

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

Anonymous ACL submission

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 maths and coding questions, (ii) a semi-automatic method to apply these perturbations, and (iii) two datasets, MORE and CORE, respectively, of perturbed maths 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.

1 Introduction

Logical reasoning in a structured and well-defined domain, such as mathematics and programming, becomes increasingly harder with the increasing presence of interspersed and diverse situations, events, and contexts formulated through natural language queries. Current state-of-the-art Large Language Models (LLM) have shown impressive performance on mathematical problems (Cobbe et al., 2021a) and reasonable performance on 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 clear 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 demanding 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 competition mathematics problems. Similar variance in LLM performance can also be observed for coding challenges (Chen et al., 2021a). Such shallow evaluations are unfit for an objective measure of the finer LLM capabilities 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 without bolstering the fundamentals. Hence, before testing the LLMs’ breadth of capabilities by delving into higher mathematics and evaluating competitive coding questions, we instead focus on depth through one fundamental question:

How robust are the capabilities of LLMs in terms of reasoning and understanding of the problem-solving process?

In this work, our goal is to provide an evaluation mechanism that provides clear insights into the robustness of the reasoning abilities of LLMs in the context of maths and coding. Following previous work towards probing language models (Ribeiro et al., 2020; Wu et al., 2023; Li et al., 2024a; Wang et al., 2024), we evaluate the robustness of LLMs’ understanding of interesting linguistic and logical structures and derive insights based on them.

Specifically, we design an adaptive dynamic evaluation benchmark through novel ontology-guided perturbations on existing problems. We introduce a novel ontology of perturbation operations that lists various changes across a diverse set of factors, which we apply to previously introduced arithmetic and coding problems. These perturbations allow

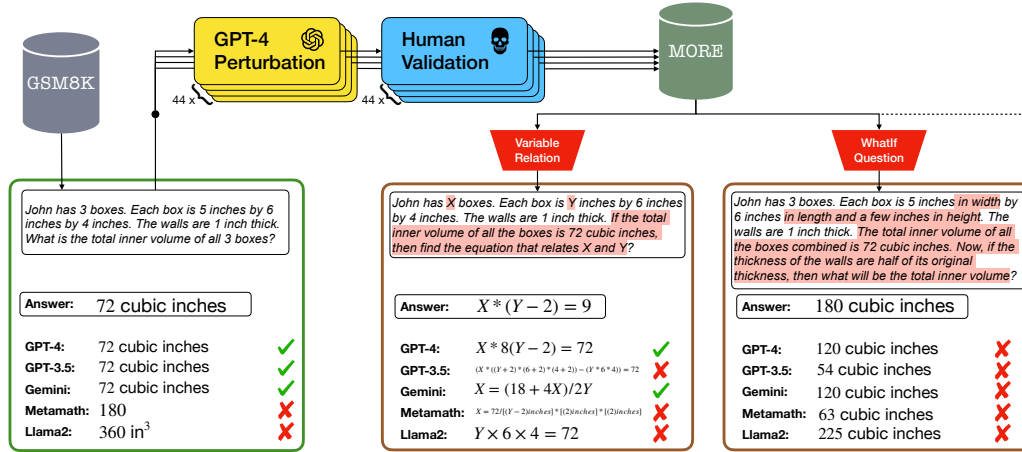


Figure 1: A semi-automated pipeline of creating MORE, from five simple questions from GSM8k. An analogous pipeline is used to create the perturbations of the coding questions from HumanEval, named CORE.

us to assess whether the model comprehends underlying concepts. For instance, while a model may correctly answer questions in a dataset like GSM8k, it might struggle when presented with a simple perturbation to the question, such as replacing numerical values in maths 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 ontology consists of 44 types of perturbations, which we apply to sample questions from GSM8K and coding questions from HumanEval, resulting in 216 and 219 perturbed questions respectively. Our evaluation of GPT-4, GPT-3.5, Metamath, Llama-code, Llama3-Instruct, and Gemini 1.5 shows that most of these models very quickly degrade under different perturbation types. Our contributions are as follows:

1. We propose a novel, extensive, and extensible ontology of perturbation operations for basic-math- and coding-based reasoning tasks expressed in natural language.
2. We present a semi-automatic method to exercise such perturbations first through GPT-4, followed by manual filtering. We generate two datasets MORE and CORE—Mathematics and Code-Oriented Robustness Evaluation, respectively—consisting of 216 maths and 219 coding questions.
3. We gain insights into the range of capabilities

and limitations on such math and coding tasks for several LLMs.

2 Related Work

A variety of datasets have been developed to assess AI reasoning capabilities across multiple domains. In causal reasoning, significant datasets include those 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 through datasets designed for 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 shown by 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 work has focused on the robustness of reasoning under various perturbations that alter reasoning question. 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 1. In the field of mathematics, contemporary works have employed techniques such as numerical or symbolic substitutions (Li et al., 2024a; Zhou et al., 2023; Meadows et al., 2023; Wang et al., 2024; Patel et al., 2021), the insertion of irrelevant distractors (Shi et al., 2023; Li et al., 2023), functional equivalence

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
MORE-CORE (Our) *	math(GSM8K), code(HumanEval)	Dynamic Ontology	GPT-4, Human (Q,A)	R L C T F S V	44

Table 1: 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

(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), or sensitivity to reasoning context (Wu et al., 2023). To our knowledge, well-established perturbation methods beyond the domain of mathematics 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) target at 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.

3 The Ontology of Perturbations

3.1 The Need for Ontology-based Perturbations

We plan to first identify a set of factors upon which the solution of a structured reasoning problem (expressed in natural language) may depend

on (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) discusses ways of perturbation, by identifying a set of factors which 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

225 have shown, the same process can be enabled to
226 perform data augmentations.

227 3.2 The Ontology

228 Extending SVAMP (Patel et al., 2021)-like pertur-
229 bations, we propose a set of high-level categories
230 that are applicable to a broad class of reasoning
231 tasks, expressed in natural language. We primarily
232 identified the following hierarchy (see Table 2):

233 **Level I: Aspect.** There are two aspects to these
234 perturbations: (i) *structural perturbation* and (ii)
235 *representational perturbation*. *Structural perturba-*
236 *tion* covers all perturbations that probe the underly-
237 ing reasoning path (or structure) in different ways,
238 by slightly varying the logic behind the question or
239 probing intermediate steps, seeking explanations.
240 *Representational perturbations* involves modifica-
241 tion of the encoding of the question or solution
242 while preserving the underlying logic of the origi-
243 nal question.

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

249 **Level III: Dimension.** This is a further refine-
250 ment that defines the exact target dimensions (the
251 WHAT) in the reasoning process (question, reason-
252 ing, computation, answer expression etc.) to which
253 the perturbations are applied.

254 **Level IV: Category.** This level captures the
255 method (the HOW) through which the higher-level
256 *Dimension* perturbation is achieved. These meth-
257 ods are domain dependent and, thus, their imple-
258 mentations vary from maths to coding problems.

259 4 Curation of MORE and CORE

260 Our objective is to assess the resilience of LLMs
261 to perturbations of maths and coding questions
262 along various dimensions. Thus, as seed datasets,
263 we use GSM8K (Cobbe et al., 2021b)—a collec-
264 tion of mathematical problems demanding rigor-
265 ous arithmetic and logical reasoning—and Hu-
266 manEval (Chen et al., 2021a) for coding. Five
267 questions¹ from GSM8K are perturbed using our
268 ontological framework (see Appendix I) to gener-
269 ate MORE. On the other hand, we sampled five

¹Maths 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

270 coding problems from HumanEval dataset (Chen
271 et al., 2021a) that were perturbed using the ontol-
272 ogy explained in Appendix I. These perturbations
273 are aimed at modifying the problems in terms of
274 complexity and representation to assess the robust-
275 ness of the LLMs to these ontological categories
276 of perturbations. Fig. 2 shows examples of three
277 perturbed questions and answers from MORE and
278 CORE. Examples and definitions of all the remain-
279 ing perturbations are present in Appendix I. We use
280 a three-staged combination of automatic generation
281 from GPT-4 (OpenAI, 2023) with human verifica-
282 tion and annotation to create MORE and CORE: (i)
283 perturbed question generation (§4.1), (ii) filtering
284 and validation of generated questions (§4.2), and
285 (iii) annotating final answers (§4.3).

286 4.1 Perturbed Question Generation

287 In the first stage, our objective is to create perturbed
288 questions from the source GSM8K/HumanEval
289 questions for each perturbation type. We write
290 prompt templates for each perturbation type and
291 fill them with a source question to create the in-
292 put prompt to GPT-4. Each template captures the
293 essence of the respective perturbation type (Ap-
294 pendix I.2, Appendix I.3, Appendix I.4) to instruct
295 GPT-4 on how to perturb the source question.

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

298 Instruction: Rewrite the original mathematical
299 context below based on the #Rewrite Require-
300 ment#.

301 Your output should only be #Rewritten Context#.

302 #Original Context#: John has 3 boxes. Each box
303 is 5 inches by 6 inches by 4 inches. The walls are
304 1 inch thick.

305 #Original Query#: What is the total inner volume
306 of all 3 boxes?

307 #Rewrite Requirement#: 1. Remove some con-
308 straints or information from the original context.
309 2. Make sure the rewritten question can still be
310 solved, but the answer is simpler.

311 #Rewritten Context#:

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

Aspect (Level I)	Target (Level II)	Dimension (Level III)	Category (Level IV)	Math	Code	
Structural Perturbation <i>Def:</i> Modification on specific aspects of logic or concepts to alter the reasoning process required to reach the answer	Logic <i>Def:</i> Modifications to the reasoning framework or logic underpinning a problem.	Granularity Adjustment <i>Def:</i> Fine-grained sub-tasks of the original question	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:</i> Target at logical structure of the original	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:</i> Target at values or entities	G13. Computation Demand G14. Change Value G15. Change Operation	Value Big Change Subject Change Calculation	Generalize Parameter Parameter Content Variable Type	
	Concept <i>Def:</i> Examination and Analysis of the underlying concepts and principles of a problem	Formulation Adjustment <i>Def:</i> Reformulate question for solution form to be an abstract expression.	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	
		Question Understanding <i>Def:</i> Interpretation of the information inside the question	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:</i> Assessment of the problem-solving processes	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	
	Representational Perturbation <i>Def:</i> Preservation of the underlying logic and conceptual framework, but modification of the encoding or representation	Critical Thinking <i>Def:</i> Identification of noise, inaccuracies and inconsistencies	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	
			Question Format <i>Def:</i> Direct modification on the encoding of the question while keeping the logical structure intact	Format Change <i>Def:</i> Rephrasing the question in a different format	G36. Setting Rephrase G37. Change Sequence G38. Close Format G39. Data Restructuring	Change Setting Change Sequence True False Value Structuring
		Answer Format <i>Def:</i> Indirect modification on the output form	Format Comparison <i>Def:</i> Comparing two problem of different forms	G40. Identical Problem	Identical Question	Identical Code
		Format Constraint <i>Def:</i> Add constraint on the solution	G41. Reasoning Format G42. Reasoning Style G43. Alternative Answer G44. New Rule	Binary Coded X Language Alternative Answer Define Rules	No Keyword X Language Alternative Answer Simple Name	

Table 2: Our proposed ontology framework with domain, dimension, mathematical and code realization categories.

4.2 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 C.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. Failed questions are regenerated and re-evaluated, with persistent failures 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 MORE 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 annotator and verified by two others.

4.3 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.

4.4 Statistics of MORE and 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 MORE and CORE, respectively. We specify the detailed statistics in Appendix B and the details of the five

Original	Logic Alteration	Concept Analysis	Format Constraint
GSM8K		MORE	
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? 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.	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? 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 $X*(Y-2)=9$.	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? One possible answer: Calculate by using: total inner volume = total outer volume - volume of the wall. Yes.	Question - Reasoning Format: Answer the following question with only base-2 coded values. ... What is the total inner volume of all 3 boxes? 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.
HumanEval		CORE	
Question: <pre>def gcd(a:int, b:int): """Return the greatest common divisor of two integers a and b""" >>>gcd(25, 15) 5"""</pre> Answer: <pre>while b: a, b = b, a % b return abs(a)</pre>	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> Answer: Both a and b share 7 as their greatest common divisor.	Question - Step Necessity: Justify the necessity and functionality of operation "abs(a)" in fulfilling the requirement. {##Original Question##} {##Original Answer##} One possible answer: Ensures GCD result is non-negative, aligning with mathematical expectations.	Question - Reasoning Format: Fulfill the coding requirement below without using python keyword "while" inside the answer. {##Original Question##} {##Original Answer##} Answer: <pre>if b == 0: return abs(a) else: return gcd(b, a)</pre>

Figure 2: 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

selected question from each dataset in Appendix F and Appendix G respectively.

5 Experiments

5.1 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 valid ways to optimize the given solution. To address these challenges, manual evaluation is necessary. To empirically justify this, we prompted GPT-4 for automated answer evaluation, yielding an agreement of 88.76% with human annotation on the answers of GPT-4 to MORE questions.

5.2 Experimental Setup

We evaluated five prominent closed- and open-sourced LLMs on our benchmark. The closed-sourced LLMs are GPT-4, GPT-3.5, and Gemini 1.5. 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 coding and maths, 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 C.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 in Appendix J on the evaluation settings.

5.3 Experimental Results and Analyses

General Performance Analysis. The results 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 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 finetuning 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 concepts. 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

Aspect	Target	Original	Structural								Representational				Weighted Avg.	
			Logic					Concept			Q. 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
Maths (MORE)	GPT-4	100	100	80	90.91	60	78.30	85	65	48	64.62	90	60	84.00	65	74.21
	GPT-3.5	80	75	27.5	54.55	25.71	38.68	55	45	12	35.38	35	40	36.00	5	35.75
	Gemini	80	90	50	81.82	37.14	56.60	60	20	16	30.77	55	20	48.00	30	46.15
	Llama	60	50	12.5	18.18	5.71	17.92	35	60	4	30.77	5	60	16.00	5	26.24
	Metamath	80	70	15	27.27	11.43	25.47	30	25	4	18.46	35	80	44.00	20	21.27
	Average	80	77	37	54.55	27.90	43.39	53	43	16.8	36.00	44	52	45.60	25	40.72
Coding (CORE)	GPT-4	80	90	37.5	46.67	50	52.29	65	80	44	61.54	65	40	60.00	55	56.7
	GPT-3.5	80	73.68	35	40	29.41	40.74	60	75	40	56.92	50	40	48.00	45	47.09
	Gemini	80	80	32.5	53.33	23.53	41.28	65	75	44	60.00	45	40	44.00	35	47.32
	Llama	60	45	12.5	33.33	11.76	21.10	50	50	8	33.85	25	40	28.00	20	36.61
	CodeLlama	60	80	40	40	11.76	38.53	35	35	28	32.31	40	0	32.00	40	26.34
	Average	72	73.74	31.5	42.67	25.29	38.79	55	63	32.8	48.92	45	32	42.40	39.00	42.81

Table 3: Model performance on maths and coding across various *Dimensions* (Level III of ontology). All the average reported is weighted average.

	Mod.	Q. Simp.	R. Adj.	C. Adj.	S. Man.	Avg.
MORE	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
CORE	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 4: The impact of incorporating the original question and answer into the prompt on the performance of *logic* Target within the MORE and CORE. The reported average is weighted average.

	Mod.	Q. Simp.	R. Adj.	C. Adj.	S. Man.	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 5: The impact of using prompting techniques on the performance of *Logic* Target within the MORE and CORE. Self-C stands for Self-Consistency prompting (Wang et al., 2022) and POT stands for Program of Thought (Chen et al., 2022)

analysis in Appendix D

Incorporation of Original Answer in Prompt. In Table 4, 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 *Symbolic Manipulation*, 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 5, 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 an number or code as answer, this involves manipulation of abstract maths 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 ChatGPT 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.

6 Discussions

6.1 Difficulty Change by Perturbations

The performance drop may stem from an increased scale of reasoning or a higher level of abstract reasoning required. To explore this, 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. Difficulty was also evaluated through A/B testing and by recording human performance and response times across various perturbation categories as detailed in Appendix E. Table 7 conducted human evaluation on 44 perturbation types, 11 increased the number of reasoning steps needed and 10 required deeper reasoning compared to 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—11.6% for increased reasoning steps and 3.9% for deeper reasoning. Further, There was also a more than thirty percent change in model performance for perturbed questions of equal difficulty, indicating that increased complexity have minor impacts on model performance, the major performance gap may still come from lack of robustness of LLMs.

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 6: 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 original question

6.2 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. We document the aspects we have modified, the rationale behind these changes in Appendix I.1.

6.3 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, as detailed in Table 7, indicate that GPT-4 can successfully filter out challenging perturbation categories, achieving a perturbation success rate of over 90%.

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, MORE and 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 MORE and 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. Our dataset MORE and CORE will be publicly available online.

8 Limitations

Despite our attempt to construct a novel systematic ontology to evaluate an LM’s "real" robustness and reasoning capabilities in structured reasoning tasks,

557	it may not precisely reflect LLM’s true ability due	Emily M. Bender, Timnit Gebru, Angelina McMillan-	603
558	to several factors.	Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? . <i>Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency</i> .	604
559	Incompleteness In our endeavor to develop a	Lukas Berglund, Meg Tong, Max Kaufmann, Mikita	605
560	comprehensive ontology for evaluating Language	Balesni, Asa Cooper Stickland, Tomasz Korbak, and	606
561	Models’ (LMs) responses to perturbed questions	Owain Evans. 2023. The reversal curse: Lms trained on "a is b" fail to learn "b is a" . <i>ArXiv</i> ,	607
562	across various reasoning scenarios, we recognize	abs/2309.12288.	608
563	significant limitations. Firstly, despite our efforts,	Alexander Bondarenko, Magdalena Wolska, Stefan	609
564	the ontology may not fully capture all essential	Heindorf, Lukas Blübaum, Axel-Cyrille Ngonga	610
565	aspects of reasoning abilities, lacking in breadth	Ngomo, Benno Stein, Pavel Braslavski, Matthias Ha-	611
566	and depth. Secondly, the complexity within each	gen, and Martin Potthast. 2022. Causalqa: A benchmark for causal question answering . In <i>International Conference on Computational Linguistics</i> .	616
567	reasoning category can vary significantly. For in-	Mark Chen, Jerry Tworek, Heewoo Jun, Qiming	617
568	stance, within the <i>Computation Demand</i> category,	Yuan, Henrique Ponde de Oliveira Pinto, Jared Ka-	618
569	adjusting the number of digits in mathematical op-	plan, Harri Edwards, Yuri Burda, Nicholas Joseph,	619
570	erations allows us to modulate the reasoning chal-	Greg Brockman, Alex Ray, Raul Puri, Gretchen	620
571	lenge. However, creating a benchmark that exhaust-	Krueger, Michael Petrov, Heidy Khlaaf, Girish Sas-	621
572	ively encompasses all facets of reasoning behavior	try, Pamela Mishkin, Brooke Chan, Scott Gray,	622
573	is an unattainable goal. Such an exhaustive com-	Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz	623
574	pileation is beyond the scope of any single study	Kaiser, Mohammad Bavarian, Clemens Winter,	624
575	and necessitates collective efforts from the broader	Philippe Tillet, Felipe Petroski Such, Dave Cum-	625
576	research community.	mings, Matthias Plappert, Fotios Chantzis, Eliza-	626
577	Scalability The size of our dataset is constrained	beth Barnes, Ariel Herbert-Voss, William Hebg-	627
578	due to the human in the loop required for its prepa-	Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie	628
579	ration. Each question generated by GPT-4 needs to	Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain,	629
580	be meticulously reviewed to ensure it is solvable	William Saunders, Christopher Hesse, Andrew N.	630
581	and accurately reflects the intended perturbation	Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan	631
582	specific to its category, without introducing unin-	Morikawa, Alec Radford, Matthew Knight, Miles	632
583	tended modifications. Furthermore, confirming the	Brundage, Mira Murati, Katie Mayer, Peter Welinder,	633
584	accuracy of answers is a critical step, as many ques-	Bob McGrew, Dario Amodei, Sam McCandlish, Ilya	634
585	tions do not yield answers that exactly match a	Sutskever, and Wojciech Zaremba. 2021a. Evaluating large language models trained on code .	635
586	predefined format. This verification process limits		636
587	our ability to expand the dataset on a large scale,		637
588	as it relies on manual effort.		638
589	9 Potential Risks	Mark Chen, Jerry Tworek, Heewoo Jun, Qiming	639
590	Not applicable.	Yuan, Henrique Ponde de Oliveira Pinto, Jared Ka-	640
591	10 Ethical Considerations	plan, Harri Edwards, Yuri Burda, Nicholas Joseph,	641
592	Not applicable.	Greg Brockman, Alex Ray, Raul Puri, Gretchen	642
593	References	Krueger, Michael Petrov, Heidy Khlaaf, Girish Sas-	643
594	Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui	try, Pamela Mishkin, Brooke Chan, Scott Gray,	644
595	Zhang, and Wenpeng Yin. 2024. Large language models for mathematical reasoning: Progresses and challenges . <i>ArXiv</i> , abs/2402.00157.	Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz	645
596		Kaiser, Mohammad Bavarian, Clemens Winter,	646
597		Philippe Tillet, Felipe Petroski Such, Dave Cum-	647
598	Jacob Austin, Augustus Odena, Maxwell Nye, Maarten	mings, Matthias Plappert, Fotios Chantzis, Eliza-	648
599	Bosma, Henryk Michalewski, David Dohan, Ellen	beth Barnes, Ariel Herbert-Voss, William Hebg-	649
600	Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and	Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie	650
601	Charles Sutton. 2021. Program synthesis with large language models . <i>ArXiv</i> , abs/2108.07732.	Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain,	651
602		William Saunders, Christopher Hesse, Andrew N.	652
		Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan	653
		Morikawa, Alec Radford, Matthew Knight, Miles	654
		Brundage, Mira Murati, Katie Mayer, Peter Welinder,	655
		Bob McGrew, Dario Amodei, Sam McCandlish, Ilya	656
		Sutskever, and Wojciech Zaremba. 2021b. Evaluating large language models trained on code .	657
			658
		Wenhu Chen, Xueguang Ma, Xinyi Wang, and	659
		William W. Cohen. 2022. Program of thoughts prompting: Disentangling computation from rea-	660
			661

662	soning for numerical reasoning tasks. <i>ArXiv</i> ,	Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng	715
663	abs/2211.12588.	Kong, and Wei Bi. 2024a. Gsm-plus: A comprehensive benchmark for evaluating the robustness of llms as mathematical problem solvers. <i>ArXiv</i> ,	716
664	Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,	abs/2402.19255.	717
665	Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias	Xiaoyuan Li, Wenjie Wang, Moxin Li, Junrong Guo,	718
666	Plappert, Jerry Tworek, Jacob Hilton, Reiichiro	Yang Zhang, and Fuli Feng. 2024b. Evaluating mathematical reasoning of large language models: A focus on error identification and correction.	719
667	Nakano, Christopher Hesse, and John Schulman.		720
668	2021a. Training verifiers to solve math word problems. <i>ArXiv</i> , abs/2110.14168.		721
669			722
670	Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,	Zekun Li, Baolin Peng, Pengcheng He, and Xifeng	724
671	Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias	Yan. 2023. Do you really follow me? adversarial instructions for evaluating the robustness of large language models. <i>ArXiv</i> , abs/2308.10819.	725
672	Plappert, Jerry Tworek, Jacob Hilton, Reiichiro		726
673	Nakano, Christopher Hesse, and John Schulman.		727
674	2021b. Training verifiers to solve math word problems. <i>ArXiv</i> , abs/2110.14168.		
675		Jordan Meadows, Marco Valentino, Damien Teney, and	728
676	Ishita Dasgupta, Andrew Kyle Lampinen, Stephanie	André Freitas. 2023. A symbolic framework for evaluating mathematical reasoning and generalisation with transformers.	729
677	C. Y. Chan, Antonia Creswell, Dharshan Kumaran,		730
678	James L. McClelland, and Felix Hill. 2022. Language models show human-like content effects on reasoning. <i>ArXiv</i> , abs/2207.07051.		731
679		Philipp Mondorf and Barbara Plank. 2024. Beyond accuracy: Evaluating the reasoning behavior of large language models - a survey. <i>ArXiv</i> , abs/2404.01869.	732
680			733
681	Aniruddha Deb, Neeva Oza, Sarthak Singla, Dinesh		734
682	Khandelwal, Dinesh Garg, and Parag Singla. 2023.	OpenAI. 2023. Gpt-4 technical report. <i>arXiv preprint</i>	735
683	Fill in the blank: Exploring and enhancing llm capabilities for backward reasoning in math word problems. <i>ArXiv</i> , abs/2310.01991.	<i>arXiv:2303.08774</i> .	736
684		Arkil Patel, S. Bhattamishra, and Navin Goyal. 2021.	737
685		Are nlp models really able to solve simple math word problems? In <i>North American Chapter of the Association for Computational Linguistics</i> .	738
686	Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon,		739
687	Pengfei Liu, Yiming Yang, Jamie Callan, and Gra-		740
688	ham Neubig. 2022. Pal: Program-aided language models. <i>ArXiv</i> , abs/2211.10435.		
689		Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin,	741
690	Dan Hendrycks, Collin Burns, Steven Basart, Andy	and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of nlp models with checklist. <i>arXiv preprint</i>	742
691	Zou, Mantas Mazeika, Dawn Xiaodong Song, and	<i>arXiv:2005.04118</i> .	743
692	Jacob Steinhardt. 2020. Measuring massive multitask language understanding. <i>ArXiv</i> , abs/2009.03300.		744
693		Baptiste Rozière, Jonas Gehring, Fabian Gloeckle,	745
694	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul	Sten Sootla, Itai Gat, Xiaoqing Tan, Yossi Adi,	746
695	Arora, Steven Basart, Eric Tang, Dawn Song, and	Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom	747
696	Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset.	Kozhevnikov, I. Evtimov, Joanna Bitton, Manish P	748
697		Bhatt, Cristian Cantón Ferrer, Aaron Grattafiori, Wen-	749
698	Yinya Huang, Ruixin Hong, Hongming Zhang, Wei	han Xiong, Alexandre D’efosseuz, Jade Copet, Faisal	750
699	Shao, Zhicheng YANG, Dong Yu, Changshui Zhang,	Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier,	751
700	Xiaodan Liang, and Linqi Song. 2023. Clomo: Counterfactual logical modification with large language models. <i>ArXiv</i> , abs/2311.17438.	Thomas Scialom, and Gabriel Synnaeve. 2023. Code llama: Open foundation models for code. <i>ArXiv</i> ,	752
701		abs/2308.12950.	753
702			754
703	Albert Qiaochu Jiang, Alexandre Sablayrolles, Arthur	Soumya Sanyal, Zeyi Liao, and Xiang Ren. 2022. Robustlr: A diagnostic benchmark for evaluating logical robustness of deductive reasoners.	755
704	Mensch, Chris Bamford, Devendra Singh Chap-		756
705	lot, Diego de Las Casas, Florian Bressand, Gi-		757
706	anna Lengyel, Guillaume Lample, Lucile Saulnier,	Tomohiro Sawada, Daniel Paleka, Alexander Havrilla,	758
707	L’elio Renard Lavaud, Marie-Anne Lachaux, Pierre	Pranav Tadepalli, Paula Vidas, Alexander Kranias,	759
708	Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang,	John J. Nay, Kshitij Gupta, and Aran Komatsuzaki.	760
709	Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. <i>ArXiv</i> , abs/2310.06825.	2023. Arb: Advanced reasoning benchmark for large language models. <i>ArXiv</i> , abs/2307.13692.	761
710			762
711	Divyansh Kaushik, Amrith Setlur, Eduard Hovy, and	Freda Shi, Xinyun Chen, Kanishka Misra, Nathan	763
712	Zachary C Lipton. 2021. Explaining the efficacy of counterfactually augmented data. <i>International Conference on Learning Representations (ICLR)</i> .	Scales, David Dohan, Ed Huai hsin Chi, Nathanael	764
713		Scharli, and Denny Zhou. 2023. Large language models can be easily distracted by irrelevant context.	765
714		In <i>International Conference on Machine Learning</i> .	766
			767

768	Saurabh Srivastava, B AnnaroseM, V AntoP, Shashank Menon, Ajay Sukumar, T AdwaithSamod, Alan Philippe, Stevin Prince, and Sooraj Thomas. 2024. Functional benchmarks for robust evaluation of reasoning performance, and the reasoning gap . <i>ArXiv</i> , abs/2402.19450.	A Recommendations	824
769		Based on our findings, we make the following recommendations as strategies to address the weaknesses we identified in the logical reasoning competencies of LLMs.	825
770		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 to perturbations by fine-tuning it with datasets in a variety of formats and adding augmented instructions.	826
771		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.	827
772			828
773			829
774	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 . <i>ArXiv</i> , abs/2401.00757.	B Dataset Details	845
775		In particular, there are a total of 5 maths questions for each category except <i>Change Subject</i> and <i>Reverse Engineering</i> , which have 3 and 4 questions, respectively, in MORE. Likewise, all but <i>Reverse Engineering</i> perturbation—with 4 questions—have 5 coding questions in CORE.	846
776			847
777			848
778			849
779	Siyuan Wang, Zhuohan Long, Zhihao Fan, Zhongyu Wei, and Xuanjing Huang. 2024. Benchmark self-evolving: A multi-agent framework for dynamic llm evaluation . <i>ArXiv</i> , abs/2402.11443.		850
780			851
781			852
782			853
783	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 . <i>ArXiv</i> , abs/2203.11171.	C Experiment Details	852
784		C.1 Model Details	853
785		• we use version "2023-09-01-preview" for both GPT-4 and GPT-3.5.	854
786		• Llama3-Instruct https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct	855
787	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 .	• MetaMath-70B-V1.0 https://huggingface.co/meta-math/MetaMath-70B-V1.0	856
788		• CodeLlama-70B-Instruct https://huggingface.co/codellama/CodeLlama-70b-Instruct-hf	857
789			858
790			859
791			860
792	Shijie Xia, Xuefeng Li, Yixin Liu, Tongshuang Wu, and Pengfei Liu. 2024. Evaluating mathematical reasoning beyond accuracy . <i>arXiv preprint arXiv:2404.05692</i> .	C.2 Filtering Criteria	864
793		(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.	865
794			866
795			867
796	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 . <i>arXiv preprint arXiv:2304.12244</i> .		868
797			869
798			870
799			871
800			872
801	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 . <i>ArXiv</i> , abs/2309.12284.		873
802			874
803			875
804			876
805			877
806	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 .		878
807			879
808			880
809			881
810			882
811	Zhongshen Zeng, Pengguang Chen, Shu Liu, Haiyun Jiang, and Jiaya Jia. 2023. Mr-gsm8k: A meta-reasoning benchmark for large language model evaluation .		883
812			884
813			885
814			886
815	Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. 2021. Minif2f: a cross-system benchmark for formal olympiad-level mathematics . <i>ArXiv</i> , abs/2109.00110.		887
816			888
817			889
818			890
819	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 <i>AAAI Conference on Artificial Intelligence</i> .		891
820			892
821			893
822			894
823			895

(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.

D Fine-Grained Analysis

As illustrated in Table 3, the introduction of perturbed questions poses significant challenges to all models in both maths and coding contexts. Specifically, GPT-4’s accuracy decreased from 100% to 74.2% and from 80% to 56.7% in math and coding scenarios respectively. This trend of performance degradation is even more pronounced in other LLMs, with all experiencing a decline exceeding 30 points in their weighted average performance across both the mathematical and coding datasets. For instance, GPT-3.5 witnessed a dramatic performance reduction from 80% to 35.75% on the mathematical dataset and from 80% to 47.09% for the coding dataset.

Notably, closed-source models consistently outperform open-source models in every tested dimension. Additionally, it has been observed that models which have undergone fine-tuning on task-specific data—such as, CodeLlama for coding problems and Metamath for math problems—show enhanced performance in the areas of *logic alteration* and *representational perturbations* as compared to the Llama2-Chat model. However, this fine-tuning process appears to compromise Llama2’s capabilities within the *concept analysis* domain. This observation suggests that the focus of fine-tuned, task-specific data on deriving a fixed solution might limit a model’s broader capacity for reasoning, thereby affecting its ability to analyze and comprehend the underlying problem-solving process.

(Level II) Target-wise Performance. Following Table 3, LLMs generally showed better results on *logic alteration* questions, which involve concrete reasoning steps in problem-solving. Despite this, even the state-of-the-art models struggled with certain perturbed versions of these questions. This indicates that while current models may possess general task-solving skills and abstract reasoning

ability, there is still a limitation in their reasoning robustness when faced with altered logic. On the other hand, *concept analysis* questions, which demand a deeper understanding of mathematical concepts and problem-solving frameworks, resulted in lower success rates. This suggests that while current models can find correct answers, they may lack a systematic logical framework for problem-solving and struggle with analyzing and understanding different concepts necessary to answer the question.

GPT-4, in particular, demonstrated superior performance across all categories, showing increased resilience to changes in question format and expected responses. This contrasts with other models, which performed poorly on tasks involving representational perturbations, hinting at a limitation in transferring their reasoning processes to different formats. Interestingly, the average performance decline across domains was similar for both math and coding contexts, with the notable exception of the *concept analysis* domain, where the drop in math performance was 21% greater than in coding. This discrepancy suggests that LLMs may possess a more profound understanding of problem-solving frameworks in coding contexts compared to mathematical ones.

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 documented changes in the number of steps required to solve perturbed questions compared to the original, noting whether the number of steps increased, or remained roughly the same. Additionally, they assessed whether the perturbed variants demanded more higher level 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 physi-

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 7: 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.

cal 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

A merchant wants to make a choice of

973

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?

Answer: If he purchases jewelry, he will make a profit of 2.5% which is $5000 * (2.5/100) = 125$. If he purchases electronic gadgets, he will make a profit of 1.2% which is $8000 * (1.2/100) = 96$. 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 $2 * 1 = 2$ 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 $4 * 3 * 2 = 24$ cubic inches. So in total the inner volume of the 3 boxes is $3 * 24 = 72$ 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
```

974

975

977

978

976

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

```

994 3 """For a given string, flip
995 lowercase characters to uppercase
996 and uppercase to lowercase.
997
998 4
999 5 >>> flip_case('Hello')
1000 6 'hELLO'
1001 7 """
8 return string.swapcase()

```

G.2 Question 2

```

1003 1 def greatest_common_divisor(a: int, b:
1004 int) -> int:
1005 2
1006 3 """ Return a greatest common divisor
1007 of two integers a and b
1008 4
1009 5 >>> greatest_common_divisor(3, 5)
1010 6 1
1011 7 >>> greatest_common_divisor(25, 15)
1012 8 5
1013 9 """
1014 10
1015 11 while b:
1016 12     a, b = b, a % b
1017 13 return abs(a)

```

G.3 Question 3

```

1018 1 def derivative(xs: list):
1019 2
1020 3 """ xs represent coefficients of a
1021 polynomial.
1022 xs[0] + xs[1] * x + xs[2] * x^2 +
1023 ....
1024 Return derivative of this polynomial
1025 in the same form.
1026 4
1027 5 >>> derivative([3, 1, 2, 4, 5])
1028 6 [1, 4, 12, 20]
1029 7 >>> derivative([1, 2, 3])
1030 8 [2, 6]
1031 9 """
1032 10
1033 11 if len(xs) == 1: return [0]
1034 12 if len(xs) == 0: return []
1035 13 return [(i * x) for i, x in
1036 enumerate(xs)][1:]

```

G.4 Question 4

```

1037 1 def sum_squares(lst):
1038 2
1039 3 """
1040 4 This function will take a list of
1041 integers. For all entries in the
1042 list, the function shall square the
1043 integer entry if its index is a
1044 multiple of 3 and will cube the
1045 integer entry if its index is a
1046 multiple of 4 and not a multiple of
1047 3. The function will not
1048 change the entries in the list whose
1049 indexes are not a multiple of 3 or
1050 4. The function shall then return
1051 the sum of all entries.
1052 5
1053 6 Examples:
1054 7
1055 8

```

```

9 For lst = [1,2,3] the output should
10 be 6
11 For lst = [] the output should be 0
12 For lst = [-1,-5,2,-1,-5] the
13 output should be -126
14 """
15 result = []
16 for i in range(len(lst)):
17     if i%3 == 0:
18         result.append(lst[i]**2)
19     elif i% 4 == 0 and i%3 != 0:
20         result.append(lst[i]**3)
21     else:
22         result.append(lst[i])
23 return sum(result)

```

G.5 Question 5

```

1055 1 def is_nested(string):
1056 2
1057 3 """
1058 4 Create a function that takes a
1059 string as input which contains only
1060 square brackets.
1061 The function should return True if
1062 and only if there is a valid
1063 subsequence of brackets
1064 where at least one bracket in the
1065 subsequence is nested.
1066 Examples:
1067 [[]] output: True
1068 [][] output: False
1069 [] output: False
1070 [[][]] output: True
1071 [[]][[] output: True
1072 """
1073 5
1074 6 stack = []
1075 7 depth = 0
1076 8 for i, char in enumerate(string):
1077 9     if char == '[':
1078 10         stack.append([''])
1079 11         if depth > 0:
1080 12             depth -= 1
1081 13         elif char == ']':
1082 14             if len(stack) > 0:
1083 15                 stack.pop()
1084 16                 depth += 1
1085 17             if depth >= 2:
1086 18                 return True
1087 19             if len(stack) == 0:
1088 20                 depth = 0
1089 21 return False
1090 22
1091 23
1092 24
1093 25
1094 26
1095 27
1096 28
1097 29
1098 30

```

H Ontology

The summary of our proposed ontological categories is shown in Table 2.

I Ontology of Perturbations

I.1 Principles behind Ontology

Consider a maths question:

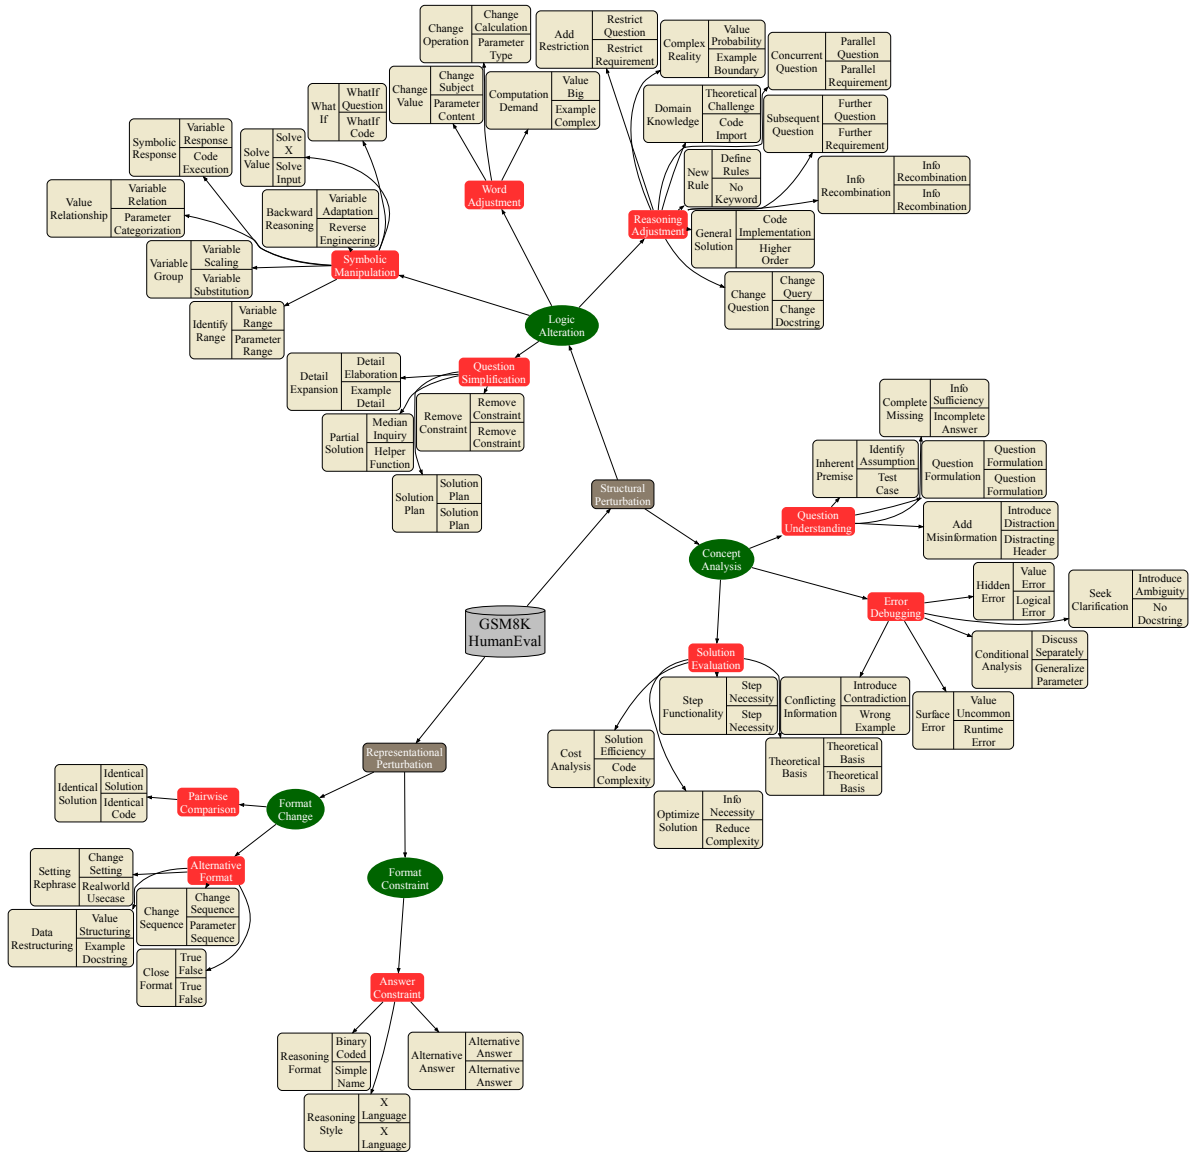


Figure 3: The ontology of the perturbations.

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:
1148 int) -> int:
1149     """ Return the greatest common
1150     divisor of two integers a and b
1151     Example:
1152     >>> greatest_common_divisor(3, 5)
1153     1
1154     >>> greatest_common_divisor(25, 15)
1155     5
1156     """
1157
```

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 """
1171 >>> greatest\_common\_divisor(3, 5)
1172 1
1173 >>> greatest\_common\_divisor(25, 15)
1174 5
1175 """
1176
```

(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 (Fig. 3) 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) **Question Simplification:** 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
4     all uppercase characters to
5     lowercase.
6
7     >>> change_case('Hello')
8     'hello'
9     """
```

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) ->
2     str:
3     """Checks if a given character
4     is uppercase or lowercase, and
5     flips its case."""
6
7     if char.isupper():
8         return char.lower()
9     elif char.islower():
10        return char.upper()
11    else:
12        return char
13
14 def flip_case(string: str) -> str:
15
16     """For a given string, flip
17     lowercase characters to
18     uppercase and uppercase to
19     lowercase by using the helper
20     function above to achieve the
21     requirement
22     >>> flip_case('Hello')
23     'hELLO'
24     """
```

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
4     character in the provided string
5     .
6     This function takes a string as
7     an argument and returns a new
8     string with each character's
9     case inverted.
10    Uppercase letters are converted
11    to lowercase, and lowercase
12    letters are converted to
13    uppercase.
14
15    Solution Plan:
16    1. Create a result variable to
17    hold the updated string.
18    2. Iterate through each
19    character in the string.
20    3. Check if the character is
21    uppercase; if so, convert it to
22    lowercase and add it to the
23    result.
24    4. If the character is lowercase
25    , convert it to uppercase and
26    add it to the result.
27    5. After iterating through all
28    characters, return the result.
29    """

```

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
3     a polynomial.
4     xs[0] + xs[1] * x + xs[2] * x^2
5     + ....
6     Return derivative of this
7     polynomial in the same form.
8
9     >>> derivative([3, 1, 2, 4, 5])
10    calculates the derivative as
11    [1*1, 2*2, 3*4, 4*5] resulting
12    in [1, 4, 12, 20].
13
14    >>> derivative([1, 2, 3])
15    calculates the derivative as
16    [1*2, 2*3] resulting in [2, 6].
17    """

```

(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 exter-

nally 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:
2     int) -> str:
3
4     """For a given string, flip
5     lowercase characters to
6     uppercase and uppercase to
7     lowercase. Only flip the case
8     for characters at indices which
9     are multiples of the provided
10    index.
11    Note: If the index provided is
12    2, only the characters at the 2
13    nd, 4th, 6th positions and so on
14    will have their cases flipped.
15
16    >>> flip_case('Hello', 2)
17    'HeLLo'
18    """

```

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

```

```

5
6     >>> flip_case_count('Hello')
7     ('hELLO', 5)
8     """

```

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)
2     -> Tuple[str, int]:
3
4     """For a given string, not only
5     should you flip lowercase
6     characters to uppercase and
7     uppercase to lowercase. You
8     should also output another Title
9     case where only the first
10    letter of each word is
11    capitalized"""
12
13    """>>> flip_case_and_count('
14    Hello')
15    ('hELLO', 'Hello')
16    """

```

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

1494 in the **question header**.

1495 **Changed from G.3:**

```
1496 def calc_derivative(xs: list):
1497
1498     """ xs represent coefficients of
1499     a polynomial.
1500     xs[0] * (exp(x))^0 + xs[1] * (
1501     exp(x))^1 + xs[2] * (exp(x))^2 +
1502     ....
1503     Return derivative of this
1504     polynomial in the same form.
1505     >>> derivative([3, 1, 2, 4, 5])
1506     [1, 4, 12, 20]
1507     >>> derivative([1, 2, 3])
1508     [2, 6]
1509     """
```

1510 **G9. Info Recombination:** Combine the fundamental concepts or frameworks from another question with the original question.

1511 *Info Recombination (Math):* Graft **mathematical structure** from another question and combine with the original question.

1512 **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?

1513 *Info Recombination (Code):* Merge the requirement from another coding question with existing question. **Changed from G.1:**

```
1514 def flip_case_and_odd_sum(string:
1515 str) -> tuple:
1516     """
1517     Given a string, flip lowercase
1518     characters to uppercase and
1519     uppercase to lowercase.
1520     Also return the odd letters that
1521     are in even positions of the
1522     original string.
1523     string index starts from 0,
1524     alphabet index start from 1. Aa
1525     is 1, Bb is 2..
1526     Examples:
1527     >>> flip_case_and_odd_sum('Hello
1528 ')
1529     ('hELLO', 'o')
1530     """
```

1531 **G10. Domain Knowledge:** Introduce a specific knowledge in math or code and merge it with the question.

1540 *Theoretical Challenge:* Incorporate a specific theorem into the question so that perturbed question requires a new **toolbox** to solve.

1541 **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?

1542 *Code Import:* The requirement requires to use a specific python library to solve the problem.

1543 **Changed from G.1:** Rewrite the function below to take in batch input parameters and use the multicore cpu for efficiency.

1544 **G11. Complex Reality:** Add an aspect of complexity in the real world scenario.

1545 *Value Probability:* Introduce concept of uncertainty to deterministic values and calculate the estimation. The perturbed question will require **toolbox** (knowledge) of probability.

1546 **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?

1547 *Example Boundary:* Add boundary examples along with the existing **examples**. The boundary examples contains input that does not met requirement specified in the docstring.

1548 **Changed from G.3:** Write a function to fulfill the requirement and all the examples inside the docstring

```
1549 def derivative(xs: list):
1550
1551     """ xs represent coefficients of
1552     a polynomial.
1553     xs[0] + xs[1] * x + xs[2] * x^2
1554     + ....
1555     Return derivative of this
1556     polynomial in the same form. The
1557     solution should pass all the
1558     test cases specified below
1559     """
```

```

1574 7 # Regular case
1575 8 >>> derivative([3, 1, 2, 4, 5])
1576 9 [1, 4, 12, 20]
1577 10 # Smaller case
1578 11 >>> derivative([1, 2, 3])
1579 12 [2, 6]
1580 13 # Special case with empty list
1581 14 >>> derivative([])
1582 15 []
1583 16 # Boundary case, the shortest
1584 17 polynomial
1585 18 >>> derivative([1])
1586 19 [0]
1587 20 # Boundary case, all-zero
1588 21 polynomial
1589 22 >>> derivative([0.0, 0.0, 0.0])
1590 [0, 0]
1591 """

```

1592 **G12. General Solution:** Provide the solution in a
1593 more general scenario.

1594 *Code Implementation:* Develop a code func-
1595 tion to solve the question in general.

1596 **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?

1597 *Higher Order:* Write a higher order function
1598 that can solve the coding problem in general.

1599 **Changed from G.2**

```

1600 1 def greatest_common_divisor(numbers:
1601 list[int]) -> int:
1602 """
1603 Calculates the greatest common
1604 divisor (GCD) of a list of
1605 integers.
1606 Returns the GCD as an integer.
1607
1608 Examples:
1609 - For numbers = [20, 40, 60],
1610 the function returns 20.
1611 - For numbers = [35, 14], the
1612 function returns 7.
1613 """

```

1614 (iii) **Computation Adjustment:** While retaining
1615 the **Logical Structure**, this type aims to change
1616 one single reasoning step of the original question.

1617 **G13. Computation Demand:** Change the value to
1618 complex values that put a high demand on
1619 computation.

1620 *Value Big:* Significantly increasing the magni-
1621 tude of values that pose a challenge for calcu-
1622 lations.

1623 **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?

1626 *Generalize Parameter:* Extend the current pa-
1627 rameter into different python object types

1628 **Changed from G.2:**

```

1629 1 def find_common_divisor(value1:
1630 Union[int, float, str], value2:
1631 Union[int, float, str]) -> float
1632 :
1633 """
1634 Takes two values (int, float, or
1635 float in string format) and
1636 finds the largest float that
1637 divides both into integers.
1638 Inputs can be a mix of types.
1639 Returns the divisor as a float.
1640
1641 Examples:
1642 print(find_common_divisor("0.5",
1643 1)) # 0.5
1644 print(find_common_divisor(0.25,
1645 "1.25")) # 0.25
1646 """

```

1647 **G14. Change Value:** Change the content of the
1648 value to a different one.

1649 *Change Subject:* If there are multiple men-
1650 tions in the question, Exchange **values** of
1651 names or references in the question.

1652 **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?

1653 *Parameter Content:* Change the format or
1654 meaning of the input parameter.

1655 **Changed from G.3:**

```

1657 1 def derivative(polynomial: str):
1658 2
1659 3 """ 'polynomial' is a string
1660 that stands for polynomial for
1661 form
1662 coefficients_0 + coefficients_1
1663 * x + coefficients_2 * x^2 +
1664 ....
1665 This function will return the
1666 derivative of the aforementioned
1667 polynomial in the same format.
1668
1669 >>> derivative('3 +1x + 2x^2 + 4
1670 x^3 + 5x^4')
1671 '1 + 4x + 12x^2 + 20x^3'
1672 >>> derivative('1 - 2x + 3x^2')
1673 '-2 + 6x'
1674 """

```

1675 **G15. Change Operation:** Change one operation re-
1676 garding 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) **Symbolic Manipulation:** This dimension test the abstract reasoning ability of under the same logical structure of the original question. This dimension focus on solving the general version of the original reasoning problem, rather than focus 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. Values 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 lowercase to
  uppercase or uppercase to
  lowercase. The function will
  only flip the case of the
  specific value in the string.
```

1760
1761
1762

```

4 >>> flip_case('Hello', 'h')
5 'hello'
6 """

```

1763
1764
1765
1766
1767
1768
1769
1770
1771

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.

1772
1773
1774
1775
1776
1777
1778
1779

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

1780
1781

```

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

```

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792

G20. What If: 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 following:
2 >>> original_function(masked_input)
3 'HELLO'
4
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)`"

```

1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816

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?

1817
1818
1819
1820
1821
1822

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()

```

1823
1824
1825
1826
1827
1828

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?

1829
1830
1831
1832
1833
1834

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
3             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
3     lowercase characters to
4     uppercase and uppercase to
5     lowercase.
6     """
```

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

Missing Info: 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
3     lowercase characters to
4     uppercase and uppercase to
5     lowercase.
6     >>> flip_case('Hello')
7     'hELLO'
8     """
9     [masked code paragraph]
10     if char.isupper():
11         result += char.lower()
12     else:
13         result += char.upper()
14     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): $5000*(2.5/100) = \$125$ $8000*(1.2/100) = \$96$ $\$125 > \96 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
3     lowercase characters to
4     uppercase and uppercase to
5     lowercase.
6     >>> flip_case('hello')
7     'HELLO'
8     """
```

(ii) **Solution Evaluation:** 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(
2     coefficients, derivative=None,
3     index=0):
4     """
5     This function calculates the
6     derivative of a polynomial using
7     recursion.
8     coefficients: List of
9     coefficients of the polynomial.
10    derivative: List to store the
11    coefficients of the derivative.
12    Initially None.
13    index: Current index in the
14    coefficients list.
15
16    The base case of the recursion
17    is when the index is equal to
18    the length of the coefficients
19    list.
20    """
21    if index > 0:
22        derivative_coefficient =
23        index * coefficients[index]
24        derivative.append(
25        derivative_coefficient)
26    return derivative_polynomial(
27    coefficients, derivative, index
28    + 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,
2   b: int) -> int:
3
4     """ Return a greatest common
5     divisor of two integers a and b
6     >>> greatest_common_divisor(3,
7     5)
8     1
9     >>> greatest_common_divisor(25,
10    15)
11    5
12    """
13    while b:
14        a, b = b, a % b
15    # What is the functionality of `
16    abs()`
17    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
3     lowercase characters to
4     uppercase and uppercase to
5     lowercase.
6     >>> flip_case('Hello')
7     'hELLO'
8     """

```

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
3     lowercase characters to
4     uppercase and uppercase to
5     lowercase.
6     >>> flip_case('Hello')
7     'hELLO'
8     """

```

(iii) **Spot Error:** 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

2077 be solved without clarification.

2078 **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?

2079 *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.

```
2088 1 def flip_case(string: str) -> str:
2089 2     """
2090 3     function('Hello') == 'hELLO'
2091 4     function('Python 3.8') == '
2092 5     pYTHON 3.8'
2093 6     function('123abcXYZ') == '123
2094 7     ABCxyz'
2095 8     function('MixedCASE123') == '
2096 9     mIXEDcase123'
2097 10    function('ALLUPPERCASE') == '
2098 11    alluppercase'
2099 12    """
```

2100 **G32. Conditional Analysis:** Based on different possible situations of the question, the answer should separately presented.

2101 *Discuss Separately:* Introduce new **information** containing variables or conditions that require the answer to be discussed separately based on conditions or variables.

2102 **Changed from F.1:**

2103 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?

2104 *Incomplete requirement:* Left some condition unspecified in the docstring. **Changed from G.1:**

```
2108 1 def flip_case(ch: str) -> str:
2109 2     """For a given string, all the
2110 3     letters inside the string should
2111 4     be changed. flip lowercase
2112 5     characters to uppercase.
2113 6     """
```

```
2117 4     >>> flip_case('h')
2118 5     'H'
2119 6     """
```

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:**

```
2129 1 def flip_case(string: str) -> str:
2130 2     """For a given string, flip
2131 3     lowercase characters to
2132 4     uppercase and uppercase to
2133 5     lowercase.
2134 6     >>> flip_case('Hello')
2135 7     'hello'
2136 8     """
```

2137 **G34. surface Error:** Introduce an obvious error that can be spot without reasoning. *Value Uncommon:* Change the **values** so that it seems wired or unusual by commonsense knowledge standards.

Changed from F.2:

2138 Can you spot anything unusual for 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

```
2146 1 def flip_case(string, str) -> str:
2147 2     """For a given string, flip
2148 3     lowercase characters to
2149 4     uppercase and uppercase to
2150 5     lowercase.
2151 6     >>> flip_case('Hello')
2152 7     'hELLO'
2153 8     """
2154 9     return string.swapcase()
```

2159 **G35. Hidden Error:** Introduce a hidden error that
 2160 need logical reasoning to spot. *Value Error:*
 2161 Change the **values** so that the question does
 2162 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?

Value Error: Introduce the change in the code that will cause a Value Error in python.

Changed from G.1

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

2177 I.4 Representational Perturbation – Format 2178 Change

2179 This dimension is inspired by metamath (Yu et al.,
 2180 2023b). It involves changing the **question repre-**
 2181 **sentation** by modifying the question encoding or
 2182 specify the representation of the answer in different
 2183 ways while keeping the underlying logical structure
 2184 and conceptual framework of the original question
 2185 intact. The objective is to verify whether the LLM
 2186 can still provide correct answers even when the for-
 2187 mat or presentation of the question changes. This
 2188 tests the model’s ability to reason irrespective of
 2189 how it’s presented. It also tests models’ instruction
 2190 following ability where the **answer representation**
 2191 must follow a certain format.

(i) **Alternative Format:**

2193 **G36. Setting Rephrase:** Rephrase the question in
 2194 another setting. *Change Setting:* Rephrase by
 2195 changing the application setting and values
 2196 inside the information, while keeping the core
 2197 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) ->
2     str:
3     """
4     Imagine you're working on a
5     document and you've mistaken the
6     case in the text you write. You
7     wrote all the lower case
8     letters in uppercase and vice
9     versa, suppose you want to
10    correct all of them using python
11    """
```

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:
2     int) -> int:
3     """ Return a common divisor that
4     is the largest of two integers
5     xray and yoyo
6     >>> munchee_bunchee(3, 5)
7     1
8     >>> munchee_bunchee(25, 15)
9     5
10    """
```

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.

2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291

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,
4   b: int) -> int:
5     """ Return a greatest common
6     divisor of two integers a and b
7     >>> greatest_common_divisor(3,
8     5)
9     1
10    >>> greatest_common_divisor(25,
11    15)
12    5
13    """
14
15 Solution:
16
17     while a:
18         a, b = a % b, a
19     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) ->
2   str:
3   """
4   For any specified sequence of
5   alphabetical characters,
6   interspersed with spaces,
7   numerical digits, and various
8   symbols, implement a
9   sophisticated transformation
10  algorithm designed to
11  selectively convert each
12  alphabetical character from its
13  current case representation,
14  either lowercase or uppercase,
15  to its diametrically opposite
16  case representation. This
17  algorithm ensures that every
18  character initially presented in
19  lowercase is meticulously
20  transmuted to uppercase, and
21  conversely, every character
22  originally in uppercase is
23  converted to lowercase, while
24  meticulously preserving the
25  integrity and original
26  positioning of spaces, numerical
27  digits, and any other non-
28  alphabetical symbols, leaving
29  these elements unaltered within
30  the sequence.
31  >>> function('Hello')
32  'hELLO'
33  """

```

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

2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333

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

```

(ii) **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,
2                             b: int) -> int:
3     """ Return a greatest common
4     divisor of two integers a and b
5     >>> greatest_common_divisor(3,
6     5)
7     1
8     >>> greatest_common_divisor(25,
9     15)
10    5
11    """

```

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
3 the case of each character:
4 lowercase to uppercase and
5 uppercase to lowercase.
6 }

```

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 rules and knowledge and use it inside the solution.

Define Rules: Define a new mathematical rule that will change how the **toolbox** (common-sense 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 to
  lowercase.
3 >>> flip_case('Hello')
4 'hELLO'
5 """
```

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 maths dataset MORE 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

in solving some other primary perturbation types which we call *Primary types*. For instance, consider the process of solving perturbed questions generated as outlined in **G20**. The initial step for the model involves identifying the value of an unknown variable from its answer. Subsequently, the model calculates how this value alters the final answer. This initial step demands skills similar to those described in **G21**. Consequently, we anticipate that enhanced perturbation types will be challenging to answer. Following Table 10, across all models, *primary types* exhibit higher overall performance as compared to *enhanced types*. Furthermore, it is observed that open-source models do not experience as significant a performance drop as closed-source models when handling enhanced types. This can be attributed to the fact that open-source models already demonstrate near-zero performance in answering primary type questions. Therefore, their inability to answer enhanced questions does not result in a notable decrease in performance.

Performance across Question Difficulty. In our experimentation with various LLMs, we consistently employed the Chain of Thought (CoT) methodology to derive the ultimate answer. This prompts a natural inquiry: *Does the performance of LLMs exhibit any correlation with the number of steps needed to arrive at the final answer?* Surprisingly, in our extensive experiments (as illustrated in Figure 4), we did not discern any definitive correlation or discernible trend. Instead, performance appears to diminish based on the inherent difficulty of the original question in GSM8K. Put differently, if an LLM fails to provide an accurate response to the initial question, its performance similarly falters when confronted with perturbed questions.

L Detailed Results

The detailed results across the perturbation categories for all the models are illustrated in Tables 8 and 9.

M Inclusivity of skill set

Dependence between Perturbation Types. In our ontology, some specialized perturbation types, which we refer to as *Enhanced types*, require skill

Dimension	Category	GPT-4	GPT-3.5	Gemini	Metamath	Llama2-Chat
Original		5	4	4	4	3
Question Simplification	Remove Constraint	5	5	5	5	4
	Partial Solution	5	3	3	3	2
	Solution Plan	5	4	5	2	4
	Detail Expansion	5	3	5	4	0
Reasoning Adjustment	Add Restriction	3	1	2	1	0
	Subsequent Question	4	1	3	0	0
	Concurrent Question	4	2	4	1	1
	Change Question	5	2	3	1	1
	Info Recombination	4	1	3	0	1
	Domain Knowledge	4	2	4	2	2
	Complex Reality	3	0	1	1	0
	General Solution	5	2	0	0	0
Computation Adjustment	Computation Demand	4	2	4	1	0
	Change Value	1	1	1	1	0
	Change Operation	5	3	4	1	2
Symbolic Manipulation	Symbolic Response	4	3	3	0	0
	Value Relationship	3	1	2	0	0
	Variable Group	3	1	2	0	0
	Backward Reasoning	2	1	1	1	0
	WhatIf	3	1	0	1	0
	Solve Value	5	2	4	1	0
	Identify Range	1	0	1	1	2
Question Understanding	Inherent Premise	5	2	2	0	1
	Complete Missing	5	4	5	2	4
	Question Formulation	3	1	2	1	1
	Add Misinformation	4	4	3	3	1
Solution Evaluation	Optimize Solution	3	3	2	2	4
	Step Functionality	1	0	0	0	2
	Theoretical Basis	4	4	1	2	4
	Cost Analysis	5	2	1	1	2
Error Debugging	Seek Clarification	1	2	1	0	0
	Conditional Analysis	3	0	2	0	0
	Conflicting Information	2	0	1	0	0
	Surface Error	4	1	0	1	1
	Hidden Error	2	0	0	0	0
Alternative Format	Setting Rephrase	4	3	2	4	1
	Change Sequence	5	2	3	3	0
	Close Format	4	2	4	0	0
	Data Restructuring	5	0	2	0	0
Pairwise Comparison	Identical Problem	3	2	1	4	3
Answer Constraint	Reasoning Format	4	0	2	0	0
	Reasoning Style	4	0	2	0	0
	Alternative Answer	2	0	0	2	0
	New Rule	3	1	2	2	1

Table 8: Number of examples correctly predicted by each model on MORE. There are a total of 5 questions for each category except "Change Value", which only has 2 questions.

Dimension	Category	GPT-4	ChatGPT	Gemini	CodeLlama	Llama2-Chat
Original		4	4	4	3	3
Question Simplification	Remove Constraint	4	4	4	4	2
	Partial Solution	5	3	4	5	2
	Solution Plan	5	4	4	3	2
	Detail Expansion	4	3	4	4	3
Reasoning Adjustment	Add Restriction	0	0	2	2	0
	Subsequent Question	2	2	1	1	3
	Concurrent Question	3	1	0	2	0
	Change Question	2	2	2	2	1
	Info Recombination	2	1	1	1	0
	Domain Knowledge	3	4	3	4	0
	Complex Reality	3	2	3	3	0
	General Solution	0	2	1	1	1
Computation Adjustment	Computation Demand	1	1	2	2	1
	Change Value	2	1	1	2	1
	Change Operation	4	4	5	2	3
Symbolic Manipulation	Symbolic Response	4	3	1	2	1
	Value Relationship	1	1	1	1	0
	Variable Group	3	1	1	0	1
	Backward Reasoning	2	2	3	1	0
	WhatIf	3	1	0	0	0
	Solve Value	1	1	0	0	0
	Identify Range	3	1	2	0	2
Question Understanding	Inherent Premise	2	3	2	1	1
	Complete Missing	3	1	3	1	2
	Question Formulation	4	4	4	2	3
	Add Misinformation	4	4	4	3	4
Solution Evaluation	Optimize Solution	2	2	3	2	2
	Step Functionality	5	5	4	2	2
	Theoretical Basis	5	3	4	0	4
	Cost Analysis	4	5	4	3	2
Error Debugging	Seek Clarification	2	2	2	3	0
	Conditional Analysis	1	1	1	0	0
	Conflicting Information	1	0	0	0	0
	Surface Error	4	4	4	2	1
	Hidden Error	3	3	4	2	1
Alternative Format	Setting Rephrase	3	2	2	1	3
	Change Sequence	4	3	2	3	1
	Close Format	3	1	2	2	0
	Data Restructuring	3	4	3	2	1
Pairwise Comparison	Identical Problem	2	2	2	0	2
Answer Constraint	Reasoning Format	2	2	2	2	1
	Reasoning Style	3	2	2	3	1
	Alternative Answer	3	3	2	2	0
	New Rule	3	2	1	1	2

Table 9: Number of examples correctly predicted by each model on CORE. There are a total of 5 questions for each category.

Domain	Enhanced Type	Primary Type	GPT-4	GPT-3.5	Gemini	Metamath	Llama2-Chat
Logic Alteration	Backward Reasoning	Solve Value	20	0	20	-20	0
	Value Relationship	Symbolic Response	20	40	40	0	0
	Variable Group	Symbolic Response	20	40	20	0	0
	Identify Range	Symbolic Response	60	60	40	-20	-40
	What If	Solve Value	40	20	80	0	0
	Solution Plan	Detail Expansion	0	20	0	-40	80
Concept Analysis	Seek Clarification	Conditional Analysis	40	-40	20	0	0
	Optimize Solution	Cost Analysis	20	20	20	20	-60
	Conflicting Information	Complete Missing	20	60	20	40	80
	Optimize Solution	Step Functionality	20	0	0	40	-40
	Hidden Error	Step Functionality	0	0	20	0	0
Average			25.45	16.36	30.91	0	12.73

Table 10: Performance drop in Enhanced vs. Primary Type questions on MORE. The value equals (accuracy of Primary - accuracy of Enhanced), so positive entries indicate higher performance for Primary Type questions.




Figure 4: Model performance for each question. The blue color indicates the model predicted correctly for the original question, and orange means the opposite. '3', '4', '5', '7', '8' stands for the number of steps in the gold answer for the perturbed question.