and check the correctness of the answer.

Changed from F.1:

A merchant wants to make a choice of purchase between 2 purchase plans: jewelry worth \$5,000 or electronic gadgets worth \$8,000. His financial advisor speculates that the jewelry market will go up 2.5% while the electronic gadgets market will rise 1.2% within the same month. If the merchant is looking to maximize profit at the end of this month by making a choice, how much profit would this be? Evaluate the correctness of this answer with respect to the above question: \$96.

True False: Check if a given code answer can solve the requirement in docstring. **Changed from G.2:** Evaluate whether the solution below is the correct solution for the coding question, True or False?

```
1 Function:
2
3 def greatest_common_divisor(a: int, b: int) -> int:
4
5     """ Return a greatest common divisor of two integers a and b
6     >>> greatest_common_divisor(3, 5)
7     1
8     >>> greatest_common_divisor(25, 15)
9     5
10    """
11
12
13 Solution:
14
15 while a:
16     a, b = a % b, a
17     return b
```

G39. Data Restructuring: Change the layout, organization of the data presented in the question.

Value Structuring: Arrange the variables inside the question in a tabular format.

Changed from F.2:

Variable	Value
a	3
b	5
c	6
d	4
e	1

John has 'a' boxes. Each box is 'b' inches by 'c' inches by 'd' inches in dimensions. The walls are 'e' inch thick. What is the total inner volume of all the 'a' boxes?

Complex Docstring: Elaborate the documentation string by exhaustively detailing more conditional pathway within the code.

Changed from G.1:

```
1 def function(string: str = None) -> str:
2     """
3     For any specified sequence of alphabetical characters, interspersed with
4     spaces, numerical digits, and various symbols, implement a sophisticated
5     transformation algorithm designed to selectively convert each alphabetical
6     character from its current case representation, either lowercase or
    uppercase, to its diametrically opposite case representation. This algorithm
    ensures that every character initially presented in lowercase is
    meticulously transmuted to uppercase, and conversely, every character
    originally in uppercase is converted to lowercase, while meticulously
    preserving the integrity and original positioning of spaces, numerical
    digits, and any other non-alphabetical symbols, leaving these elements
    unaltered within the sequence.
    >>> function('Hello')
    'hELLO'
    """
```

(ii) **Format Comparison:**

G40. Identical Problem: Check if the two question or code are identical in describing or solving the same problem.

Identical Question: If two questions requires exactly the same framework or thinking procedure to solve.

Changed from F.1:

Question 1: A merchant wants to make a choice of purchase between 2 purchase plans: jewelry worth \$5,000 or electronic gadgets worth \$8,000. His financial advisor speculates that the jewelry market will go up 2.5% while the electronic gadgets market will rise 1.2% within the same month. If the merchant is looking to maximize profit at the end of this month by making a choice, how much profit would this be? Question 2: An investor is unsure of which investment to make: gold valued at \$10,000 or stocks valued at \$15,000. His financial consultant predicts that the gold market will inflate by 3.5% while the stock market will increase by 2.2% over the next quarter. If the investor wants to achieve the highest return on his investment at the end of this quarter, how much would his initial investment be? Does Question 1 and Question 2 require identical steps to answer?

Identical Code: Are the two solutions to the question identical in terms of their functionality?

Changed from G.3 Is function_1 and function_2 identical in terms of its functionality?

```
1 Code 1:
2 def function(xs: list):
3     return [(i * x) for i, x in enumerate(xs)][1:]
4 Code 2:
5 def function(xs: list):
6     derivative = [i * xs[i] for i in range(1, len(xs))]
```

I.5 Answer Format

(i) **Answer Constraint:** This dimension add a constraint on the solution so that it should conduct reasoning under the constraint

G41. Reasoning Format: The format for the final answer should be converted in a certain way.

Binary Coded: Answer the final question in base-n.

Changed from F.2:

Answer the following question with only base-2 coded values. Question: John has 3 boxes.

Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes?

No Keyword: The solution should not use a specific python keyword. For example, “for” or “while”.

Changed from G.2: Answer the coding function below without using python keywords: "while", "for" in the solution

```
1 def greatest_common_divisor(a: int, b: int) -> int:
2
3     """ Return a greatest common divisor of two integers a and b
4     >>> greatest_common_divisor(3, 5)
5         1
6     >>> greatest_common_divisor(25, 15)
7         5
8     """
```

G42. Reasoning Style: The reasoning steps should be performed in a certain style.

X Language (Math): Give the answer in certain language from Spanish, Chinese, Bengali, English, French

Changed from F.2: Answer the following question with only Chinese language, because I do not understand English.

Question: John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes?

X Language (Code): Give the code answer in another coding language.

Changed from G.1: Answer the coding question below;

```
1 func flipCase(str string) string {
2 // flipCase takes a string and flips the case of each character: lowercase to
   uppercase and uppercase to lowercase.
3
4 }
```

G43. Alternative Answer: Find the alternative solutions to existing solution.

Alternative Answer (Math) : Give an alternative solution that is different from the standard **reasoning steps**, but arrives at the same correct final answer.

Changed from F.2:

Give an different step-by-step solution to calculate the answer to the following question. Make sure the solution is different from the solution below. Question: John has 3 boxes. Each box is 5 inches by 6 inches by 4 inches. The walls are 1 inch thick. What is the total inner volume of all 3 boxes? Solution: The walls subtract a $(1 + 1) = 2$ inches from each dimension. So, each box has a reduced width of $(5 - 2) = 3$ inches, reduced length of $(6 - 2) = 4$ inches and reduced height of $(4 - 2) = 2$ inches. So the inner volume of each box is $3 * 4 * 2 = 24$ cubic inches. The total inner volume of 3 boxes are $3 * 24 = 72$ cubic inches. Alternative Step by Step Solution:

Alternative Answer (Code): Find an alternative solution to existing coding solution. **Changed from G.1:**

```
1 Find a different solution other than:
2 def flip_case(string: str) -> str:
3
4     return string.swapcase()
```

G44. New Rule: Integrate a new rule into the original question that requires the solution follow the new rule. This type tests the model’s ability to adapt to new ruels and knowledge and use it inside the solution.

Define Rules: Define a new mathematical rule that will change how the **toolbox** (commonsense knowledge) is applied during calculation.

Changed from F.2:

In a parallel universe, John has 3 boxes. Each box has peculiar dimensions: 5 quarks by 6 quarks by 4 quarks with walls that are 1 quark thick. In this universe, the total inner volume of a box is calculated by using the Illusory Volume operation, represented as IV. The IV operation is defined as: $(\text{length} * \text{width} * \text{height}) - (\text{number_of_walls} * \text{thickness_of_each_wall})$. What is the total inner volume of all 3 boxes?

Simple Name: The generated code should only have variables names in a certain format.

Changed from G.1: Answer the coding question below and only use 6 letter word for each variable names inside the solution

```
1 def flip_case(string: str) -> str:
2     """For a given string, flip lowercase characters to uppercase and uppercase
3     to lowercase.
4     >>> flip_case('Hello')
5     'hELLO'
6     """
```

Overall, these dimensions in the “Format Change” and “Format Constraint” Domain are designed to challenge the LLMs in ways that reveal their limitations and strengths in maintaining accuracy and functionality under modified or challenging conditions.

J Evaluation Details

We used separate prompt templates for open source and close source models because close source models sometimes give the final answer directly and omit reasoning steps even if prompted with "Let's think step by step". To ensure the model performs Chain of Thought Reasoning, we use the following prompt template for GPT-4, GPT-3.5, and Gemini to generate the answer:

Solve the question step by step before giving the final answer. Do not directly give the final answer.

Question

Reasoning Step:

For Metamath, CodeLlama and Llama2-Chat, we use the following:

Below is an instruction that describes a task. Write a response that appropriately completes the request.

Instruction: Question

Response: Let's think step by step.

The temperature of GPT-4 and GPT-3.5 was set to 0.7 (the default setting in OpenAI playground) for *Concept Analysis question* and 0.1 for *Logic Alteration questions and Format Change questions*. Similarly, the temperature for Llama, ChatGPT, and Gemini were set to 0.8 and 0.1 for *Concept Analysis and Logic Alteration questions and Format Change questions*, respectively.

K Experiment Details

Prompt to incorporate original answer:

Given the original question and its answer, Solve the question that is a perturbed variant of the original question. Solve the #perturbed question# step by step before giving the final answer. Do not directly give the final answer.

#Original Question#: original question

#Original Answer#: original answer

#Perturbed Question#: perturbed question

Self Consistency Prompting: We randomly picked one question answer pair in the same category from our math dataset GSMORE and prepend it to the front of the perturbed question as a one shot demonstration. Then we use the below prompt template for Self Consistency prompting. We sample the generation three times and get the final answer by majority voting. In case of tie, we randomly pick an answer.

Given the oneshot demonstration of a question and its final answer, Solve the #question# step by step before giving the final answer. Do not directly give the final answer.

#Demonstration Question#: [demonstration question](#)
#Demonstration Final Answer#: [demonstration answer](#)

#Question#: [question](#) Reasoning Step: [Reasoning Steps] Final answer: [Final answer]

Program of Thought Prompting: We use the following prompt template for the program of thought experiments:

Instruction: You are an experienced professional skilled in using python programs to solve math related problems. Solve the question below using python programs, You will only write code blocks.

Problem: [Question](#)