FRONTIER LLMs STILL STRUGGLE WITH SIMPLE REASONING TASKS

Anonymous authorsPaper under double-blind review

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

033 034 035

037

038

040

041

042 043

044

046

047

048

049

051

052

ABSTRACT

While state-of-the-art large language models (LLMs) demonstrate advanced reasoning capabilities—achieving remarkable performance on challenging competitive math and coding benchmarks—they also frequently fail on tasks that are easy for humans. This work studies the performance of frontier LLMs on a broad set of such "easy" reasoning problems. By extending previous work in the literature, we create a suite of procedurally generated simple reasoning tasks, including counting, first-order logic, proof trees, and travel planning, with changeable parameters (such as document length, or the number of variables in a math problem) that can arbitrarily increase the amount of computation required to produce the answer while preserving the fundamental difficulty. While previous work showed that traditional, non-thinking models can be made to fail on such problems, we demonstrate that even state-of-the-art thinking models consistently fail on such problems and for similar reasons (e.g., statistical shortcuts, errors in intermediate steps, and difficulties in processing long contexts). To further understand the behavior of the models, we introduce the UNPUZZLES dataset, a different "easy" benchmark consisting of trivialized versions of well-known math and logic puzzles. Interestingly, while modern LLMs excel at solving the original puzzles, they tend to fail on the trivialized versions, exhibiting several systematic failure patterns related to memorizing the originals. We show that this happens even if the models are otherwise able to solve problems with different descriptions but requiring the same logic. Our results highlight that out-of-distribution generalization is still problematic for frontier language models and the new generation of thinking models, even for simple reasoning tasks, and making tasks easier does not necessarily imply improved performance.

1 Introduction

Modern transformer-based large language models (LLMs) (Vaswani, 2017) trained using next-token prediction have achieved significant success across a wide range of tasks, especially in reasoning. For instance, OpenAI's o1 model—one of the leading reasoning models to date—"ranks in the 89th percentile on competitive programming questions (Codeforces), places among the top 500 students in the US in a qualifier for the USA Math Olympiad (AIME), and exceeds human PhD-level accuracy on a benchmark of physics, biology, and chemistry problems (GPQA)".

On the other hand, researchers continue to uncover surprisingly simple reasoning problems that still confuse even the most advanced LLMs. These include tasks such as counting characters in words, comparing numbers like 9.11 and 9.9 (Xie, 2024), making simple inferences about family relationships (Nezhurina et al., 2024), and solving various classes of arithmetic and logic problems (see, e.g., McLeish et al., 2024; Zhang et al., 2022). Many of these failures are identified in isolation, making it difficult to find common underlying issues. Moreover, some studies focus on earlier model generations, leaving it open whether these failures persist in state-of-the-art (SOTA) models.

In this work, we study the performance of several high-quality, open and closed-source language models, both traditional (GPT-40, Gemini 1.5 Pro 2.0 Flash, and 2.5 Pro, Gemma 3 27B, Claude 3.5 and 3.7 Sonnet) and thinking variants (OpenAI o1 and o3, Gemini 2.0 Flash Thinking and 2.5 Pro,

https://openai.com/index/learning-to-reason-with-llms/

Original Puzzle

054

055 056

058

059 060

061 062

063

064

065 066 067

068

069

071

073

074 075

076

077

079

081

083

085

087

880

089

090

091

092

093

094

095

096

097

098

100

101

102

103

104

105

106 107 13 purple, 15 yellow, and 17 maroon chameleons are found on an island. When two different-coloured chameleons meet in a pair, they both turn into the third color. Is it possible that, after some pairwise meetings, all the chameleons are the same color?

Answer: no (proof by contradiction related to a problem invariant)

Unpuzzle

15 purple, 15 yellow, and 17 maroon chameleons are found on an island. When two different-coloured chameleons meet in a pair, they both turn into the third color. Is it possible that, after some pairwise meetings, all the chameleons are the same color??

Answer: yes (purple and yellow chameleons all pair up)

Context-shifted Unpuzzle

There are 31 Spurs fans, 31 Arsenal fans, and 49 Chelsea fans. Every time fans of two different sports teams meet, they realize they are both wrong and become fans of the third team. Is it possible that, after a certain number of pairwise meetings, everyone is a fan of one team?

Answer: yes (Spurs fans and Arsenal fans all pair up)

Figure 1: Chameleons go on a date: a puzzle, corresponding unpuzzle, and a context-shifted unpuzzle

DeepSeek R1), across a broad range of "easy" reasoning problems. We begin by examining four simple reasoning tasks: (1) character and word counting, (2) first-order logic evaluation and negation, (3) math word problems based on proof trees, and (4) travel planning problems. Rather than using fixed datasets, we generate problems randomly and procedurally, incorporating tunable parameters such as paragraph length in word counting and the number of cities in travel planning—that adjust the amount of computation required to produce an answer while preserving the underlying reasoning difficulty. With appropriately chosen parameters, these tasks may be tedious for humans but remain straightforward. On the other hand, frontier LLMs consistently fail on such tasks, with underlying causes including statistical shortcuts, errors in intermediate steps, and difficulties in dealing with long contexts. While previous work showed that earlier SOTA models fail on similar tasks, here we demonstrate that even the next generation of LLMs, the so-called thinking models fail when the tasks become long enough. To our knowledge, no earlier papers demonstrated that thinking models are similarly subject to such performance degradation; essentially, we provide evidence that many of the claims in the literature of decreasing LLM performance with task difficulty will apply to thinking models as well. Concurrently with our work, Shojaee et al. (2025) evaluate LLMs on four puzzles with programmable complexity, and show that thinking LLMs completely fail beyond a certain critical complexity threshold. However, the experiment design has been found lacking (Opus and Lawsen, 2025; Chan, 2025).

In addition to the aforementioned tasks, we introduce a new dataset for evaluating language models called UNPUZZLES. This is a small dataset in two parts: The first has 97 well-known logical puzzles and brainteasers that are commonly found on the internet, as well as their *trivialized* versions which we refer to as "*unpuzzles*". Each unpuzzle is created manually by making minimal textual edits to the original puzzle in order to remove the difficulty and render the answer obvious. We demonstrate that while SOTA models all perform well on the original (difficult) puzzles, they exhibit poor performance on the corresponding (easy) unpuzzles. The second part focuses on a subset of 64 unpuzzles (with numerical answers that can be machine-evaluated) and adds a "context-shifted" version of the unpuzzle where the language, setting, or vocabulary is changed but the logical structure is preserved. These context-shifted trivial problems can be used to test if a model has the ability of solving the simpler problems, which in turn helps us examine the reasons behind failures to solve the unpuzzle problems. Figure 1 for an example; the unpuzzle only differs by two characters, and the solution logic of the unpuzzle and the context-shifted version are identical. ²

While many existing works evaluate LLM reasoning robustness by perturbing problems while maintaining the same difficulty level (Mirzadeh et al., 2024; McCoy et al., 2024), our study instead

²All data will be released.

shows that *decreasing* difficulty can also lead to much worse performance. A key failure mode we observe is that LLMs tend to "overthink" easy problems, often erroneously reusing reasoning steps corresponding to the more complex puzzle solutions — a phenomenon we term *reasoning delirium*. Further, these failures are not because the models do not know how to reason about easy problems: every model we tested performed better on the context-shifted unpuzzles than the original ones, indicating that failure was at least in part due to memorization of the original puzzle.

In summary, we make the following contributions: (1) we conduct a comprehensive evaluation of frontier LLMs across a wide range of simple reasoning problems; (2) we connect failure modes to their potential causes; (3) we present a new set of procedurally generated reasoning tasks with tunable parameters that are challenging for high-quality LLMs; (4) we introduce the UNPUZZLES dataset that confuses frontier LLMs, exposing memorization artifacts. Our work demonstrates that the qualitative trend of performance degradation still exists even for the latest thinking models, even though quantitative results have improved. We hope the new benchmarks and our methodology for identifying failures will improve the assessment of reasoning capabilities of future model generations.

2 Related work

108

109

110

111

112

113 114

115

116

117

118

119

120

121122123

124 125

126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

153 154

155

156

157

158

159

160

161

There is a long line of research focused on identifying tasks that challenge modern LLMs and developing new benchmarks. In this paper, we review the studies most relevant to the tasks we investigate. Transformer-based LLMs are known to struggle with seemingly simple tasks such as counting (Ouellette et al., 2023; Yehudai et al., 2024; Barbero et al., 2024) and copying (Liu et al., 2024; Barbero et al., 2024), due to issues related to tokenization, architecture, and embeddings. They also perform poorly on tasks requiring multi-step reasoning, such as arithmetic, logic puzzles, and dynamic programming (Dziri et al., 2024). The difficulty of solving simple logic problems has been explored in Yang et al. (2023); Parmar et al. (2024); Han et al. (2022), where these tasks are often framed as translation problems from natural language to first-order logic. Other works, such as Valmeekam et al. (2024a;b), construct planning benchmarks using Planning Domain Definition Language (PDDL), while Xie et al. (2024) develops a travel planning benchmark in real-world scenarios. These studies show that existing LLMs are far from saturating these datasets. Additionally, reasoning benchmarks with large amounts of irrelevant content have been proposed (Shi et al., 2023; Mirzadeh et al., 2024) to test models' long-context generalization capabilities. Most of these benchmarks are fixed and often combine the core challenge (e.g., logic or planning) with secondary challenges, such as understanding PDDL or real-world commonsense reasoning for travel, making it difficult to pinpoint the exact sources of failure. Furthermore, fixed benchmarks are difficult to extend or generalize and are prone to saturation or overfitting as LLMs improve. In contrast, our work takes a principled approach by simplifying problems to isolate failure causes. Our tasks are randomly and procedurally generated, allowing for easy adjustments to their distribution and difficulty (at a superficial level), ensuring they remain challenging for future LLMs. The work most relevant to ours is that of Opedal et al. (2025), which evaluates the Out-of-Distribution (OOD) generalization ability of LLMs through MathGAP, a framework that procedurally generates arithmetic problems by representing them as sequences of logical forms, with solutions structured as proof trees. Compared to Opedal et al. (2025), our work takes a broader perspective by examining a wider range of tasks and identifying multiple critical failure modes beyond OOD generalization. In a concurrent work to ours, Shojaee et al. (2025) evaluate LLMs on four puzzles with controllable "complexity" and show that the accuracy of all models completely collapses beyond a certain complexity threshold. It has been pointed out that their experiment design is somewhat flawed, including unsolvable problems and potentially ignoring token limits (Opus and Lawsen, 2025; Chan, 2025).³. Nonetheless, we observe qualitatively similar performance with accuracy decreasing as a function of the task tediousness.

The idea of perturbing existing benchmarks to test the robustness of LLM reasoning has been explored in several prior works. Mirzadeh et al. (2024) introduce a variant of the GSM8K benchmark for mathematical reasoning, modifying numerical values and adding irrelevant information, both of which lead to a performance drop in common models. Similarly, Jiang et al. (2024) evaluate LLMs on conjunction and syllogistic fallacies by perturbing well-known problems—changing names, inserting celebrity references, adding irrelevant content, and replacing quantifiers with synonyms—revealing

³While the work of Opus and Lawsen (2025) was initially published as a joke, some of the flaws discussed are legitimate concerns

evidence of "token bias" in LLMs. The negative effect of adding irrelevant context to math word problems was also shown recently by Xu et al. (2025). These studies primarily focus on perturbing original problems while maintaining or increasing their difficulty. In contrast, our UNPUZZLES benchmark takes the opposite approach: we make minimal edits to the wording but drastically *reduce* problem difficulty. A related work by Williams and Huckle (2024) introduces a benchmark of 30 easy problems that LLMs fail on, 12 of which are logical puzzles. Our evaluation on puzzles is considerably more comprehensive. Finally, the failure modes identified in UNPUZZLES also relate to findings from McCoy et al. (2024), which demonstrate that LLM accuracy is heavily influenced by the likelihood of task formulations, inputs, or outputs appearing in the training data.

3 PROCEDURALLY GENERATED REASONING TASKS

This section presents our collection of simple reasoning tasks, including several extensions of tasks from existing literature. Each task is procedurally generated, allowing a near-infinite number of new problems to be generated, and defined by parameters that control the difficulty or complexity. One of our goals was to design tasks that are easy (albeit tedious) for humans, but become unsolvable by frontier models when the difficulty parameters are large enough; all our results demonstrate this feature. For brevity, each task is described informally; full descriptions, usually with pseudocode, are in the appendix.

Throughout, we abbreviate Google's Gemini 1.5 Pro, 2.0 Flash, 2.0 Flash Thinking, and 2.5 Pro with G1.5, G2.0F, G2.0FT, and G2.5P, respectively. We also abbreviate Anthropic's Claude 3.5 and 3.7-sonnet (run without thinking tokens), OpenAI's o1 and GPT-4o, DeepSeek's R1, and Gemma 3 27B by C3.5, C3.7, o1, 4o, R1, and Gm3, respectively; see the appendix for the specific versions. Unless specified otherwise, for every task and every choice of parameters, we average the performance of the models across 20 randomly sampled tasks.

3.1 Tasks

Character and word counting Until somewhat recently, many LLMs infamously could not count the number of r's in "strawberry." This task extends this task to simultaneous word or character counting. The WORD COUNTING task requires the model to simultaneously count the number of occurrences of each word in a list of size k from a paragraph of length m. The task obviously becomes more difficult as k and m increase. The CHARACTER COUNTING task only requires counting a single character, which already proves difficult for the models. The paragraphs are extracted from the WikiText-2 dataset and are either selected to have minimum size m = 50 (with maximum size m = 150) or minimum size m = 150 (with maximum size m = 150).

First-order logic tasks We evaluate models on two fundamental logic tasks: evaluating propositional logical statements and negating first-order logical statements. A logic formula can be represented with a tree with logic operators as nodes and propositions and predicates as leaves. An atomic proposition is a simple, binary-valued variable, usually represented P or Q, whereas a predicate represents a property about an individual: for example, P(x) indicates that individual x has property P. We include the standard logical operators \vee , \wedge , \Leftrightarrow , \Rightarrow , \neg , $\forall x \in X$, and $\exists x \in X$, (respectively, or, and, equivalent, implies, negation, for all, and exists), where the last two are quantifying operators that require a domain to be specified. Exploiting the tree structure, we can sample a logic formula recursively. The complexity is controlled by the maximum depth d and the total number n of predicates and atomic propositions to sample for leaves. We either choose n=16 (16 predicates, 16 atomic propositions, and 8 domains) or n = 8 (8, 8, and 4, respectively). The final parameter is what vocabulary we use for the leaves: we created three categories: random 20 character strings, capital letters (reflecting the training data), and words that describe motion pictures. We consider the tasks of (1) LOGIC EVALUATION - identifying which of four value assignments evaluates to true, and (2) LOGIC NEGATION - identifying the negation of a logic formula from four options. See Appendix G.2 and G.3 for more details and examples.

Math word problems based on proof trees We extend the MATHGAP task of Opedal et al. (2025), which uses a tree-based representation of proofs to generate mathematical word problems. Each problem is represented as a sequence of *logical forms* under the formalism from Opedal et al. (2024).

Word Counting: Given	a
text paragraph, count the	oc-
currences of every word	in
a k-long list.	
Parameters: number	of
words to count k, minimu	m

paragraph size m.

	k	m	o1	4o	R1	Gm3	C3.7	G2.0F	G2.0FT	G2.5P
1	1	50	1.00	1.00	0.95	0.95	1.00	0.95	0.95	1.00
ı		150	1.00	1.00	0.90	0.65	0.80	0.90	0.90	1.00
ı	3	50	1.00	0.85	0.90	0.70	0.70	0.90	0.55	1.00
ı		150	0.95	0.50	0.55	0.15	0.50	0.65	0.05	1.00
ı	6	50	1.00	0.75	0.95	0.45	0.80	0.70	0.50	1.00
ı		150	0.95	0.30	0.35	0.05	0.40	0.25	0.00	0.90
l	m	(o1 4	40 F	R1 G1	m3 C3	3.7 G	2.0F G	2.0FT G	2.5P
J	50	1.0	00 0.	15 0.0	0.5	.10 0.	30	0.15	0.15	0.60
	15	0 1.0	00 0.	15 0.0	00 0.	.05 0.	25	0.15	0.10	0.45

Figure 2: **Top**: The pass@5 performance on the word counting task vs. the number of words to count k and minimum paragraph size m. **Bottom:** The pass@5 performance for the single character counting task vs minimum paragraph size m. While o1 performs well on word counting for the parameters in the table, it eventually fails with a sub 40% accuracy with $k \geq 3$ and $m \geq 2000$.

A logical form is a truth statement about the world, typically describing an arithmetic relationship, such as "Alice has 3 more apples than Bob." *Inference rules* can be used to prove new logical forms from existing ones. Problems are constructed by sampling a *Proof Tree* with logical forms as nodes, leaves as axioms, and a question as the root, before programmatically converting nodes to natural language. See Appendix H.4 for details. MATHGAP includes only four logical forms with one being non-commutative (*transfer*, e.g. "Alice gave Bob 5 apples"). We extend MATHGAP in two ways:

- **Diversity:** We increase the diversity of logical forms and inference rules by adding nine statement types, six of which are non-commutative. Examples include "A eats 5 apples", "A and B switch the apples they have". Such statements make it more difficult for the model to keep track of the intermediate states before computing the final answer. See Appendix H.1 and H.5 for the full list and an example. Problem parameters are tree depth d and inclusion of diverse logical forms.
- Irrelevant statements: We generate additional statements involving people irrelevant to the original problem and shuffle them into the original statements, such as "A is very generous and enjoys sharing food with others". See Appendix H.3 for the complete list. The problem parameters are the number of additional people and the number of additional sentences.

Travel Planning This task presents the model with a list of cities and various connecting modes of transit and asks the model to design a travel itinerary satisfying multiple constraints. This work is similar to that of Xie et al. (2024). For each task, we randomly generated a directed graph where the S nodes represent cities and the edges represent connections. Each edge carries a subset of A transportation modes, each with a randomly sampled cost. Based on this graph, we construct our travel planning problem, which consists of a word-based graph description and the constraints. The constraints include the starting and ending cities, a limit on the total travel cost, and N, the number of unique cities the traveler must visit. The problem parameters are S, A, and N. See Appendix I for further details.

3.2 RESULTS AND FAILURE ANALYSIS

Shortened problem descriptions, parameters, and evaluation results are shown in Figures 2, 3, 4, and 5. In most cases, increasing the "tediousness" of each task through the available parameters leads to a drop in performance. We summarize some of the common failure patterns here and refer the reader to the appendix for examples:

• Accumulation of errors. In general, if each computational step has a small probability of error, increasing the number of steps exponentially increases the probability of overall failure, even if the model follows the correct approach to calculating the solution. Examples of such failures are evident in the Counting task when increasing paragraph length, Logic and ProofTree problems when increasing tree depth, and Travel Planning when increasing the number of unique cities to visit. We emphasize that the tasks do not become more difficult, just more tedious, and hypothesize that models fail due to missing a piece of information, or making an error in one

Logic Evaluation: Given a propositional logic formula and four value assignments, identify which assignment evaluates to true.

Parameters: formula tree depth d, number of unique atomic propositions n.

d	n	o1	4o	R1	Gm3	C3.7	G2.0F	G2.0FT	G2.5P
4	16	1.00	0.75	1.00	0.95	0.98	0.90	0.70	1.00
								0.75	
8	16	0.73	0.50	0.77	0.30	0.27	0.27	0.45	0.78
							0.37		0.85
12	16	0.35	0.17	0.35	0.33	0.35	0.32	0.38	0.52
	8	0.43	0.35	0.38	0.25	0.38	0.33	0.30	0.55

Logic Negation: Given a propositional logic formula, identify its negation from four options.

Parameters: formula tree depth *d*, vocabulary for propositions,

predicates, and do-

mains.

d	names	o1	4o	R1	Gm3	C3.7	G2.0F	G2.0FT	G2.5P
4	letters	0.95	0.82	0.15	0.85	0.95	0.85	0.97	1.00
	movies	1.00	0.72	0.15	0.75	0.97	0.95	1.00	0.97
	random 20	1.00	0.80	0.15	0.70	1.00	0.95	0.95	0.95
8	letters	0.97	0.68	0.17	0.60	0.93	0.88	0.80	1.00
	movies	0.95	0.47	0.88	0.57	0.88	0.90	0.72	1.00
	random 20	0.90	0.55	0.90	0.47	0.82	0.95	0.53	0.95
12	letters	0.75	0.40	0.82	0.45	0.80	0.82	0.57	0.93
	movies	0.80	0.23	0.82	0.35	0.82	0.68	0.55	0.88
	random 20	0.62	0.38	0.68	0.30	0.62	0.82	0.55	0.90

Figure 3: **Top**: Accuracy for the logic evaluation task vs. tree depth d and number of possible unique predicates n. **Bottom**: Accuracy for the logic negation task vs. depth d and the vocabulary used for propositions, predicates, and domains (random 20 denotes random character strings of length 20).

ProofTree with diverse statements: Given a diverse set of logical statements, answer questions that require deduction sampled from a proof tree with a bounded depth and number of leaves. **Parameters:** max tree depth d, whether to include diverse logical forms, $\ell \in \{True, False\}$

d	ℓ	01	40	R1	Gm3	C3.7	G2.0F	G2.0FT	G2.5P
3	F	1.00	0.85	1.00	0.95	0.90	1.00	0.85 0.70	1.00
	T	0.90	0.95	0.90	0.65	0.95	0.65	0.70	
6	F	0.60	0.35	0.95	0.30	0.60	0.30	0.20	0.50
								0.20	
9	F	0.35	0.05	0.55	0.15	0.20	0.30	0.10	0.25
	T	0.55	0.15	0.55	0.15	0.15	0.35	0.15	0.35

ProofTree with irrelevant information: Answer proof tree questions that include irrelevant information. **Parameters:** max tree depth d, number of irrelevant people P, number of irrelevant sentences S.

P	S	o1	4o	R1	Gm3	C3.7	G2.0F	G2.0FT	G2.5P
1	0	0.50	0.58	0.75	0.35	0.50	0.25	0.40	0.60
	60	0.45	0.40	0.50	0.25	0.20	0.10	0.50	0.50
2	0	0.50	0.38	0.75	0.20	0.55	0.20	0.50	0.50
	60	0.45	0.26	0.45	0.15	0.15	0.20	0.30	0.40
4	0	0.40	0.16	0.65	0.15	0.20	0.35	0.35	0.35
	60	0.25	0.16	0.25	0.00	0.10	0.10	0.10	0.35
6	0	0.40	0.00	0.55	0.00	0.15	0.20	0.15	0.30
	60	0.30	0.00	0.35	0.05	0.00	0.05	0.20	0.35

Figure 4: Pass@5 scores for the proof tree tasks. **Top:** results for the diverse logic rules task, where we vary the depth d and whether the diverse rules are included. **Bottom:** results for the irrelevant sentences task, where we vary P, the number of irrelevant people, and S, the number of irrelevant sentences.

or more reasoning steps along the way. See Appendix H.7 for an example of o1 failing on a ProofTree problem due to an error in a reasoning step.

• Long context. Given a long context, models may struggle to sharply attend to relevant quantities, resulting in higher error rates. We observe this phenomenon most clearly in the case of ProofTree problems with irrelevant sentences, where models may overlook essential statements surrounded by irrelevant ones. Similarly, in word counting, models may overlook word instances in long paragraphs. Note that context length typically increases jointly with the parameters controlling the amount of computation.

324
325
326
327
328
328
329
Travel Planning: Create a travel itinerary using a city connection graph that adheres to a list of constraints.

Parameters: num. cities S, num. transportation modes

A, num. unique cities N

S	N	01	4o	R1	Gm3	C3.7	G2.0F	G2.0FT	G2.5P
10	5	1.00	0.34	0.90	0.05	0.70	0.00	0.10 0.00	0.60
	8	1.00	0.00	0.45	0.00	0.35	0.00	0.00	0.35
20	5	1.00	0.06	0.55	0.00	0.45	0.00	0.05	0.30
	8	0.75	0.00	0.05	0.00	0.10	0.00	0.00	0.15

Figure 5: Travel planning: pass@5 performance results. We always have A=4.

- Statistical shortcuts and educated guesses. Rather than executing computation, models sometimes prefer to exploit simplifications or take educated guesses. This often happens on logic tasks: models look for quick ways to eliminate choices, which is only effective for smaller problems. See Appendix G.4 for examples of simplifications and guesses on logic questions. We observe similar guessing behavior for all models in the Travel Planning task, where the models typically fail by (1) randomly sampling a few solutions and concluding that the problem is infeasible, or (2) proposing a solution with hallucinated parameters that satisfy constraints (see Appendix I.6 for examples). In all cases, increasing the problem size makes it more difficult for models to succeed through educated guessing.
- Poor state tracking. Some of our tasks require the models to track an increasingly complicated state, and we find that models fail as the size of that state increases. For example, while models can reliably count a single word, their performance drops considerably when simultaneously counting multiple words. Similarly, in ProofTree problems, introducing more complicated state changes and irrelevant agents (that are nonetheless tracked by the models) degrades performanc.
- Poor out-of-distribution (OOD) generalization. Changing the vocabulary in the logic tasks highlights errors due to poor OOD generalization. Universally, the performance is the best when the problem variables are single letters (which is likely the format of logical problems in the training data), and worst for random 20-character strings. We also attribute some of the ProofTree failures to poor OOD generalization, as some of the introduced statements (such as "A and B switch their apples") are not common in math word problems.
- Copying errors. We observe that logical reasoning can suffer from copying errors, which come in two flavors. First, when the formulas involve a smaller number of atomic propositions (n=8), the same atomic propositions appear in more subformulas, making those subformulas more easily confusable. Second, for formulas involving random 20-character variables, models will abbreviate those strings to single letters. This is related to OOD generalization, as the random token sequences are not likely to appear in training data.
- Tokenization. The character counting performance of all models is significantly lower than word
 counting, which suggests that tokenization is an issue for this task. Additionally, using random
 20-character strings as variable names in logic tasks likely results in multi-token variables, which
 may be more difficult for models to copy correctly in reasoning traces.

Overall, we show that LLM performance scales poorly in problem parameters related to the amount of computation and storage, even on problems which are self-contained and quite easy for humans.

4 Unpuzzles

We introduce the UNPUZZLES dataset, which consists of 97 well-known logical puzzles that are commonly found on the internet, and their *trivialized* versions, manually constructed by formulating textually similar questions that remove difficulty. While the puzzles typically require reasoning and background math knowledge, the answers to the unpuzzles are intended to be immediately obvious by common sense. See Appendix C for more details and dataset creation instructions and some examples. As we will show, all language models perform much better on the puzzles than on the unpuzzles, suggesting that they rely on memorized input patterns to generate answers rather than performing true logical reasoning.

Model	G1.5	G2.0F	G2.5P	Gm3	C3.5	C3.7	4o	o1	о3	R1
Puzzle										
Unpuzzle	17.5	38.1	62.9	34.0	27.8	48.5	19.6	59.8	74.2	59.8

Table 1: Percentage of correct answers on puzzles and unpuzzles.

Model	G1.5	G2.0F	G2.5P	Gm3	C3.5	C3.7	4o	o1
Context corruption (CC)	80	59	34	56	63	41	76	38
CC, correct	7	6	4	2	12	4	13	6
CC, incorrect with delirium	40	36	20	25	26	14	31	22
CC, incorrect (other)	33	16	10	29	25	23	32	10

Table 2: Number of unpuzzle solutions (out of 97) containing "context corruption." We further subcategorize corrupt solutions as (i) correct: leading to a correct final answer; (ii) incorrect with delirium: leading to an incorrect final answer with a solution that corresponds nearly exactly to the solution of the original puzzle; (iii) incorrect (other): leading to an incorrect final answer for other reasons. R1 and o3 are omitted since the answers we obtained often did not include full reasoning.

To provide further evidence of memorization, we created a dataset of *context-shifted* (CS) unpuzzles. Each CS unpuzzle is textually different from the corresponding unpuzzle but retains the same logical structure; that is, its answer is equally obvious. Performing poorly on an unpuzzle and well on a corresponding CS unpuzzle would suggest that the failure is due to the memorization of the puzzle text and solution, rather than inherent inability to reason about the problem. We generated CS unpuzzles automatically for a subset of 64 unpuzzles with simple numerical or categorical answers (as opposed to, e.g., puzzles asking for a strategy). We prompted models (o1 and Gemini 2.0 Flash) to rewrite each unpuzzle and change the language and setting, but keep the same logical structure and answer. We then verified and optionally edited the results. See Appendix C.2 for details. Figure 1 shows an example of a puzzle, unpuzzle, and CS unpuzzle.

4.1 EVALUATION

Correctness We generated the solution to each puzzle and unpuzzle independently using each model. We first verified whether the final answer to each is correct or not (regardless of whether the solution leading to the answer is correct). The evaluation was performed manually by four human annotators, since the answers to some puzzles are strategies rather than simple values. Each answer was assessed by a single annotator, or by consensus of all annotators if marked ambiguous.

Context corruption Next, we characterize the extent to which the poor performance on the unpuzzles is a consequence of memorization of the original puzzles. We define "context corruption" in an unpuzzle solution as erroneous or superfluous content (e.g. assumptions or reasoning steps) inappropriately recalled from the original puzzle or its solution. We evaluated each unpuzzle solution according to whether it contains context corruption or not. The most extreme behavior is when the models provided a solution that is nearly identical to the puzzle solution, sometimes without acknowledging that the unpuzzle is different – we call this category "delirium." We omit o3 and R1 as they often just responded with the final answer, making the degree of context corruption unclear, though their erroneous answers usually correspond to the answer to the original puzzle. This evaluation was performed by four human annotators and summarized in Table 2. We observe that memorization artifacts from the original puzzle and its solution are found in most cases, and even thinking models simply output the solution to the original puzzle about a fifth of the time. See Figure 6 for an example of context corruption in o1's solution to "Chameleons Go on a Date." See Appendix D for more illustrative examples as well as some amusing answers.

Context-shifted evaluation We evaluated models on the size-64 subset with corresponding CS unpuzzles. The results are shown in Table 3. We note that all models perform better on the context-shifted version of the dataset, which offers further evidence that the poor performance on the unpuzzles is due to the wording (and memorization of the original puzzles), rather than models' inherent inability to reason about the problems.

Model	G2.0F	G2.0FT	G2.5P	Gm3	C3.7	4o	o1	о3	R1
Puzzle Score	67	66	72	52	67	58	77	73	80
Unpuzzle Score	53	36	55	41	55	33	50	75	67
Context-shifted Score	70	48	66	50	63	52	59	80	73

Table 3: Percentage accuracy of all tested models on the original puzzles, the unpuzzles, and the context-shifted unpuzzles, for a subset of 64 problems with numerical or categorical answers. We see that every model performs better on the context-shifted unpuzzles than the unpuzzles, indicating that similarity to the puzzles degrades performance. In some cases, performance on the context-shifted unpuzzles is higher than for the original puzzles.

o1 solution to the Chameleons Go on a Date unpuzzle (excerpt)

A succinct way to see that they can all end up the same color (and in fact which color that must be) is via the following "invariant" argument: [...]

Define the quantity $S = P + 2Y \pmod{3}$. [...]

Since the initial $S = 0 \pmod{3}$, the only way to remain consistent with that invariant and end up with all chameleons of one hue is if they all become Maroon (because only that final state also gives S = 0). Conclusion: Yes. It is possible for all of the chameleons to end up the same color; specifically, they can all turn Maroon.

Figure 6: Chameleons go on a date (Figure 1): solution by o1. The answer to the original puzzle is 'no' through a proof by contradiction. The answer to the unpuzzle is 'yes'. o1 provides the correct final answer, but an incorrect solution; it finds that the contradiction does not hold and answers 'yes' without a constructive proof.

The UNPUZZLES dataset complements the procedural evaluations by providing another benchmark that is easy for humans and difficult for LLMs. It illustrates that the good performance of the models on the original (difficult) puzzles is at least in part a consequence of memorization of internet data, rather than true problem-solving abilities.

5 DISCUSSION

In a society that is increasingly utilizing frontier language models, understanding the capabilities and weaknesses of these models is becoming more and more important.ressive results across a variety of We have presented a comprehensive set of procedurally-generated parametric problems that are inherently easy (if tedious) for humans, and designed to assess LLM failures due to statistical shortcut learning, long context, long reasoning chains, and OOD generalization. As we demonstrate, these problems can be made difficult enough to make all SOTA LLMs fail. One suggestion from our paper is that LLMs should be evaluated not only by the most difficult problem they can solve, but also by the simplest problem they struggle with.

In addition, we have provided a small human-curated UNPUZZLES dataset of trivialized versions of math and logic puzzles commonly found on the internet. Our manual analysis shows that all models perform significantly worse on the unpuzzles than on the original puzzles, in most cases due to memorization of web data. Similarly to other recent works, our results suggest that LLMs mimic training data rather than performing true reasoning, making it relatively easy to find out-of-distribution problems where the models fail, and this problem is also present at the newest thinking models. This suggests that users remain careful when relying on the output of LLMs.

The main limitation of our work is that most of the experiments were run on closed-source models, which limits our ability to understand shortcomings beyond observing trends in the experiments and inspecting reasoning traces when available. We hope that our benchmarks will be useful in assessing and improving the reasoning capabilities of future generations of models.

REFERENCES

Federico Barbero, Andrea Banino, Steven Kapturowski, Dharshan Kumaran, João Guilherme Madeira Araújo, Alex Vitvitskyi, Razvan Pascanu, and Petar Veličković. Transformers need glasses!

- information over-squashing in language tasks. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
 - Lawrence Chan. Beware General Claims about "Generalizable Reasoning Capabilities" (of Modern AI Systems). https://www.lesswrong.com/posts/5uw26uDdFbFQgKzih/beware-general-claims-about-generalizable-reasoning, 2025. Accessed: 2025-07-07.
 - DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL https://arxiv.org/abs/2501.12948.
 - Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Sean Welleck, Peter West, Chandra Bhagavatula, Ronan Le Bras, et al. Faith and fate: Limits of transformers on compositionality. *Advances in Neural Information Processing Systems*, 36, 2024.
 - Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, et al. Folio: Natural language reasoning with first-order logic. *arXiv preprint arXiv:2209.00840*, 2022.
 - Bowen Jiang, Yangxinyu Xie, Zhuoqun Hao, Xiaomeng Wang, Tanwi Mallick, Weijie J Su, Camillo J Taylor, and Dan Roth. A peek into token bias: Large language models are not yet genuine reasoners. *arXiv* preprint arXiv:2406.11050, 2024.
 - Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024.
 - R Thomas McCoy, Shunyu Yao, Dan Friedman, Mathew D Hardy, and Thomas L Griffiths. Embers of autoregression show how large language models are shaped by the problem they are trained to solve. *Proceedings of the National Academy of Sciences*, 121(41):e2322420121, 2024.
 - Sean McLeish, Arpit Bansal, Alex Stein, Neel Jain, John Kirchenbauer, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, Jonas Geiping, Avi Schwarzschild, and Tom Goldstein. Transformers can do arithmetic with the right embeddings. In *Advances in Neural Information Processing Systems*, volume 38, pages 108012–108041, 2024.
 - Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models, 2016.
 - Ning Miao, Yee Whye Teh, and Tom Rainforth. Selfcheck: Using llms to zero-shot check their own step-by-step reasoning. *arXiv* preprint arXiv:2308.00436, 2023.
 - Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv preprint arXiv:2410.05229*, 2024.
 - Marianna Nezhurina, Lucia Cipolina-Kun, Mehdi Cherti, and Jenia Jitsev. Alice in wonderland: Simple tasks showing complete reasoning breakdown in state-of-the-art large language models. *arXiv* preprint arXiv:2406.02061, 2024.
 - Andreas Opedal, Alessandro Stolfo, Haruki Shirakami, Ying Jiao, Ryan Cotterell, Bernhard Schölkopf, Abulhair Saparov, and Mrinmaya Sachan. Do language models exhibit the same cognitive biases in problem solving as human learners? In *Proceedings of the 41st International Conference on Machine Learning*, pages 38762–38778, 2024.
 - Andreas Opedal, Haruki Shirakami, Bernhard Schölkopf, Abulhair Saparov, and Mrinmaya Sachan. MathGAP: Out-of-distribution evaluation on problems with arbitrarily complex proofs. In *The Thirteenth International Conference on Learning Representations*, 2025.
 - C Opus and A Lawsen. The illusion of the illusion of thinking: A comment on shojaee et al.(2025). *arXiv preprint arXiv:2506.09250*, 2025.
 - Simon Ouellette, Rolf Pfister, and Hansueli Jud. Counting and algorithmic generalization with transformers. *arXiv preprint arXiv:2310.08661*, 2023.

- Mihir Parmar, Nisarg Patel, Neeraj Varshney, Mutsumi Nakamura, Man Luo, Santosh Mashetty, Arindam Mitra, and Chitta Baral. Logicbench: Towards systematic evaluation of logical reasoning ability of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13679–13707, 2024.
 - Rod Pierce. About Math is Fun. https://www.mathsisfun.com/aboutmathsisfun.html. Accessed: October 2024.
 - Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H Chi, Nathanael Schärli, and Denny Zhou. Large language models can be easily distracted by irrelevant context. In *International Conference on Machine Learning*, pages 31210–31227. PMLR, 2023.
 - Parshin Shojaee, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity. *arXiv* preprint arXiv:2506.06941, 2025.
 - Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. Planbench: An extensible benchmark for evaluating large language models on planning and reasoning about change. *Advances in Neural Information Processing Systems*, 36, 2024a.
 - Karthik Valmeekam, Kaya Stechly, and Subbarao Kambhampati. Llms still can't plan; can lrms? a preliminary evaluation of openai's o1 on planbench. *arXiv preprint arXiv:2409.13373*, 2024b.
 - A Vaswani. Attention is all you need. Advances in Neural Information Processing Systems, 2017.
 - Sean Williams and James Huckle. Easy problems that Ilms get wrong. *arXiv preprint arXiv:2405.19616*, 2024.
 - Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, and Yu Su. Travelplanner: A benchmark for real-world planning with language agents. *arXiv* preprint arXiv:2402.01622, 2024.
 - Zikai Xie. Order matters in hallucination: Reasoning order as benchmark and reflexive prompting for large-language-models. *arXiv preprint arXiv:2408.05093*, 2024.
 - Xin Xu, Tong Xiao, Zitong Chao, Zhenya Huang, Can Yang, and Yang Wang. Can LLMs solve longer math word problems better? In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=C9ju8QQSCv.
 - Yuan Yang, Siheng Xiong, Ali Payani, Ehsan Shareghi, and Faramarz Fekri. Harnessing the power of large language models for natural language to first-order logic translation. *arXiv preprint arXiv:2305.15541*, 2023.
 - Gilad Yehudai, Haim Kaplan, Asma Ghandeharioun, Mor Geva, and Amir Globerson. When can transformers count to n? *arXiv preprint arXiv:2407.15160*, 2024.
 - Honghua Zhang, Liunian Harold Li, Tao Meng, Kai-Wei Chang, and Guy Van den Broeck. On the paradox of learning to reason from data. *arXiv preprint arXiv:2205.11502*, 2022.

A LICENCES FOR EXISTING ASSETS

A.1 MODELS

Below, we've tabulated the specific models and licences we have used.

- OpenAI The specific models we used from OpenAI are gpt-4o-2024-08-06, o1-2024-12-17, and o3-2025-04-16, which were abbreviated by 4o, o1, and o3 in the text. Terms of Use can be found at https://openai.com/policies/row-terms-of-use/.
- Anthropic We used claude-3-5-sonnet-20240620 and claude-3-7-sonnet-20250219, which were abbreviated C3.5 and C3.7. Terms of Service can be found at https://privacy.anthropic.com/en/articles/9190861-terms-of-service-updates.
- Gemma 3: We used the 27b-it model, which has open weights and permits responsible commercial use. Terms of Service are given at https://gemma3.app/terms-of-service.
- DeepSeek DeepSeek's R1 model and weights are licenced under the MIT licence DeepSeek-AI (2025).
 - Gemini The Gemini 2.0 Flash, 2.0 Flash Thinking, and 2.5 Pro had API names of gemini-2.0-flash-001, gemini-2.0-flash-thinking-exp and gemini-2.5-pro-exp-03-25. Terms of service can be found at https://ai.google.dev/gemini-api/terms.

A.2 DATA

We list the websites used to collect math and logic puzzles and their licences and terms of use below. Please see the released dataset for per-puzzle attributions.

- Wikipedia (https://www.wikipedia.org/): CC BY-SA 4.0 Creative Commons Attribution-ShareAlike 4.0 International https://creativecommons.org/licenses/by-sa/4.0/
- www.mathisfun.com.copyright Rod Pierce, cited as instructed on the website (Pierce)
- https://puzzles.nigelcoldwell.co.uk/copyright Nigel Coldwell.
- https://geeksforgeeks.org/, Terms of Use https://www.geeksforgeeks.org/legal/intellectual-property-rights-legal/

B PROCEDURAL LOGIC RESULTS WITH CONFIDENCE INTERVALS

We now include the procedural logic results including simple Gaussian error bars; these were omitted from the main body due to space constrains. In particular, results for OpenAI's o3 model are included. You can find the results in Figures 7-13.

m	o1	о3	40	R1	g3	C3.7	G2.0F	G2.0FT
		0.8 ± 0.18 0.45 ± 0.23						

Figure 7: Full results, with confidence intervals, for the Character Count task

C UNPUZZLING

This section provides more detail about the UNPUZZLES and their auto-evaluation.

C.1 Dataset creation instructions

The following are instructions given to humans to trivialize puzzles:

\overline{k}	m	o1	о3	40	R1	g3	C3.7	G2.0F	G2.0FT
1	50	1.0±0.00	0.95±0.10	1.0±0.00	0.95±0.10	0.95±0.10	1.0±0.00	0.95±0.10	0.95±0.10
	150	1.0 ± 0.00	1.0 ± 0.00	1.0 ± 0.00	0.9 ± 0.14	0.65 ± 0.22	0.8 ± 0.18	0.9 ± 0.14	0.9 ± 0.14
3	50	1.0 ± 0.00	0.95 ± 0.10	0.85 ± 0.16	0.9 ± 0.14	0.7 ± 0.21	0.7 ± 0.21	0.9 ± 0.14	0.55 ± 0.23
	150	0.95 ± 0.10	1.0 ± 0.00	0.5 ± 0.23	0.55 ± 0.23	0.15 ± 0.16	0.5 ± 0.23	0.65 ± 0.22	0.05 ± 0.10
6	50	1.0 ± 0.00	0.95 ± 0.10	0.75 ± 0.20	0.95 ± 0.10	0.45 ± 0.23	0.8 ± 0.18	0.7 ± 0.21	0.5 ± 0.23
	150	0.95 ± 0.10	1.0 ± 0.00	0.3 ± 0.21	0.35 ± 0.22	0.05 ± 0.10	0.4 ± 0.22	0.25 ± 0.20	0.0 ± 0.00

Figure 8: Full results, with confidence intervals, for the Word Count task

d	n	o1	о3	4o	R1	g3	C3.7	G2.0F	G2.0FT
4	medium	1.00±0.00	1.00 ± 0.00	0.75 ± 0.11	1.00 ± 0.00	0.95 ± 0.06	0.98 ± 0.03	0.90 ± 0.08	0.70 ± 0.12
	small	0.98 ± 0.03	1.00 ± 0.00	0.68 ± 0.12	0.98 ± 0.03	0.93 ± 0.06	0.97 ± 0.05	0.97 ± 0.05	0.75 ± 0.11
8	medium	0.73 ± 0.12	0.93 ± 0.06	0.50 ± 0.13	0.77 ± 0.11	0.30 ± 0.12	0.27 ± 0.12	0.27 ± 0.12	0.45 ± 0.13
	small	0.80 ± 0.10	0.98 ± 0.03	0.45 ± 0.13	0.78 ± 0.11	0.38 ± 0.13	0.28 ± 0.12	0.37 ± 0.13	0.42 ± 0.13
12	medium	0.35 ± 0.12	0.45 ± 0.13	0.17 ± 0.10	0.35 ± 0.12	0.33 ± 0.12	0.35 ± 0.12	0.32 ± 0.12	0.38 ± 0.13
	small	0.43 ± 0.13	0.33 ± 0.12	0.35 ± 0.12	0.38 ± 0.13	0.25 ± 0.11	0.38 ± 0.13	0.33 ± 0.12	0.30 ± 0.12

Figure 9: Full results, with confidence intervals, for the Logic Evaluation task

Task: Trivialize a puzzle Make a minimal edit to a well-known logical puzzle such that the solution becomes trivial. Either choose a puzzle from the given list or add a new puzzle. Suitable puzzles should be known to all language models, meaning that they readily provide you with the solution. Prefer puzzles where the solution is simply stated or can be checked with a simple question, for example one with a yes/no or an integer. Many famous puzzles can be modified to have simple solutions. Create an unpuzzle: modify the puzzle such that there is a trivial solution and the original solution is no longer necessary or even correct. Ideally, the simple question that verified the original puzzle should have a different answer. Check that large models still use the original solution to erroneously solve the modified puzzle or give the original (incorrect) answer. If not, repeat from step 3. Examples:

• Puzzle: There are 100 lockers in a row, all initially closed. A person walks down the row and opens every locker. Then, another person walks down the row and closes every second locker (starting from the second locker). Next, a third person walks down the row and changes the state (opens it if it's closed or closes it if it's open) of every third locker (starting from the third locker). This continues until 100 people have walked down the row. At the end, how many lockers are open?

Unpuzzle: There are 100 lockers in a row, all initially closed. A person walks down the row and opens every locker. Then, another person walks down the row and closes every second locker (starting from the second locker). At the end, how many lockers are open?

Explanation: The original puzzle requires that one finds the number of times each locker door's state is changed, which in turn requires the number of prime factors. This puzzle can be checked by asking a simple, integer-valued question. On the other hand, the unpuzzle has an obvious solution, as every second door is closed. The reasoning steps needed for the original puzzle are not required at all. (Gemini gives the same answer for both: 10)

• Puzzle: You have 12 coins, and one is counterfeit, being either heavier or lighter than the others. You have a balance scale and can use it three times. How can you identify the counterfeit coin and determine if it is heavier or lighter?

Unpuzzle: You have 12 coins, and they are all counterfeit. You have a balance scale and can use it three times. How can you identify all the counterfeit coins?

Explanation: The original puzzle requires careful reasoning through all possible results from the weighing. The unpuzzle has a laughably trivial solution. We could also modify the puzzle to ask "how many weighings are required to determine which is the counterfeit coin?".

d	name	o1	о3	4o	R1	g3	C3.7	G2.0F	G2.0FT
4	letters	0.95 ± 0.07	1.00 ± 0.00	0.82 ± 0.12	0.15 ± 0.11	$0.85 {\pm} 0.11$	0.95 ± 0.07	0.85 ± 0.11	0.97±0.05
	movies	1.00 ± 0.00	1.00 ± 0.00	0.72 ± 0.14	0.15 ± 0.11	0.75 ± 0.14	0.97 ± 0.05	0.95 ± 0.07	1.00 ± 0.00
	random 20	1.00 ± 0.00	0.97 ± 0.05	0.80 ± 0.13	0.15 ± 0.11	0.70 ± 0.15	1.00 ± 0.00	0.95 ± 0.07	0.95 ± 0.07
8	letters	0.97 ± 0.05	1.00 ± 0.00	0.68 ± 0.15	0.17 ± 0.12	0.60 ± 0.16	0.93 ± 0.08	0.88 ± 0.11	0.80 ± 0.13
	movies	0.95 ± 0.07	1.00 ± 0.00	0.47 ± 0.16	0.88 ± 0.11	0.57 ± 0.16	0.88 ± 0.11	0.90 ± 0.10	0.72 ± 0.14
	random 20	0.90 ± 0.10	0.97 ± 0.05	0.55 ± 0.16	0.90 ± 0.10	0.47 ± 0.16	0.82 ± 0.12	0.95 ± 0.07	0.53 ± 0.16
12	letters	0.75 ± 0.14	0.88 ± 0.11	0.40 ± 0.16	0.82 ± 0.12	0.45 ± 0.16	0.80 ± 0.13	0.82 ± 0.12	0.57 ± 0.16
	movies	0.80 ± 0.13	0.88 ± 0.11	0.38 ± 0.16	0.82 ± 0.12	$0.35{\pm}0.15$	0.82 ± 0.12	0.68 ± 0.15	0.55 ± 0.16
	random 20	0.62 ± 0.16	0.62 ± 0.16	$0.35{\pm}0.15$	0.68 ± 0.15	0.30 ± 0.15	0.62 ± 0.16	0.82 ± 0.12	0.55 ± 0.16

Figure 10: Full results, with confidence intervals, for the Logic Negation task

d	diverse	o1	о3	4o	R1	g3	C3.7	G2.0F	G2.0FT
3	False	1.00 ± 0.00	0.70 ± 0.21	$0.85 {\pm} 0.16$	1.00 ± 0.00	0.95 ± 0.10	0.90 ± 0.14	1.00 ± 0.00	0.85 ± 0.16
	True	0.90 ± 0.14	0.60 ± 0.22	0.95 ± 0.10	0.90 ± 0.14	0.65 ± 0.22	0.95 ± 0.10	0.65 ± 0.22	0.70 ± 0.21
6	False	0.60 ± 0.22	0.55 ± 0.23	0.35 ± 0.22	0.95 ± 0.10	0.30 ± 0.21	0.60 ± 0.22	0.30 ± 0.21	0.20 ± 0.18
	True	0.70 ± 0.21	0.90 ± 0.14	0.25 ± 0.20	0.65 ± 0.22	0.15 ± 0.16	0.55 ± 0.23	0.30 ± 0.21	0.20 ± 0.18
9	False	0.35 ± 0.22	0.75 ± 0.20	0.05 ± 0.10	0.55 ± 0.23	0.15 ± 0.16	0.20 ± 0.18	0.30 ± 0.21	0.10 ± 0.14
	True	0.55 ± 0.23	0.55 ± 0.23	0.15 ± 0.16	0.55 ± 0.23	0.15 ± 0.16	0.15 ± 0.16	0.35 ± 0.22	0.15 ± 0.16

Figure 11: Full results, with confidence intervals, for the MathGap Diverse task

C.2 CONTEXT-SHIFTED UNPUZZLES

We generated the context-shifted unpuzzles by first identifying a subset of 64 unpuzzles with simple categorical or integer answers (e.g. asking "what is the minimum number of crossings?" instead of "How can we move all items across the river?"). We used the following method for automatically shifting the context for the unpuzzles

- 1. We prompt a strong model with "I will give you a puzzle and a solution. I would like you to provide a single rewrite of the puzzle that changes the language and setting but keeps the logical structure and the answer the same; think carefully, highlighting the logical structure present in the puzzle," followed by a templated response specifying the domain the answer must lie in (the categories or an integer).
- 2. We verify that the new puzzle has the same solution as the original unpuzzle. If not, return to step 1).
- 3. We query the same model with the new unpuzzle; if the correct answer is not returned, return to step 1).
- 4. Verify that the context-shifted puzzle has the correct logical structure.

We found that models differed on the unpuzzles they could context-shift successfully, so we recommend using a few models simultaneously (we used o1 and Gemini 2.5 Flash). Of the context-shifted unpuzzles produced this way, 75% required minimal or no modification. One could use this method to generate large numbers of context-shifted puzzles.

C.3 AUTO-EVALUATION

Prompting models to disambiguate between the different levels of delirium is difficult. However, we had some success automatically evaluating correctness of the unpuzzle solution if we have access to a ground-truth unpuzzle solution.

the first question: is the unpuzzle solution correct or not? Our approach involved asking the model two questions. The first (following Miao et al. (2023)) asks a critic model whether correct solution "supports," "contradicts," or "is not directly related to" the model's response. The second presents the unpuzzle with the correct solution and asks whether the model's response had different reasoning, regardless of its correctness (we frequently saw that models would say that any reasoning not aligning with the original puzzle's solution was incorrect). We only conclude that the model's response

P	S	01	03	4o	R1	g3	C3.7	G2.0F	G2.0FT
1	0	0.50 ± 0.23	0.60 ± 0.22	0.30 ± 0.21	0.75 ± 0.20	0.35 ± 0.22	0.50 ± 0.23	0.25 ± 0.20	0.40 ± 0.22
	60	0.45 ± 0.23	0.45 ± 0.23	0.10 ± 0.14	0.50 ± 0.23	0.25 ± 0.20	0.20 ± 0.18	0.10 ± 0.14	0.50 ± 0.23
2	0	0.50 ± 0.23	0.55 ± 0.23	0.20 ± 0.18	0.75 ± 0.20	0.20 ± 0.18	0.55 ± 0.23	0.20 ± 0.18	0.50 ± 0.23
	60	0.45 ± 0.23	0.40 ± 0.22	0.05 ± 0.10	0.45 ± 0.23	0.15 ± 0.16	0.15 ± 0.16	0.20 ± 0.18	0.30 ± 0.21
4	0	0.40 ± 0.22	0.60 ± 0.22	0.05 ± 0.10	0.65 ± 0.22	0.15 ± 0.16	0.20 ± 0.18	0.35 ± 0.22	0.35 ± 0.22
	60	0.25 ± 0.20	0.40 ± 0.22	0.05 ± 0.10	$0.25{\pm}0.20$	0.00 ± 0.00	0.10 ± 0.14	0.10 ± 0.14	0.10 ± 0.14
6	0	0.40 ± 0.22	0.60 ± 0.22	0.10 ± 0.14	0.55 ± 0.23	0.00 ± 0.00	0.15 ± 0.16	0.20 ± 0.18	0.15±0.16
	60	0.30 ± 0.21	0.30 ± 0.21	0.05 ± 0.10	$0.35{\pm}0.22$	0.05 ± 0.10	0.00 ± 0.00	0.05 ± 0.10	0.20 ± 0.18

Figure 12: Full results, with confidence intervals, for the MathGap Irrelevant task

S	N	o1	о3	4o	R1	g3	C3.7	G2.0F	G2.0FT
10	5	0.95±0.10	0.55 ± 0.23	0.05 ± 0.10	0.90 ± 0.14	0.05 ± 0.10	0.45 ± 0.23	0.00 ± 0.00	0.00 ± 0.00
	8	0.65 ± 0.22	$0.65{\pm}0.22$	0.00 ± 0.00	0.45 ± 0.23	0.00 ± 0.00	0.15 ± 0.16	0.00 ± 0.00	0.00 ± 0.00
20	5	0.75 ± 0.20	0.70 ± 0.21	0.00 ± 0.00	0.55 ± 0.23	0.00 ± 0.00	0.25 ± 0.20	0.00 ± 0.00	0.05 ± 0.10
	8	0.50 ± 0.23	0.75 ± 0.20	0.00 ± 0.00	0.05 ± 0.10	0.00 ± 0.00	0.05 ± 0.10	0.00 ± 0.00	0.00 ± 0.00

Figure 13: Full results, with confidence intervals, for the Travel task

is correct if the two questions were answered "supports" and "no;" the prompt details are in the appendix.

Each row of figure 14 shows the performance of using the given model as a critic to evaluate the responses from every model: each cell gives (false positive rate, false negative rate), where positive means the unpuzzle solution is correct. In general, the false negative rates were significantly lower than the false positive rates, so autoevaluation gives a conservative estimate of performance. Also noteworthy is the complete lack of symmetry: o1 is much better at judging than being judged, though overall Claude seems to make the best critic. We hope that the autoevals have enough fidelity to allow the unpuzzles to be used for model improvement. Given the original unpuzzle, the correct unpuzzle_solution, and the model's response, we prompted the model twice with the following question:

```
Autoevaluation prompt template 1

Here is a simple question:
{unpuzzle}

This simple question has the simple solution:
<correct solution>
{unpuzzle_solution}
</correct solution>

Is the reasoning in the following solution different,
regardless of it is correct or not?
<second solution>
{response}
</second solution>
Please answer with <answer>yes or no</answer>
```

The second question is the following:

	Model								
Critic	G1.5	C3.5	4o	o1					
G1.5	(35.3, 2.4)	(32.0, 4.1)	(38.9, 1.2)	(36.7, 7.7)					
C3.5	(11.8, 0.0)	(24.0, 8.1)	(27.8, 3.7)	(40.0, 5.1)					
4o	(64.7, 0.0)	(80.0, 1.4)	(72.2, 0.0)	(78.3, 5.1)					
o1	(23.5, 1.2)	(28.0, 6.8)	(22.2, 2.5)	(43.3, 7.7)					

Figure 14: (False positive, False Negative) percentages for autoevaluation. Each row corresponds to using a different critic model for evaluation.

Autoevaluation prompt template 2 The following are 2 solutions to a simple puzzle. Solution 1: {unpuzzle_solution} Solution 2: {response} Compare the key points from both solutions step by step and then check whether Solution 1 'supports', 'contradicts' or 'is not directly related to' the conclusion in Solution 2. Please answer with <answer>'supports', 'contradicts', or 'is not directly related to' </answer>

We only declared the model response as correct if we received answers "no" and "supports."

D UNPUZZLES: EXAMPLES OF CONTEXT CORRUPTION

We give examples of answers containing various degree of context corruption. As an example of reasoning delirium, we show the solution of o1 to the Conjunctive Fallacy problem (Figure 15) in Figure 16. As an example of context corruption leading to a correct final solution, we show o1's full solution to Chameleons Go on a Date (Figure 1) in Figure 17. As an example of context corruption with an incorrect final solution, we show G1.5's full solution to the same problem in Figure 18.

Conjunctive Fallacy (original puzzle)

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more probable? A) Linda is a bank teller. B) Linda is a bank teller and is active in the feminist movement. **Answer**: A (conjunctive fallacy)

Conjunctive Fallacy (unpuzzle)

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. She is active in the feminist movement. Which is more probable? A) Linda is a bank teller. B) Linda is active in the feminist movement. **Answer**: B (given)

Figure 15: Conjunctive fallacy: puzzle and corresponding unpuzzle.

o1's Solution to Conjunctive Fallacy

This question is a version of the well-known "Linda problem," originally posed by Tversky and Kahneman to illustrate the conjunction fallacy. In the classic form, the two options are: A) Linda is a bank teller. B) Linda is a bank teller and is active in the feminist movement. The correct answer is always option A (that she is simply a bank teller), because any single event (being a bank teller) must be at least as probable as that same event combined with another (being both a bank teller and a feminist). Adding more conditions (in this case, being a feminist as well) can never make an event more likely, even if that extra detail sounds more "representative" of Linda's personality.

Why it feels counter-intuitive (the "conjunction fallacy"): • People often choose the more "representative" description—bank teller plus feminist—because it fits Linda's profile better.

• However, mathematically, $P(A \text{ and } B) \le P(A)$. Hence, it cannot be more probable to be both a bank teller and a feminist than it is to be just a bank teller.

In short, the more probable statement is simply "Linda is a bank teller."

Figure 16: o1 solution to the Conjunction Fallacy problem, illustrating complete context corruption, or delirium. The model provides the solution to the original problem, acknowledging the difference only by referring to the "classic form" of the problem.

o1's solution to Chameleons Go on a Date

A succinct way to see that they can all end up the same color (and in fact which color that must be) is via the following "invariant" argument:

1) Label the number of Purple, Yellow, and Maroon chameleons by P, Y, M, respectively. 2) Notice that in any "meeting" of two different-colored chameleons, P, Y, M change according to one of these three rules:

```
P+Y \rightarrow 2M, so (P,Y,M) becomes (P-1,Y-1,M+2). Y+M \rightarrow 2P, so (P,Y,M) becomes (P+2,Y-1,M-1). P+M \rightarrow 2Y, so (P,Y,M) becomes (P-1,Y+2,M-1).
```

- 3) Define the quantity $S = P + 2Y \pmod{3}$. One checks that under each of the above transformations, S remains unchanged (that is, S is an invariant modulo 3).
- 4) At the start, we have P=15, Y=15, M=17. Then $S=15+2*15=15+30=45=0 \pmod{3}$.
- 5) If eventually all chameleons turn into a single color, say all Purple (P=47) or all Yellow (Y=47) or all Maroon (M=47), then we can compute S in each scenario:
- All Purple: $P = 47 \implies S = 47 + 2 \cdot 0 = 47 = 2 \pmod{3}$.
- All Yellow: $Y = 47 \implies S = 0 + 2 \cdot 47 = 94 = 1 \pmod{3}$.
- All Maroon: $M = 47 \implies S = 0 + 2 \cdot 0 = 0 = 0 \pmod{3}$.

Since the initial $S = 0 \pmod{3}$, the only way to remain consistent with that invariant and end up with all chameleons of one hue is if they all become Maroon (because only that final state also gives S = 0).

Conclusion: Yes. It is possible for all of the chameleons to end up the same color; specifically, they can all turn Maroon.

Figure 17: o1's solution to Chameleons Go on a Date, illustrating context corruption with a correct final answer

E UNPUZZLES MODEL RESPONSE HIGHLIGHTS

We saw many "creative" responses for the UNPUZZLES. Figures 19 and 20 show amusing examples involving socks, ducks, and chickens.

Occasionally, a model will get delirious due to a different puzzle. Figure 21 shows an example of o1 copying a solution not from the original puzzle, but from an unrelated puzzle that also happens to involve socks.

Gemini 1.5's Solution to Chameleons Go on a Date unpuzzle

Let p, y, and m be the number of purple, yellow, and maroon chameleons, respectively. Initially, we have p = 15, y = 15, and m = 17.

When two different-colored chameleons meet, they both turn into the third color. This means the changes in the numbers of chameleons are always in increments of 2. We look at the remainders when the numbers are divided by 3.

Initially, we have $p \equiv 0 \pmod{3}$, $y \equiv 0 \pmod{3}$, and $m \equiv 2 \pmod{3}$.

Consider the remainders when dividing the number of each color by 3 at each step. Let $p \equiv p' \pmod{3}$, $y \equiv y' \pmod{3}$, and $m \equiv m' \pmod{3}$. The possible values for (p', y', m') are (0, 0, 2).

If a purple and a yellow chameleon meet, we have (p-1,y-1,m+2). Then the remainders are (2,2,1). If a purple and a maroon chameleon meet, we have (p-1,y+2,m-1). Then the remainders are (2,2,1). If a yellow and a maroon chameleon meet, we have (p+2,y-1,m-1). Then the remainders are (2,2,1).

We can notice a pattern. The remainders after any interaction will always satisfy $p'+y'+m'\equiv 2\pmod{3}$. The possible remainder sets are: (0,0,2),(0,2,0),(2,0,0),(1,1,0),(1,0,1),(0,1,1),(2,2,1),(2,1,2),(1,2,2). We need all chameleons to have the same color. This means two of the colors must be zero. The only combinations with two zeros are (0,0,2),(0,2,0), and (2,0,0). We start with (0,0,2). If all the chameleons are purple, we would have (1,0,0), which has a remainder sum of 1, which is not possible. If all the chameleons are yellow, we would have (0,1,0), which has a remainder sum of 1, which is not possible. If all the chameleons are maroon, we would have (0,0,1), which has a remainder sum of 1, which is not possible.

Since all final states are not achievable, it is not possible for all chameleons to have the same

My final is: no.

Figure 18: G1.5's solution to Chameleons Go on a Date, illustrating context corruption leading to an incorrect final answer.

Sometimes the models seem aware that the problem is different, but they seem pretty intent on sticking to the training distribution and make up flimsy excuses to rehash the original solution, as in the River Crossing puzzle shown in Figure 22.

Matching Socks 2 (original puzzle)

You are about to leave for holiday, but you forgot socks! You race back to your room, but the power is off so you can't see sock colors. Never mind, because you remember that in your drawer there are ten pairs of identical green socks, ten pairs of identical black socks, and eleven pairs of identical blue socks, but they are all mixed up. How many of your socks do you need to take before you can be sure to have at least one pair matching in color? **Answer:** 4 (by the pigeonhole principle)

Matching Socks 2 (unpuzzle)

You are about to leave for holiday, but you forgot socks! You race back to your room, but the power is off so you can't see sock colors. Never mind, because you remember that in your drawer there are ten pairs of green socks, but they are all mixed up. All socks have an identical shape (there is no distinction between socks for left and right feet). How many of your socks do you need to take before you can be sure to have at least one pair matching in color? **Answer:** Since all socks are green, taking two clearly suffices.

Gemini 1.5's Solution

The worst-case scenario is that you pick one sock from each of the ten pairs. This means you could pick ten socks and still not have a matching pair. However, the very next sock you pick *must* create a pair with one of the socks you already have.

So you need to pick 10 + 1 = 11 socks.

My final answer is: 11

Figure 19: Matching Socks 2: puzzle, unpuzzle, and an amusing answer by Gemini 1.5.

Counting Ducks (original puzzle)

There are two ducks in front of a duck, two ducks behind a duck and a duck in the middle. How many ducks are there? **Answer:** 3

Matching Socks 2 (unpuzzle)

There are two ducks in front of the chicken, two ducks behind the chicken, and a duck to the left of the chicken. How many ducks are there? **Answer:** 5

o1's Solution

A classic way to see why the answer is three is to realize that "the chicken" named in the riddle is actually just one of the three ducks (it has the nickname "Chicken"). Line the three ducks up:

- Duck 1 (sometimes called "Chicken"),
- Duck 2,
- Duck 3.

From the perspective of each duck, you can say there are "two ducks in front of it," "two ducks behind it," and so on—yet you never actually need more than three ducks. The riddle's wording double-counts the same birds from different points of view. Hence, the total number of ducks is 3.

Figure 20: Counting ducks: puzzle, unpuzzle, and an answer by o1.

F DETAILS FOR COUNTING

The counting tasks both use data from the Salesforce/wikitext and wikitext-2-v1 dataset Merity et al. (2016) hosted on Hugging Face. We will briefly describe each of them and provide task cards.

More Socks (original puzzle)

A man has 53 socks in his drawer: 21 identical blue, 15 identical black and 17 identical red. The lights are out and he is completely in the dark. How many socks must he take out to make 100 percent certain he has at least one pair of black socks? **Answer:** 38

More Socks (unpuzzle)

A man has 53 socks in his drawer: 21 identical blue, 15 identical black and 17 identical red. There is no distinction between left and right socks. The lights are out and he is completely in the dark. How many socks must he take out to have one sock for each foot? **Answer:** 2

o1's Solution

 A straightforward way to see why the answer is 4 is to use the "worst-case" or pigeonhole principle argument:

- There are three colors (blue, black, red).
- If you pick 3 socks, in the worst case you could have one of each color (1 blue, 1 black, 1 red) and thus no matching pair.
- As soon as you pick a 4th sock, by the pigeonhole principle at least two of them must be the same color.

Hence, you must take out 4 socks to guarantee having a matching pair for your two feet.

Figure 21: More Socks: puzzle, unpuzzle, and an answer by o1.

F.1 CHARACTER COUNTING

The character counting task is defined by paragraph bounds m_l and m_u . We randomly choose a wikitext snippit from all paragraphs that adhere to the length limits. We then count all the characters and randomly pick one in the top 10 most frequent.

Character counting prompt template

```
I will provide you a block of text. Please count the number of times
the character "{sampled_char}" appears in the text.
Give your answer using the format:

"The character appears #your answer# times."

Think step by step.
Here is the text.
{sampled_paragraph}
```

Riv

River Crossing with a Weight Constraint (original puzzle)

Four people (Alex, Brook, Chris and Dusty) want to cross a river in a boat that can only carry 100kg. Alex weighs 90kg, Brook weighs 80kg, Chris weighs 60kg and Dusty weighs 40kg, and they have 20kg of supplies. What is the minimum number of river crossings they need to get across?

River Crossing with a Weight Constraint (unpuzzle)

Four people (Alex, Brook, Chris and Dusty) want to cross a river in a boat that can only carry 300kg. Alex weighs 90kg, Brook weighs 80kg, Chris weighs 60kg and Dusty weighs 40kg, and they have 20kg of supplies. What is the minimum number of river crossings they need to get across? **Answer:** 1

Claude 3.5's Solution

Let's approach this step-by-step:

- 1. First, we need to consider the total weight: 90 + 80 + 60 + 40 + 20 = 290 kg This is within the 300kg limit, but they can't all go at once because the boat can only carry 300kg.
- 2. The most efficient way to cross is to always have the boat full or nearly full on each trip across, and to send the lightest person back with the boat.
- 3. Here's an efficient crossing strategy: Trip 1: Chris (60kg) and Dusty (40kg) cross (100kg total) Trip 2: Dusty returns alone Trip 3: Alex (90kg) and Brook (80kg) cross (170kg total) Trip 4: Chris returns alone Trip 5: Chris (60kg), Dusty (40kg), and supplies (20kg) cross (120kg total)
- 4. This strategy requires 5 crossings in total.
- 5. We can't reduce this number because: We need at least 2 trips to get everyone across (the boat must return at least once) We can't fit everyone in just two trips because Alex and Brook together are already 170kg, leaving only 130kg for Chris, Dusty, and the supplies, which isn't enough Therefore, 5 is the minimum number of crossings required.

My final answer is: 5

Figure 22: River Crossing with a Weight Constraint: puzzle, unpuzzle, and an answer by Claude 3.5.

F.2 WORD COUNTING

For the word counting task, we begin the same way by sampling a paragraph that obeys the length restrictions. We then compute the word frequencies, always asking the model to find the top k most frequent words. Because the wikitext data have white spaces around each word and all characters are lower case, each word always has the same tokenization.

The prompt template is give below.

```
I will provide you a block of text. Please count the number of
times each word in the list [word 1, word 2,...,word k] appears
in the text.
Give your answer using the format:
"The words appear [ your answer for the first word ,
your answer for the second word , ... ] times."
Think step by step.
Here is the text
{text}
```


G DETAILS FOR THE LOGIC TASKS

This section provides pseudocode for generating tasks for logic evaluation and logic negation tasks.

A logic formula can be represented by a tree where nodes are logical operators and leaves are atomic propositions. The nodes have a certain truth value depending on the value of their children. The

standard nodes have three types: connective, unitary, and quantifying. Connective nodes have two children (left and right), and unitary and quantifying nodes have one child. Throughout this section, T and F denote True and False, respectively.

There are two types of leaves:

Atomic propositions (often denoted by single capital letters, e.g., P, Q, etc.) are either true
or false.

 • Predicates represent a property about an individual. For example, for predicate P, we have P(x) = T if the individual x has the property P. We expect P(x) to have different values as x changes.

 There are seven operators, described in the following table (other logical primitives, e.g., the exclusive or, may be derived from the ones below).

Name	Symbol	Type	Description
and	Λ	connective	True if both children are True
or	V	connective	True if at least one child is True
implies	\rightarrow	connective	Only False if $T \to F$
equals	\Leftrightarrow	connective	True if the left and right child are equal
not	_ ¬	unitary	The opposite value of its child
universal quantification	$\forall x \in X$	quantifying	True if the child evaluates to True for
			every value x in domain X .
existential quantification	$\exists x \in X \text{ s.t.}$	quantifying	True if the child evaluates to True for
		- •	some value x in domain X .

Nodes of connective and unitary types are only defined by their symbol. A quantifying node is defined by its symbol and the domain it operates on. For simplicity, we will simply number the possible domains, e.g., D_1, D_2, \ldots

The first step in constructing a logic task is to sample a logic formula. We describe how in the next section.

G.1 SAMPLING A LOGIC FORMULA

Including first-order logic requires a sampling procedure that ensures the domains have scopes that make sense. In particular, the domain of a predicate must be from one of its ancestors. To enforce this, we keep track of every used domain in each subtree and limit the domains of predicates to these domains. Once we finish sampling a subtree with a root quantifying node, we then check if the subtree actually used the domain of the root. If not, the quantifying node is removed.

The logic problems were also parameterized by the number of unique propositions, n. For n=8, we also chose the number of unique predicates and domains to be 8 and 4, respectively. For n=16, the number of unique predicates and domains were 16 and 8, respectively.

We use several different sets of names for the propositions, predicates, and domains. They include

• common letters;

random 20 character-long lower case strings;
words about movies.

For generating each prompt, a subset of the appropriate size was selected from larger sets. For example, the "movie" vocabulary uses the following words:

• Propositions: dark, dramatic, intense, thrilling, suspenseful, romantic, comedic, tragic, epic, inspiring, thought-provoking, emotional, powerful, beautiful, visually-stunning, artistic, creative, imaginative, innovative, classic, mainstream, independent, foreign, animated, biographical, historical, fictional, realistic, surreal, abstract

 Predicates: has_subtitles, is_streamable, is_theatrical_release, is_direct_to_video, is_part_of_franchise, has_sequel, has_prequel, is_remake, is_based_on_book,

is_based_on_true_story, is_animated, uses_cgi, uses_stop_motion, is_live_action, is_musical, is_comedy, is_drama, is_horror, is_action, is_sci_fi, is_fantasy, is_romance, is_thriller, is_documentary, is_historical_fiction, is_independent_film, is_big_budget, won_awards, has_famous_actors, has_original_score,, is_award_winning,

• Domains: action_movies, comedies, period_pieces, science_fiction_films, fantasy_films, horror_films, thrillers, dramas, romantic_comedies, romantic_dramas, musicals, westerns, crime_films, war_films, documentaries, biopics, animated_films, adventure_films, mystery_films, superhero_films.

Algorithm 1: Sampling a First-order Logic Formula

```
Data: maximum depth d_{\max} probability of deepening tree p_d probability of sampling a connective node p_c probability of sampling a unitary node p_u probability of sampling a quantifying node p_q the number of unique atomic propositions N_a the number of unique domains N_d Call the helper function Algorithm 3 with \mathcal{D} = \emptyset, d = 0;
```

Algorithm 2: Sampling a node

```
Data: Probabilities p_c, p_u, p_q of sampling a connective, unitary, or quantifying node List of previously used domains \mathcal{D}
Number of unique domains N_D
```

```
Number of unique domains N_D
```

```
if |\mathcal{D}| = N_d then | Choose node from (connective, unitary) with probabilities proportional to (p_c, p_u); else | Choose node from (connective, unitary, quantifying) with probabilities to (p_c, p_u, p_q); if node is connective then | Choose operator from (\land, \lor, \to, \Leftrightarrow) with probabilities (.3, .3, .3, .1); else if node is unitary then | Set operator to be \neg.; else | /* node is quantifying | */ Choose operator from (\lor, \exists) with equal probability.; | Choose new domain uniformly from \{1, \ldots, N_d\} \setminus D;
```

return operator, new domain

G.2 CONSTRUCTING THE LOGIC EVALUATION TASK

For every task, we first sample a logic formula with $p_q=0$, i.e. without quantifying nodes. We use $p_d=.8$, $p_c=.85$ and $p_u=.15$; that is, we only choose an atomic proposition 20% of the time (unless we must to adhere to the maximum depth), and of the remaining 80%, we choose a connective node 85% of the time and a \neg operator 15%. After sampling the formula, the names for all atomic propositions are chosen from a name set as described above. We then sample random value assignments for all atomic propositions until we find one that evaluates to true and three that evaluate to false. These are then presented in random order using the following prompt template. Some models (notably 01) have restrictions on the language you can use to prompt the model. In that case "think carefully step-by-step and" was removed from the last sentence.

```
1242
         Algorithm 3: Sampling Helper function
1243
         Data: maximum depth d_{\max}
1244
                 current depth d
1245
                 List of previously used domains \mathcal{D}
1246
                 Number of unique domains N_d
1247
                 Probability of going deeper p_d
1248
                 List of atomic propositions \mathcal{L}_{prop}
1249
                 List of predicates \mathcal{L}_{pred}
1250
1251
         Sample U \sim \text{Uniform}[0, 1];
         if d = d_{\max} or U \ge p_d then
1252
             if With probability 50% then
1253
                  return An atomic predicate uniformly from \mathcal{L}_{prop}
             else
1255
                  Sample a predicate uniformly from \mathcal{L}_{pred};
1256
                  Sample a domain uniformly from \mathcal{D};
1257
                  return the predicate over the domain
         else
1259
             Chose node N, with domain D_{new} if N is quantifying, using Algorithm 2;
1260
              For each child of N, sample using this algorithm with d = d + 1, \mathcal{D} = \mathcal{D} \cup \{D_{new}\};
1261
             if N is quantifying and D_{new} was not used by the descendants of N then
1262
                 return the child of N
1263
             else
1264
               \perp return N
1265
```

```
You are a logic student. I will give you a logical formula, written in propositional logic, as well as four options for values of every atomic proposition in the formula.

Logical formula: {formula}

Which of the following choices makes the logical formula evaluate to True?

A: {answer 1}
B: {answer 2}
C: {answer 3}
D: {answer 4}

Please think carefully step-by-step and provide your answer with <answer>A, B, C, or D</answer>.
```

G.3 CONSTRUCTING THE LOGIC NEGATION TASK

Similar to the Logic Evaluation Task, the negation task samples a logic formula with $p_c=.6$, $p_u=p_q=.2$ and all other sampling parameters the same. We then compute the negation using the standard rules for first-order logic, assign it to a random choice, then perturb the correct answer to arrive at the three incorrect choices.

We perturb a logic formula by selecting, uniformly at random, a single node or leaf of the tree; the perturb operation depends on node type.

• **Proposition**: we create a list of all propositions in the formula, append a new, unused proposition (so long as the total number of propositions satisfies the constraints of the problem), then replace the proposition from the list uniformly at random.

1297

1299

1300 1301

1302 1303

1304

1305

1306

1322

1323 1324

1325 1326

1327

1329

1330

- Predicate: We do an analogous procedure.
- Quantifying node: we changed it to the other type.
- Connective node: we replace with a connective node of a different type, selected uniformly at random
- Unitary node: We simply remove this node.

We apply two perturbations to generate each incorrect answer, and repeat the perturbation process to guarantee that all four choices are unique. Finally, we form a question using the following template. Some models (notably o1) have restrictions on the language you can use to prompt the model. In that case "think carefully step-by-step and" was removed from the last sentence.

```
1307
          Logic negation prompt template
1308
1309
          You are a logic student. I will give you a statement in first-order
1310
          logic, and your task is to find it's negation from a list of choices.
1311
          Logical formula: {formula}
1312
1313
          Which of the following is its negation?
1314
1315
          A: {answer 1}
          B: {answer 2}
1316
          C: {answer 3}
1317
          D: {answer 4}
1318
1319
          Please think carefully step-by-step and provide your answer with
1320
          <answer>A, B, C, or D</answer>.
1321
```

G.4 Examples of Failures

Here are some illuminating examples of incorrect responses.

Finding a shortcut but failing at simple logic Here is an example where the prompt is very long but there is an obvious simplification. All models notice this simplification and try to exploit it to find the answer, but all make very elementary mistakes after making the simplification. Given the model, s scores for d=4 trees, we know that they are able to perform the correct logical steps for a simple formula, but fail when the formula is simple but with a large sub-tree in it.

```
1331
1332
         Question
1333
         You are a logic student. I will give you a logical formula, written in
1334
         propositional logic, as well as four options for values of every atomic
1335
         proposition in the formula.
1336
1337
         Logical formula:
1338
         thrilling and ((not thrilling or X....X )) or (dark or thrilling))
1339
         Which of the following choices makes the logical formula evaluate to
1340
         True?
1341
         A) dark: False, dramatic: True, intense: True, thrilling: False,
         suspenseful: False, romantic: True, comedic: True, tragic: False.
         B) dark: False, dramatic: True, intense: True, thrilling: False,
1343
         suspenseful: True, romantic: False, comedic: False, tragic: False.
1344
         C) dark: True, dramatic: False, intense: True, thrilling: False,
1345
         suspenseful: True, romantic: True, comedic: False, tragic: False.
         D) dark: True, dramatic: True, intense: False, thrilling: True,
1347
         suspenseful: False, romantic: False, comedic: False, tragic: True.
1348
         Please provide your answer with <answer>A, B, C, or D</answer>.
1349
```

where the correct answer is D and X...X is a large sub-expression we have hidden.

```
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
```

Step-by-step reasoning:

o1's Response

1) Observe the formula's overall shape. Its outermost structure is of the form (thrilling AND ...) OR (dark OR thrilling).

In more schematic form:

 ${\tt X}$ or (dark or thrilling), where ${\tt X}$ is a very large sub-formula involving all the propositions.

- 2) Recall how OR works in propositional logic:
- \cdot If (dark or thrilling) is true, then the whole formula is true regardless of X.
- If (dark or thrilling) is false, then dark = false and thrilling = false, and in that case we would need (thrilling AND ...) to be true. But if thrilling = false, then (thrilling AND ...) is automatically false, making the entire formula false.
- 3) Check each choice's (dark or thrilling):
 - A) dark = False, thrilling = False => dark or thrilling = False
 Then the left part is "thrilling AND ..." = False AND ... = False.

 Overall result = False or False = False.
 - B) dark = False, thrilling = False => exactly the same reasoning as
 (A), yields False.
 - C) dark = True, thrilling = False => dark or thrilling = True, so the whole formula = X or True = True (no need to check X).
 - D) dark = True, thrilling = True => dark or thrilling = True, so again the whole formula = X or True = True.

Thus both (C) and (D) make (dark or thrilling) = True, and so make the entire formula evaluate to True.

Since the question asks "Which of the following choices makes the logical formula evaluate to True?" the correct answer (in standard multiple-choice fashion) is whichever of (C) or (D) is offered as valid. In fact, both of them do.

Therefore, a valid selection is:

<answer>C</answer>

The error results from eliminating all options except for C and D). This mistake causes the model to focus on the wrong part of the clause.

Gemini makes a similar mistake.

1421

1422

1423

1404 Gemini 2.0 Flash's Response 1405 1406 The given statement is of the form \$P \Leftrightarrow Q\$. The negation of 1407 \$P \Leftrightarrow Q\$ is \$\neg (P \Leftrightarrow Q)\$, which is equivalent to $(P \land Q) \land (\land P \land Q)$. 1408 1409 Let's denote the left side of the main biconditional as \$L\$ and the right 1410 side as \$R\$. So the given statement is \$L \Leftrightarrow R\$. Its negation 1411 1412 \$\neg (L \Leftrightarrow R) \equiv (L \land \neg R) \lor (\neg L \land R)\$. 1413 Looking at the options, we can see that all of them are of the form 1414 $(L \to neg R) \$ (\neg L \to R), which is incorrect. The negation of 1415 \$P \Leftrightarrow Q\$ is not equivalent to 1416 $(P \to p Q) \le (neg P \to R)$ 1417 1418 1419

The reasoning traces of Claude 3.7 often suggest that it is taking shortcuts rather than executing computations, in that explicit computations are not present in the reasoning trace. Claude typically tries to guess the best option based on the formula structure ("This is extremely complex, but I'll focus on the overall structure.")

```
1424
         Claude 3.7 Response
1425
1426
         I need to find the negation of the given logical statement. Let me work
1427
         through this step by step.
1428
         The original statement has the form:
1429
         b or (...complex formula...) -> (...complex formula...) or (h and ...)
1430
1431
         To negate this statement, I need to apply the negation rules:
1432
         - Negation of "p or q" is "not p and not q"
         - Negation of "p -> q" is "p and not q"
1433
         - Negation of "p and q" is "not p or not q"
1434
1435
         So the negation should have the form:
1436
         not b and ((...complex formula...) and not(...complex formula...)) and
1437
          (not h or ...)
1438
         Looking at the options, I notice they all start with "not b and...", which
1439
         is correct for the negation of "b or...".
1440
1441
         Among the choices, I need to check which one correctly negates all parts of
1442
         the original statement. This is extremely complex, but I'll focus on the
         overall structure.
1443
1444
         Since the original statement has the form:
1445
         b or X \rightarrow Y or (h and Z)
1446
1447
         The negation should be:
1448
         not b and X and not Y and (not h or not Z)
1449
         Looking at options B and C, they appear to have the correct structure for
1450
         the negation of the original statement. Option B includes "not b and..."
1451
         followed by the same complex formula structure with appropriate negations
1452
         applied. The final part properly includes "not h or...".
1453
         After careful examination, I believe option B correctly represents the
1454
         negation of the original statement. It properly negates all components
1455
         while maintaining the correct logical structure.
1456
1457
         <answer>B</answer>
```

Nonsensical Logic When there is no simplifying tactic, the models will try to evaluate the expression by breaking it into sub expressions where errors creep in eventually

```
1466
                                  Gemini 2.0 Flash's Response
1467
                                  Let the given statement be $S$. The statement $S$ has the form
1468
                                  (P \land Q) \land (R \land (neg e)) \land T$, where
1469
                                  P = \text{xists } x \in X3 \text{ ((E(x) \ lor (g \ h))}
1470
                                  \left( \left( \left( \left( \left( \left( x \right) \right) \right) \right) \right) \right) 
1471
                                  \inf \left( \int (f \cdot (b)) \right)
1472
                                  Q = \left( \int X \right) 
1473
                                  R = \left( (<table-cell> xists x \in X1 \setminus x1 \setminus s.t. \right) (h \in b) \lor
                                   (\forall x \in X1 E(x))) \lor b)$
1474
                                  T = \int x \in X3 (((((\exists x \in X1)
1475
                                  \text{\text{$\setminus$}} \{ \text{s.t. } C(x) \} \cap g \cap ((C(x) \in f) \cap g) \}
1476
                                  (((b \setminus E(x)) \setminus E(x) \setminus
1477
                                  \to ((c \to b) \iff (C(x) \land C(x)))) \land ((\neg(E(x) \lor
1478
                                  (C(x) \land b \to ((h \to h) \land (\exists x \in X
                                  \text{text} \{ s.t. \} H(x))))) \to g))$
1479
1480
                                  We want to find the negation of $S$, which is
1481
                                  neg S = neg((P \land Q) \land (R \land (neg e)) \land T)$.
1482
                                  Using De Morgan's laws, we have:
1483
                                  \ s = \neg(P \land Q) \lor \neg(R \land (\neg e)) \lor \neg T$
                                  \ \neq S = (\neq P \lor \neq Q) \land (\neq R \lor \neq(\neq e)) \land \neq T
1484
                                  neg S = (neg P \lor neg Q) \land (neg R \lor e) \land neg T$
1485
1486
                                  Let's negate each part:
1487
                                  \ P = \left( x \times x \right) \
                                                   1488
                                  \inf \lceil (f \rceil (b)))...
1489
1490
                                 Without a clear match, and given the high complexity, I cannot confidently
1491
                                  select an option. I suspect an error in the provided question or options.
1492
1493
                                 Final Answer: (No answer can be confidently determined due to
                                  structural mismatch)
1494
                                  <answer>A</answer>"
1495
1496
```

The logic used is hard to follow because it's nonsensical. Gemini starts trying to evaluate the whole phrase, but then essentially gives up and leaves large expressions in the (...).

Misreading the options Another common failure mode is that the model concludes that two options are the same or that no correct answer exists when it erroneously assumes two options are the same.

```
1512
          Gemini 2.0 Flash's Response
1513
1514
           "The given logical statement is:
1515
          not ((((Exists x in X5 s.t. not (k \le M(x))) and ((Exists x in X5 s.t. A(x))
          or b \ll not m) or (not (p and g \rightarrow (e \ll i)))))
1516
          and ((((f \rightarrow 1) \rightarrow j \text{ and } f \leftarrow (Exists x in X1 s.t. N(x) and H(x))) \rightarrow
1517
          ((j or j \rightarrow not m) and (p and g \rightarrow h)))
1518
          <=> ((not j -> (d <=> m -> o)) <=> n and f))) or c)
1519
1520
          We need to find the negation of this statement. Let's represent the statement
          inside the outermost 'not' as P. So the given statement is 'not (P or c)'.
1521
1522
          The negation of 'not (P or c)' is '(P or c)'.
1523
1524
          Now let's substitute P back with its original form:
1525
          (((Exists x in X5 s.t. not (k <=> M(x)))) and ((Exists x in X5 s.t. A(x)) or b
          <=> not m) or (not (p and q \rightarrow (e \le i))))) and (((f <math>\rightarrow l) \rightarrow j and f \le i
1526
          (Exists x in X1 s.t. N(x) and H(x))) -> ((j or j -> not m) and
1527
          (p and g -> h)))
1528
          <=> ((not j -> (d <=> m -> o)) <=> n and f))) or c
1529
1530
          Comparing this result with the given options:
1531
          Option A: (((Exists x in X5 s.t. not (k \le M(x))) and ((Exists x in X5 s.t.
1532
          A(x) or b \ll not m or (not (p and g \rightarrow (e \ll i))))) and ((((f <math>\rightarrow 1) \rightarrow j
1533
          and f \ll (Exists x in X1 s.t. N(x) and H(x)))
1534
          -> ((j or j -> not m) and
1535
          (p \text{ and } g \rightarrow h))) \iff ((not j \rightarrow (d \iff m \rightarrow o)) \iff n \text{ and } f))) \text{ or } c
          This option is exactly the same as the expression we derived for the
1536
          negation.
1537
1538
          Option B: (((Exists x in X5 s.t. not (k \le M(x))) and ((Exists x in X5 s.t.
1539
          A(x) or b <=> not m) or (not (p and <math>g -> (e <=> i))))) and ((((f -> 1) -> j and
1540
          f \ll (Exists x in X1 s.t. N(x) and H(x))) -> ((j or j -> not m) and (p and q)
          -> h))) <=> ((not j -> (d <=> m -> o)) <=> n and f))) or c
1541
          This option is also exactly the same as the expression we derived for the
1542
          negation.
1543
1544
          Option C: (((Exists x in X5 s.t. not (1 \le M(x))) and ((Exists x in X5 s.t.
1545
          A(x) or k \le not n or (not (e and b -> (m <=> f))))) and ((((c -> p) -> o
          and p <=> (Exists x in X1 s.t. N(x) and H(x))) -> ((d or d -> not o) and
1546
          (o and b \rightarrow p))) <=> ((not i \rightarrow (h <=> o \rightarrow b)) <=> e and i))) or d
1547
          This option is different from the derived negation.
1548
1549
          Option D: (((Exists x in X5 s.t. not (k \le M(x))) and ((Exists x in X5 s.t.
1550
          A(x) or b \rightarrow not m) or (not (p and g \rightarrow (e \le i))))) and (((f and <math>l \rightarrow j
          and f <=> (Exists x in X1 s.t. N(x) and H(x))) -> ((j or j -> not m) and
1551
           (p \text{ and } g \rightarrow h))) \iff ((not j \rightarrow (d \iff m \text{ or o})) \iff n \text{ and } f))) \text{ or } c
1552
          This option is also different from the derived negation.
1553
1554
          Since the negation of 'not (P or c)' is '(P or c)', and both options A and
1555
          B are identical to P or c, both are correct negations. However, usually
          there should be only one correct answer in a multiple-choice question.
1556
          Let's re-examine the question and the options to ensure there isn't a subtle
1557
          difference I missed.
1558
1559
          There is no subtle difference. Options A and B are identical and represent
1560
          the negation of the given statement. Assuming there was a typo in the
          question and only one option should be correct, and given that they are
1561
          identical, we can choose either one.
1562
1563
          <answer>A</answer>"
1564
```

Missing Simple Rules Perhaps surprisingly, sometimes the models fail because they cannot apply simple logical rules (De Morgan's, in this case:)

```
"First, note that the given statement is of the form
not ( Exists x in X_2 : ...)

in other words, it begins with "not (Exists x in X_2 s.t. ...)."

The standard logical negation of
not (Exists x : P(x))
is
Exists x : not P(x).
...
```

We truncated the output because the error is already present: the model negated the first term incorrectly, clearly violating the rule it just wrote out.

H DETAILS FOR PROOF TREE

This section describes the Proof Tree construction, along with the inference rules we derived.

H.1 LOGICAL FORMS

1566

1567

1568

1569 1570

1571

1572 1573

1574 1575

1576

1579 1580

1581

1585

1587 1588

1590

1591

1592

1593

1596

1598

1604

1609

1610

1611

1612 1613

1614 1615

1616

1617

1618

1619

These forms were included in the original paper:

- container (cont): "A has 5 apples."
- comparison (comp): "A has 3 more apples than B."
- transfer: "A gives B 3 apples."
- comp-eq: "The number of apples that C has more than D is equal to the difference between the number of apples that A and B have."
- partwhole: "A and B combine the fruits that they have."

Our diverse rules task added the following additional rules.

- consume (cons): "A eats 5 apples."
- increase: "The number of apples that A has increases by 2 times."
- switch: "A and B switch the apples they have."
- redistribute: "A and B redistribute their apples to ensure each has an equal amount."
- split: "A splits all the apples she owns equally between B and C."
- conditional transfer (cond-transfer): "If B has more than 2 apples, B will transfer all their apples to A."
- cumulative (cum): "The combined quantity of apples that A, B, and C have is 20."
- multi-agent comparison (multi-comp): "A has 10 more apples than B and C combined."
- sequential comparison (seq-comp): "A has 3 more apples than B and 5 less apples than C."

H.2 INFERENCE RULES

Each logical form requires inference rules that describe its implications on our knowledge of the number of apples everyone has. The inference rules from the original paper include:

```
• ContCompInference: \frac{\text{cont}(a,q_1,e) - \text{comp}(b,a,q_2,e)}{\text{cont}(b,q_1+q_2,e)}
```

Example: "Alice has 3 apples. Bob has 2 more apples than Alice. ⊢ Bob has 5 apples."

```
• ContTransferInference: \frac{\operatorname{cont}(a,q_1,e)}{a} \frac{\operatorname{transfer}(a,b,q_2,e)}{a}
1620
                                                                  cont(a,q_1+q_2,e)
1621
1622

    Example: "Alice has 3 apples. Bob gave 2 apples to Alice. ⊢ Alice has 5 apples."

                     • ContContInference: \frac{\text{cont}(a,q_1,e)}{\text{cont}(b,q_2,e)}
1623
                                                         comp(b,a,q_2-q_1,e)
1624
1625

    Example: "Alice has 3 apples. Bob has 5 apples. ⊢ Bob has 2 more apples than Alice."

                     • CompEqInference: \frac{\operatorname{cont}(a,q_1,e) - \operatorname{comp}(d,c,q_2,e)}{\operatorname{comp-eq}(b,a,d,c,e)}
1626
1627
                                                                       cont(b,q_1+q_2,e)
1628
                           - Example: "Alice has 7 apples. David has 2 more apples than Charlie. The number of
1629
                               apples that Bob has more than Alice is the same as the difference between the number
                               of apples that David and Charlie have. ⊢ Bob has 9 apples."
           To be able to make correct inferences over our new rules, we also derived the following inference
           rules.
1633
                     • ContConsInference: \frac{\text{cont}(a,q_1,e) - \text{cons}(a,q_2,e)}{\text{cont}(a,q_1,e)}
                                                           \operatorname{cont}(a,q_1-q_2,e)
1635

    Example: "A has 10 apples. A eats 3 apples. ⊢ A has 7 apples."

1637
                     • ContIncreaseInference: \frac{\text{cont}(a,q_1,e)}{a_1,a_2,\dots,a_{r-1}}
                                                                         increase(a,q_2,e)
                                                                 cont(a,q_1 \times q_2,e)
1639
                            - Example: "A has 4 apples. The number of apples that A has increases by 3 times. \vdash A
1640
                               has 12 apples."
1641
                     • ContSwitchInference: \frac{\operatorname{cont}(a,q_1,e)}{\operatorname{cont}(b,q_2,e)}
                                                                                        switch(a,b,e)
1642
                                                                cont(a,q_2,e) \quad cont(b,q_1,e)
1643
                           - Example: "A has 5 apples. B has 8 apples. A and B switch the apples they have. ⊢ A
1644
                               has 8 apples. B has 5 apples."
1645
                     • ContRedistributeInference: \frac{\cot(a,q_1,e)}{a+a_0} \frac{\cot(b,q_2,e)}{a+a_0} redistribute(a,b,e)
1646
                                                                     cont(a, \frac{q_1+q_2}{2}, e) \quad cont(b, \frac{q_1+q_2}{2}, e)
1647
1648
                           - Example: "A has 6 apples. B has 10 apples. A and B redistribute their apples to ensure
                               each has an equal amount. ⊢ A has 8 apples, and B has 8 apples."
1649
1650
                     • SplitInference: \frac{\cot(a,q_1,e)}{\cot(a,q_1,e)} \frac{\cot(b,q_2,e)}{\cot(a,q_1,e)} \frac{\operatorname{split}(a,q_4,\{b,c\},e)}{\cot(a,q_1,e)}
                                                    cont(a, q_1 - q_4, e) \quad cont(b, q_2 + \frac{q_4}{2}, e)
1652
                           - Example: "A has 12 apples. B has 4 apples. A splits all the apples she owns equally
                               between B and C. ⊢ A has 0 apples. B has 10 apples."
                     • CondTransferInference: \frac{\cot(a,q_1,e)}{\cot(a,q_1,e)} \frac{\cot(b,q_2,e)}{\cot(b,q_2,e)} \frac{\cot(b,q_2,e)}{\cot(b,q_2,e)} \frac{\cot(b,q_2,e,q_2>q_3)}{\cot(b,q_2,e)}
                                                                cont(a,q_1+q_2,e) if q_2>q_3; cont(a,q_1,e) otherwise
1656

    Example: "A has 5 apples. B has 7 apples. If B has more than 6 apples, B will transfer

1657
                               all their apples to A. \vdash A has 12 apples."
1658
                     • CumulativeToContInference: \frac{\text{cont}(a_1,q_1,e)}{\dots}
                                                                                        cont(a_{n-1},q_{n-1},e)
                                                                                                                 \operatorname{cum}(a_1,...,a_n,q,e)
1659
                                                                                       cont(a_n, q - \sum_{i=1}^{n-1} q_i, e)
1660
                           - Example: "A has 5 apples. B has 3 apples. The combined quantity of apples that A, B,
                               and C have is 15. \vdash C has 7 apples."
                     • MultiCompInference: \frac{\operatorname{cont}(a,q_1,e) - \operatorname{cont}(b,q_2,e) - \operatorname{multi-comp}(a,b,c,q_3,e)}{\operatorname{cont}(a,q_1,e) - \operatorname{cont}(a,q_1,e)}
1663
                                                                         \operatorname{cont}(c,q_1-q_2-q_3,e)
                           - Example: "A has 12 apples. B has 2 apples. A has 10 more apples than B and C
                               combined. ⊢ C has 0 apples."
                     \bullet \ \ \mathbf{SeqCompInference:} \ \ \frac{\mathsf{seq\text{-}comp}(a,b,c,q_1,q_2,e)}{\mathsf{cont}(a,q_3+q_1,e)} \ \ \frac{\mathsf{cont}(b,q_3,e)}{\mathsf{cont}(c,q_3+q_1+q_2,e)} \\
                                                                          \operatorname{cont}(c, q_3 + q_1 + q_2, e)

    Example: "A has 3 more apples than B and 5 fewer apples than C. B has 7 apples. ⊢ A

1669
                               has 10 apples. C has 15 apples."
```

H.3 DETAILS FOR PROOF TREE IRRELEVANT

1671

16721673

The irrelevant sentences are samples from the following list:

```
1674
         Irrelevant sentences template
1675
1676
         "{} is very generous and enjoys
1677
         sharing food with others.",
1678
         "{} tends to be laid-back and prefers
         staying in rather than going out.",
1679
         "{} is highly introverted and prefers
         minimal communication with others.",
1681
         "{} is very outgoing and frequently
1682
         hosts parties at home.",
1683
         "{} and {} are good friends who often
1684
         go fruit or vegetable picking together
1685
         on weekends.",
         "{} and {} have been married for {}
1687
         years.",
1688
         # Random years will be added
1689
         "{} is {} years old."
         # Random age will be added
```

H.4 CONSTRUCTING BASIC PROOF TREE AND PROMPTS

A proof tree is generated by first picking a target conclusion predicate—a "cont" (container) that states how many items a single agent possesses. Given this target, the system identifies all inference rule classes that can yield such a conclusion. Each rule class is assigned a weight, determining its likelihood of selection; higher weights correspond to a greater chance of being chosen. Specifically, "ContCompInference" is weighted at 1, "ContTransferInference" at 5, "ContContInference" at 1, and "CompEqInference" at 10. The system then randomly selects one inference rule among those whose premises can produce the target conclusion, with the probability of each rule proportional to its weight. The chosen rule provides the premises (new conclusion targets) required to derive the original predicate. Each of these premises is then handled the same way: we attempt to produce them (recursively) via suitable rules, or it marks them as leaves (facts) if no rules fit or the tree has reached its maximum size constraints. This procedure yields a proof tree where each internal node applies a randomly selected (but weighted) inference rule to derive the node's conclusion from its premises, while the leaves represent axiomatic statements used in the proof. See Algorithm 4 and Algorithm 5 for the pseudocode.

Algorithm 4: Pseudocode for Generating a Proof Tree

```
Function GenerateProofTree (max\_depth, max\_leaves, available\_agents); selected\_agent \leftarrow randomly pick 1 from <math>available\_agents; remove selected\_agent from available\_agents; quantity \leftarrow random integer in [10, \dots, 30]; entity \leftarrow random pick an entity; root\_predicate \leftarrow Cont(selected\_agent, quantity, entity); return GenerateSubtree(root\_predicate, max\_depth, max\_leaves, 0, 1, available\_agents)
```

Once the proof tree is constructed, its leaves are traversed in order and converted into sentences using natural language templates, forming the textual body of the problem. The question of the problem is derived from the logical form at the root of the proof tree.

```
1728
        Algorithm 5: Pseudocode for Generating a Subtree Tree
1729
        Data: node, max_depth, max_leaves, current_depth, current_leaves, available_agents
1730
        if current\_depth \ge max\_depth then
1731
            return node / * do not expand further at max depth
1732
        candidate\_rules \leftarrow \varnothing;
1733
        for rule class \in
1734
          {ContCompInference, ContTransferInference, CompEqInference, ContContInference} do
1735
            if rule_class.can_yield(node.conclusion, available_agents) and
1736
             (rule\_class.num\_premises + current\_leaves) \le max\_leaves then
1737
                candidate_rules.add(rule_class);
1738
1739
        if candidate_rules is empty then
1740
            return node /* no valid rules; node is leaf
                                                                                                      */
1741
        weights \leftarrow \text{map each rule class in } candidate\_rules \text{ to its weight;}
1742
        chosen\_rule\_class \leftarrow randomly select from <math>candidate\_rules using weights;
        instantiated\_rule \leftarrow
1743
         chosen_rule_class.make_rnd_instance(node.conclusion, available_agents);
1744
        node.rule \leftarrow instantiated\_rule;
1745
        current\_leaves \leftarrow current\_leaves + instantiated\_rule.num\_premises - 1;
1746
        if current\_depth < max\_depth - 1 then
1747
            for premise \in instantiated\_rule.premises do
1748
                child\_node \leftarrow Generate Subtree with data premise, max\_depth, max\_leaves,
1749
                 current\_depth + 1, current\_leaves, available\_agents, and
1750
                 node.children.add(child_node;
1751
                current\_leaves \leftarrow current\_leaves + (child\_node.num\_leaves() - 1);
1752
                for agent \in premise.agents() do
1753
                    if agent \in available agents then
                     remove agent from available_agents
1754
1755
        return node
1756
```

cturn noac

1757 1758

1759 1760

1761

1762

1763

1764

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

1777

1778

1779

1780

1781

```
Proof Tree Example with max depth 5 and max leaves 20
Lindsay has 13 apples.
    Arleth has 4 more apples than Mathew.
    Nellie has 17 apples.
        Dian has 3 more apples than Amy.
            Amy has 17 apples.
                Courtney has 14 more apples than Peggie.
                Ida has 31 apples.
                The number of apples that Peggie has more than Courtney
                is equal to the difference between the number of apples
                that Amy and Ida have.
            Dian has 20 apples.
                Dian has 13 apples.
                Prudence gives 7 apples to Dian.
        Annabelle has 14 apples.
            Lacie has 13 more apples than Federico.
            Georgia has 27 apples.
                Jose has 13 more apples than Agatha.
                Wilson has 40 apples.
                The number of apples that Agatha has more than Jose is
                equal to the difference between the number of apples
                that Georgia and Wilson have.
            The number of apples that Federico has more than Lacie is
            equal to the difference between the number of apples that
            Annabelle and Georgia have.
        The number of apples that Dian has more than Amy is equal to the
        difference between the number of apples that Nellie and
        Annabelle have.
    The number of apples that Mathew has more than Arleth is equal to the
    difference between the number 39 apples that Lindsay and Nellie have.
```

Prompt Example

Courtney has 14 more apples than Peggie. Ida has 31 apples. The number of apples that Peggie has more than Courtney is equal to the difference between the number of apples that Amy and Ida have. Dian has 13 apples. Prudence gives 7 apples to Dian. Jose has 13 more apples than Agatha. Wilson has 40 apples. The number of apples that Agatha has more than Jose is equal to the difference between the number of apples that Georgia and Wilson have. Lacie has 13 more apples than Federico. The number of apples that Federico has more than Lacie is equal to the difference between the number of apples that Annabelle and Georgia have. The number of apples that Dian has more than Amy is equal to the difference between the number of apples that Nellie and Annabelle have. Arleth has 4 more apples than Mathew. The number of apples that Mathew has more than Arleth is equal to the difference between the number of apples that Lindsay and Nellie have. How many apples does Lindsay have? Give your answer using the format: "The final answer is \$\boxed{#your answer}\$."

H.5 CONSTRUCTING PROOF TREES WITH DIVERSE STATEMENTS

In this task, given a diverse set of logical statements, the model must answer word-based questions that require deduction, sampled from a tree with a bounded depth and number of leaves. The parameters are the maximum tree depth d, and whether to include the additional logical forms.

Proof Tree with Diverse Statements example

Briana has 2 bananas. Tom has 0 bananas. If Tom has more than 1 bananas, Tom will transfer all their bananas to Briana.... Whitney and Freida redistribute their bananas to ensure each has an equal amount. Eula has 6 more bananas than Dexter and 11 fewer bananas than Bernardo.... How many bananas does Amelia have?

Give your answer using the format: The final answer is \$\boxed{#your answer}\$.

The process of constructing a proof tree with diverse statements is similar to the basic proof tree construction, with the key difference being the set of inference rules used and their assigned weights. Specifically, the weights for the inference rules are as follows: "ContCompInference" is weighted at 1, "ContTransferInference" at 1, "ContContInference" at 1, "ContEntInference" at 10, "ContSwitchInference" at 10, "ContRedistributeInference" at 10, "SplitInference" at 10, "CondTransferInference" at 10, "CumulativeToContInference" at 1, "MultiCompInference" at 10, "SeqCompInference" at 10.

In our experiments, we set the maximum number of leaves to 20. We then vary the maximum depth and the inclusion of diverse statements to evaluate the model's performance.

1837 1838

1839 1840

1841

1843

1844

1845

1846

1847

1849

1850

1851

1857

H.6 CONSTRUCTING PROMPTS WITH IRRELEVANT INFORMATION

```
Proof Tree with Irrelevant information example

Veda is very generous and enjoys sharing food with others.

Sibyl has 14 more apples than Ashley. ...

The number of apples that Ali has more than Howell is equal to the difference between the number of apples that Jacqueline and Vollie have....

Carlo tends to be laid-back and prefers staying in rather than going out....

How many apples does Destiny have?

Give your answer using the format:

The final answer is \$\textbackslash boxed\{\#\textbackslash text\{your answer\}\}\$.
```

In problems involving proof trees with irrelevant information, the problem parameters are the maximum tree depth d, the number of irrelevant people P, and the number of irrelevant sentences S. To construct prompts containing irrelevant information, we first generate the baseline proof tree with a maximum depth of 5 and a maximum of 20 leaves. Irrelevant information is then introduced through two main components: irrelevant agents and irrelevant sentences:

- *Irrelevant agents*: Irrelevant agents are created by dividing the pool of agent names into subsets that are distinct from the key agents, ensuring no overlap. These subsets are then used to generate irrelevant proof trees, employing a consistent randomization process (i.e., all the irrelevant proof trees are identical to the key proof trees, differing only in the names of the agents involved). Each irrelevant proof tree is converted into axioms and shuffled alongside the key axioms.

```
1859
1860
         Irrelevant Proof Tree Example with max depth 5 and max leaves 20
1861
         Nora has 13 apples.
1862
             Hal has 4 more apples than Jean.
1863
             Aggie has 17 apples.
                  Theron has 3 more apples than Marjorie.
                      Marjorie has 17 apples.
1866
                          Caryl has 14 more apples than Robert.
                          Philomena has 31 apples.
                          The number of apples that Robert has more than Caryl is
1868
                          equal to the difference between the number of apples that
                          Marjorie and Philomena have.
1870
                      Theron has 20 apples.
                          Theron has 13 apples.
                          Stefani gives 7 apples to Theron.
1872
                  Genevieve has 14 apples.
1873
                      Ida has 13 more apples than Angelique.
1874
                      Doris has 27 apples.
1875
                          Lorenzo has 13 more apples than Gussie.
                          Adrian has 40 apples.
1876
                          The number of apples that Gussie has more than Lorenzo
1877
                          is equal to the difference between the number of apples
1878
                          that Doris and Adrian have.
1879
                      The number of apples that Angelique has more than Ida is
1880
                      equal to the difference between the number of apples that
                      Genevieve and Doris have.
                  The number of apples that Theron has more than Marjorie is equal
                  to the difference between the number of apples that Aggie and
                  Genevieve have.
             The number of apples that Jean has more than Hal is equal to the
1885
              difference between the number of apples that Nora and Aggie have.
```

- *Irrelevant sentences*: Irrelevant sentences are generated using predefined templates (see Section H.3). To integrate the irrelevant information with the context, these sentences are randomly inserted into the shuffled list of axioms at arbitrary positions.

Irrelevant Sentence Examples Caryl is highly introverted and prefers minimal communication with others. Courtney is very generous and enjoys sharing food with others. Adella is 46 years old. Newton tends to be laid-back and prefers staying in rather than going out. Arleth is very generous and enjoys sharing food with others. Nico is very outgoing and frequently hosts parties at home. Dennis is very generous and enjoys sharing food with others. Moe and Agatha have been married for 17 years. Rubie and Angelique have been married for 16 years. Jean tends to be laid-back and prefers staying in rather than going out. Mathew is highly introverted and prefers minimal communication with others. Dalton tends to be laid-back and prefers staying in rather than going out. Joel is very generous and enjoys sharing food with others. Adrian and Perla are good friends who often go fruit or vegetable picking together on weekends. Rosina is very generous and enjoys sharing food with others. Mickie is very outgoing and frequently hosts parties at home. Elijah is very generous and enjoys sharing food with others. Bert is very generous and enjoys sharing food with others. Robert is 32 years old. Delma is highly introverted and prefers minimal communication with others. Vallie and Miriam are good friends who often go fruit or vegetable picking together on weekends. Orma is highly introverted and prefers minimal communication with others. Cornelius is very generous and enjoys sharing food with others. Marylee is very outgoing and frequently hosts parties at home. Mitchell and Doris have been married for 3 years.

In our experiments, we set the maximum depth to 5 and the maximum number of leaves to 20. We then vary the number of irrelevant agents A and the number of irrelevant sentences S to evaluate the model's performance.

H.7 Example of Failure Due to Diverse Statements

Question

Milo has 24 blueberries. The number of blueberries that Milo has increases by 19/6 times. Aubrey has 1 more blueberry than Larue. Gregory has 3 blueberries. The number of blueberries that Larue has more than Aubrey is equal to the difference between the number of blueberries that Thaddeus and Gregory have. Zane has 7 blueberries. The number of blueberries that Zane has increases by 9/7 times. Odell has 21 blueberries. Christian has 3 blueberries. If Christian has more than 5 blueberries, Christian will transfer their blueberries to Odell. Milo splits all the blueberries she owns equally among Thaddeus and Della. Zane and Odell redistribute their blueberries to ensure each has an equal amount. Thaddeus has 17 more blueberries than Zane and Artie combined. Keegan has 7 more blueberries than Aliya and 7 fewer blueberries than Everett. Aliya has 0 blueberries. Margarita has 1 blueberry. If Margarita has more than 4 blueberries, Margarita will transfer their blueberries to Artie. The number of blueberries that Everett has increases by 17/14 times. If Everett has more than 2 blueberries, Everett will transfer their blueberries to Artie. How many blueberries does Artie have? Give your answer using the format: "The final answer is \$\boxed{#your answer}\$."

```
1998
          o1's Response
1999
2000
          First, let us track each person's blueberry count step by step
2001
          (in the order given), focusing on changes that actually affect Artie:
          1) Milo starts with 24 blueberries.
2002
             The puzzle says Milo's count "increases by 19/6 times," so Milo now has
2003
             24 \times (19/6) = 24 \times 19 \div 6 = 76.
2004
          2) Aubrey has 1 more blueberry than Larue (Aubrey = Larue + 1), but no specific
          numbers result yet.
2006
          3) Gregory has 3 blueberries.
          4) Zane starts with 7 blueberries.
2007
             The puzzle says Zane's count "increases by 9/7 times," so Zane now has
2008
             7 \times (9/7) = 9.
2009
          5) Odell has 21 blueberries.
2010
          6) Christian has 3 blueberries.
2011
             The puzzle says that if Christian has more than 5, Christian would transfer
             them to Odell; since Christian has only 3, no transfer happens.
2012
         7) Milo (who has 76) "splits all the blueberries she owns equally among Thaddeus
2013
          and Della."
2014
             Each receives 76 \div 2 = 38.
2015
             Thus, Thaddeus = 38, Della = 38, and Milo goes to 0.
          8) Zane (9) and Odell (21) "redistribute their blueberries to ensure each has
2016
          an equal amount."
2017
             Together they have 9 + 21 = 30, so each ends up with 15.
2018
             Hence Zane = 15, Odell = 15.
2019
          9) "Thaddeus has 17 more blueberries than Zane and Artie combined."
2020
             Thaddeus = 38 from above. Let Artie = x. Then
2021
             38 = (Zane's 15) + x + 17
             38 = 15 + x + 17 = 32 + x
2022
             x = 6.
2023
             So Artie = 6 at this point.
2024
          10) Keegan has 7 more blueberries than Aliya and 7 fewer than Everett.
2025
              Aliya = 0, so Keegan = 7 and Everett = 14.
2026
          11) Margarita has 1 blueberry. Because that is not "more than 4," there is
          no transfer to Artie.
2027
         12) Everett's 14 blueberries "increase by 17/14 times," so Everett becomes
2028
              14 \times (17/14) = 17.
2029
          13) "If Everett has more than 2 blueberries, Everett will transfer their
2030
          blueberries to Artie.'
              Since Everett has 17 (> 2), Everett transfers all 17 to Artie, leaving
2031
              Everett at 0 and raising Artie's total to
2032
              Artie = 6 + 17 = 23.
2033
         No further steps change Artie's count. Therefore, Artie finishes with 23
2034
         blueberries.
2035
          The final answer is \frac{23}{5}.
2036
2037
```

In 7), o1 claims that Thaddeus has 38 after the splitting, ignoring the fact that Thaddeus initially had 2 (which can be inferred from previous information). This lack of understanding of the splitting process leads to an incorrect solution by o1.

H.8 FAILURE PATTERNS WITH THE INCLUSION OF IRRELEVANT INFORMATION

20382039

2040

2041

204220432044

2045

2049 2050

- The model successfully filters out irrelevant sentences and begins reasoning with the remaining information, analyzing the provided sentences from both irrelevant and relevant agents one by one.
 - (a) The model overlooks important relevant information, resulting in an incorrect solution.

2052 Question 1 2053 Franco has 5 more oranges than Sergio. Alfred has 5 more oranges 2055 than Parker. Angus and Danny have been married for 11 years. Penelope has 15 oranges. The number of oranges that Sergio has more 2056 than Franco is equal to the difference between the number of oranges 2057 that Susan and Penelope have. Oran tends to be laid-back and prefers 2058 staying in rather than going out. 2059 2060 Novella has 8 oranges. Susan is highly introverted and prefers minimal communication with others. Leann is highly introverted and 2061 prefers minimal communication with others. Hosea is very outgoing 2062 and frequently hosts parties at home. Meredith is highly introverted 2063 and prefers minimal communication with others. Novella gives 3 2064 oranges to Angus. Collie and Gust have been married for 20 years. 2065 Maxine has 2 more oranges than Celeste. 2066 How many oranges does Cornelia have? 2067 Give your answer using the format: "The final answer is \$\boxed{#your answer}\$." 2070

```
o1's Response for Question 1
```

2072 2073

2074

2075

2076

2077

2078

20792080

2081

2082

2083

2085

2086

2090

2091

2092

2093

2094

2095 2096

2097

2098

2099

2100

2102 2103

2104

2105

```
First, note that several statements in the puzzle concern
personalities or marriages but do not affect the orange counts; we
will ignore those and track only statements that specify
relationships or transfers of oranges.
1) "Franco has 5 more oranges than Sergio."
   Let S = Sergio's oranges. Then Franco's oranges F = S + 5.
2) "Penelope has 15 oranges."
   Pn = 15.
10) "Novella has 8 oranges."
   N1 = 8.
11) "Maxine has 2 more oranges than Celeste."
    Let Ce = Celeste's and Mx = Maxine's -> Mx = Ce + 2.
12) "The number of oranges that Celeste has more than Maxine is equal
to the difference between the number of oranges that Johathan and
Novella have."
    Ce - Mx = Jo - Nl, where Jo = Johathan's oranges.
    But Mx = Ce + 2 -> Ce - (Ce + 2) = -2.
    So -2 = Jo - 8 \rightarrow Jo = 6 (Johathan has 6).
18) "The number of oranges that Florida has more than Aaron is equal
to the difference between the number of oranges that Cornelia and
Johathan have."
    Florida has 13, Aaron has 10 \rightarrow 13 - 10 = 3.
    Let Co' = Cornelia's oranges;
    Johathan has 6 -> 3 = Co' - 6 -> Co' = 9.
    So before any transfers to Cornelia, she has 9.
20) "Beaulah gives 17 oranges to Cornelia."
    Cornelia had 9, then receives +17 -> Cornelia now has 26.
21) "Janet gives 17 oranges to Kattie" (irrelevant to Cornelia's
total). No further changes to Cornelia's total are stated.
Therefore, Cornelia ends up with 26 oranges.
The final answer is $\boxed{26}$.
```

o1 overlooks the critical information that "Novella gives 3 oranges to Angus," resulting in an incorrect calculation of Novella's oranges. The correct number should be 5 instead of 8, ultimately leading to an incorrect solution.

(b) Redundant information confuses the model, resulting in either an incorrect solution, an incomplete proof, or the model claiming that a proof is impossible.

Gemini 2.0 Flash's Response Let \$S\$ be the number of blueberries Sherman has, \$T\$ be the number of blueberries Troy has. \$S = T + 7\$ (1) Let \$N\$ be the number of blueberries Ned has, \$G\$ be the number of blueberries Garth has. \$N = G - 7\$ or \$G = N + 7\$ (2) Jarvis has \$J = 20\$ blueberries. Hattie has \$H = 20\$ blueberries. ... We are given insufficient information to uniquely determine the value of Yesenia. There might be an error in the problem statement or missing information. ...

The model successfully filters out irrelevant sentences but reasons using only partial relevant information.

```
Question 2
Jarred is very outgoing and frequently hosts parties at home. Annie has 2
tomatoes. William has 5 tomatoes. Lavinia is very generous and enjoys
sharing food with others. Alyce has 19 tomatoes.
Alyssa has 7 more tomatoes than Jones.
Roslyn has 19 tomatoes.
The number of tomatoes that Jones has more than Alyssa is equal to the
difference between the number of tomatoes that Bonnie and Roslyn have.
Bonnie gives 3 tomatoes to Lavinia. Faith is 19 years old.
Alyce and Myles have been married for 17 years. Charlene gives 4 tomatoes
The number of tomatoes that Clara has more than Gracie is equal to the
difference between the number of tomatoes that Ryann and Bonnie have.
Fernando and Alton are good friends who often go fruit or vegetable
picking together on weekends. How many tomatoes does Ryann have?
Give your answer using the format:
"The final answer is $\boxed{#your answer}$."
```

2188

2189

21902191219221932194

21952196219721982199

```
2160
                 o1's Response for Question 2
2161
2162
                 First, note that several "difference" statements in the puzzle equate
2163
                 one pair's difference of tomato counts to another pair's difference.
                 In particular, two key statements pin down Bonnie's and then Ryann's
2164
                 counts.
2165
                 1) From "Alyssa has 7 more tomatoes than Jones" and "The number of
2166
                 tomatoes that Jones has more than Alyssa is equal to the difference
2167
                 between the number of tomatoes that Bonnie and Roslyn have," we get:
2168
                    • Let J = Jones's tomatoes, A = Alyssa's, B = Bonnie's,
                    and R = Roslyn's (19).
2169
                    • A = J + 7.
2170
                    • J - A = B - 19.
2171
                    Substituting A = J + 7 gives J - (J + 7) = B - 19 -> -7 = B - 19
2172
                    -> B = 12
2173
                    So Bonnie ends up with 12.
                 2) From "The number of tomatoes that Clara has more than Gracie is equal
2174
                 to the difference between the number of tomatoes that Ryann and Bonnie
2175
                 have," we have:
2176
                    • Clara = 20 and Gracie = 22 (found elsewhere in the puzzle),
2177
                    • so (Clara - Gracie) = 20 - 22 = -2,
2178
                    · and that equals (Ryann - Bonnie).
                    Hence -2 = Ryann - 12 \rightarrow Ryann = 10.
2179
                 Accordingly, Ryann has 10 tomatoes.
2180
                 The final answer is \frac{10}{\$}.
2181
2182
```

o1 reasons based only on partial relevant information instead of tracking the dynamic changes in Bonnie's tomato count (e.g., "Bonnie gives 3 tomatoes to Lavinia"). This oversight leads to an incorrect solution.

3. The model bases its reasoning on irrelevant sentences.

```
2200
2201
                Ouestion 3
2202
                Ollie has 5 more onions than Antonia. Wilford has 5 more onions than Ida.
2203
2204
                Angel has 11 onions. Gregorio is 35 years old. Janelle has 9 onions.
2205
                The number of onions that Bailey has more than Stacey is equal to the
                difference between the number of onions that Gregorio and Lorraine have.
                The number of onions that Enrique has more than Reynold is equal to the
2208
                difference between the number of onions that Abner and Ivory have.
2209
                Alonzo gives 20 onions to Gregorio. Nora gives 20 onions to Abner.
2210
                Enrique tends to be laid-back and prefers staying in rather than going
2211
                out. How many onions does Gregorio have?
2212
                Give your answer using the format:
                "The final answer is $\boxed{#your answer}$."
```

```
2214
                 o1's Response for Question 3
2215
2216
                Here is the streamlined logical path:
2217
                1) Ollie has 5 more onions than Antonia.
                    - Denote Antonia's initial onions by A, so Ollie = A + 5.
2218
2219
                 9) We read "The number of onions Bailey has more than Stacey is equal to
2220
                 the difference between Gregorio and Lorraine."
2221
                    - Bailey - Stacey = 22 - 15 = 7.
2222
                    - The difference between Gregorio (call it G for short) and Lorraine (8)
                   must also be 7.
2223
                    - So either G - 8 = 7 (giving G = 15) or 8 - G = 7 (giving G = 1).
2224
                10) Finally, Alonzo gives Gregorio 20 onions. Thus Gregorio's final
2225
                 total = (his initial) + 20.
2226
                    - If Gregorio's initial were 1, then final would be 21.
2227
                    - If Gregorio's initial were 15, then final would be 35.
2228
                Because the puzzle is set up so that Gregorio's age is given as 35 (very
2229
                likely a hint that he ends up with 35 onions!), the consistent (and
2230
                usual) solution is that he started with 15 and, after receiving 20 onions,
2231
                 finishes with 35.
2232
                 The final answer is \frac{35}{5}.
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
```

I DETAILS FOR TRAVEL PLANNING

This section outlines the methodology for generating travel planning problems, including the construction of a city connection graph, the generation of feasible travel paths, and the formulation of problem descriptions. The model is asked to create a travel itinerary using a city connection graph that adheres to a list of constraints. The parameters include the number of cities in the graph S, the number of transportation modes A, and the number of unique cities N.

Travel Planning example

You are a proficient planner. Based on the provided information and query, please give me your plan as a sequence of trips in the format: [(city1, city2, transportation_method), ...]

You are planning a trip across 10 cities with up to 2 transportation methods. The cities are: ['Arlington',...]

The available transportation methods are: ['tram', 'car']

Here are the travel connections:

- From Arlington to New Orleans: car (cost: car=\$53)
- ..
- From Arlington to Fresno: tram, car (cost: tram=\$54, car=\$14)

Constraints:

2268

2270

2271

2272

2273

2274

2275

22762277

2278

2279

2281

2282

2283

2285

2287

2289

2290

2291

2293 2294

2295

2296 2297

2298

2299

2300

2301

2302

2305 2306

2307

2308

2309 2310

2311

2312

2313

2314

2315

2316

2317

2318 2319

2320

2321

- 1. Start your trip at 'Philadelphia' and end at 'Irvine'.
- 2. You cannot exceed a budget of \$163.
- 3. Visit at least 5 unique cities, including the start and end cities.

I.1 CONSTRUCTING THE CITY CONNECTION GRAPH

The travel planning process begins with the creation of a graph representing city connections. The steps are as follows:

1. Selection of Cities and Transportation Methods:

- Choose the 100 largest U.S. cities by population.
- Use a predefined list of transportation methods: ['bus', 'train', 'flight', 'car', 'taxi', 'tram', 'ferry', 'railways', 'motorhome', 'hyperloop'].
- Randomly select a subset of S cities and A transportation methods for the problem.

2. Graph Construction:

- Create a directed graph where cities are nodes, and transportation connections are edges.
- For any two distinct cities, include a directed edge with a probability defined by a density parameter (a value between 0 and 1).

3. Edge Weights and Costs:

- For each established edge, select a random number of transportation modes (uniformly between 1 and A) from the available list.
- For each mode:
 - Generate a cost range by randomly selecting:
 - * lowest_cost from [10, 50].
 - * highest_cost from [60, 100].
 - Assign an actual cost for traveling via the mode as a random integer between lowest_cost and highest_cost.

I.2 BUILDING A FEASIBLE PATH AND COMPUTING THE BUDGET

The objective is to create a travel plan that visits at least a predefined number of distinct cities and to calculate the required budget.

1. Path Construction:

2322

2323

2375

2324	 Select a random starting city. 	
2325	 Maintain a set of visited cities (visited_cities) to avoid revisiting. 	
2326	• At each step:	
2327	 Identify adjacent cities from the current city, excluding already visited ones. 	
2328	 If a "dead end" is reached (no unvisited adjacent cities), abandon the attempt an 	nd
2329	retry with a different starting city.	
2330	- Among feasible next-city options, choose the city and transportation mode with the	he
2331	lowest cost.	
2332	• Continue until the desired number of unique cities is visited, marking the final city	as
2333	the ending point.	
2334	2. Budget Calculation:	
2335	• Sum the costs of all steps in the feasible path to determine the total travel budget.	
2336	• This budget acts as a constraint for the solution.	
2337	This cauget at a to instant for the section.	
2338	I.3 GENERATING THE PROBLEM DESCRIPTION	
2340 2341	The problem description includes all details required for solving the travel planning task:	
2342 2343	 The number of selected cities and transportation modes. 	
2344	 Lists of the specific cities and transportation methods. 	
2345	• For each directed edge in the graph:	
2346	 Possible transportation modes and their associated costs. 	
2347	Problem constraints:	
2348 2349	 Designated starting and ending cities. 	
2350	 Computed budget. 	
2351	 Minimum number of distinct cities to visit. 	
2352	- Minimum number of distinct cities to visit.	
2353 2354	An example query and travel plan are provided to guide response formatting and ensure accura parsing of solutions.	ite
2355 2356	I.4 RETURNING AND STORING RESULTS	
235723582359	The problem description, or "prompt," is stored for further use. Alongside the prompt, additional datis preserved, including:	ıta
2360	The full directed graph.	
2361 2362	• The cost table.	
2363	• The feasible path.	
2364 2365	 Metadata such as the chosen start and end cities and the budget. 	
2366 2367	This stored information aids in verifying the correctness of solutions generated by LLMs.	
2368 2369	I.5 PARAMETERS	
2370 2371	In our experiments, we set the number of unique cities to 5 and the graph density to 0.5. We then value the total number of cities (S) in the graph and the number of transportation methods (A) to evaluate the set (S) in the graph and the number of transportation methods (A) to evaluate (S) in the graph (S) in the graph and the number of transportation methods (A) to evaluate (S) in the graph and the number of transportation methods (A) to evaluate (S) in the graph and the number of transportation methods (A) to evaluate (S) in the graph and the number of transportation methods (A) to evaluate (A) in the graph and the number of transportation methods (A) to evaluate (A) in the graph and (A) i	
2372 2373	the model's performance.	
2374	I.6 Failure Patterns	

We outline some common failure patterns below:

2377

2378

23792380

2381

2382

2384

2385

2386

2387 2388

2389

2391

2392

2394

2395

2396

239723982399

2400

2401

2402

2403

2404 2405

240624072408

2409

2410 2411

2412

2413

2414

2415

2416

2417

2418 2419

2420

2421 2422

2423

2424

2425

2426

2427

2428

2429

1. The model repeatedly samples travel options but ultimately concludes that the problem is infeasible after several attempts.

```
40 response
To meet the constraints of starting at 'Greensboro', ending at
'Minneapolis', not exceeding a budget of $68, and visiting at least 5
unique cities, we need to carefully select travel options. Here's a
potential travel plan using the provided travel connections:
 ''python
('Greensboro', 'Boise', 'flight'),
                                     # Cost: $22
('Boise', 'Louisville', 'flight'), # Cost: $21
('Louisville', 'Mesa', 'motorhome'), # Cost: $30
('Mesa', 'Minneapolis', 'flight'), #Cost:$56
Total cost: $22 + $21 + $30 + $56 = $129 (This exceeds the budget)
Attempting to adjust for budget:
Upon further investigation and adjustment to remain under budget:
'''python
('Greensboro', 'Boise', 'flight'),
                                      # Cost: $22
('Boise', 'Louisville', 'flight'), # Cost: $21
('Louisville', 'Minneapolis', 'motorhome') # Cost: $26
Total cost: $22 + $21 + $26 = $69 (This now exceeds the adjusted budget
by $1) The provided constraints and connections do not allow a feasible
plan under the budget given the stringent constraints listed, without
exceeding the budget or revisiting cities which could potentially
minimize costs. Please let me know if any corrections to constraints or
route connections are permitted.
```

```
Claude 3.7 response
Cost: $28 + $27 + $39 + $39 + $34 + $59 = $226  (still over budget)
One final attempt:
'''python
[('Irvine', 'Pittsburgh', 'flight'),
('Pittsburgh', 'Kansas City', 'train'),
('Kansas City', 'Memphis', 'taxi'),
('Memphis', 'Atlanta', 'flight'),
('Atlanta', 'Charlotte', 'train'),
('Charlotte', 'Austin', 'motorhome'),
('Austin', 'Oklahoma City', 'taxi')]
I've made several attempts to find a valid route within budget, but the
constraints make it challenging. The best route I found is:
'''python
[('Irvine', 'Spokane', 'motorhome'),
('Spokane', 'Atlanta', 'train'),
('Atlanta', 'Charlotte', 'train'),
('Charlotte', 'Kansas City', 'train'),
('Kansas City', 'Pittsburgh', 'train'),
('Pittsburgh', 'Wichita', 'train'),
('Wichita', 'Oklahoma City', 'motorhome')]
```

2431

2432

2464 2465

246624672468

2. Hallucination: the model either hallucinates the total cost or the number of unique cities visited.

2433 2434 Claude 3.7 response 2435 2436 I need to plan a trip from Charlotte to Chicago with at least 5 2437 unique cities, staying within a \$100 budget. 2438 2439 Let me analyze the connections and find an efficient route: 2440 1. Starting from Charlotte, I'll look for affordable connections. 2441 2. I need to find a path that visits at least 5 cities (including 2442 Charlotte and Chicago). 2443 3. Each leg needs to be cost-effective to stay under \$100. 2444 2445 After analyzing the possible routes, here's my plan: 2446 '''python 2447 [('Charlotte', 'Oklahoma City', 'motorhome'), 2448 ('Oklahoma City', 'Austin', 'ferry'), 2449 ('Austin', 'Aurora', 'motorhome'), ('Aurora', 'Chicago', 'hyperloop')] 2450 2451 2452 Cost calculation: 2453 - Charlotte to Oklahoma City (motorhome): \$16 2454 - Oklahoma City to Austin (ferry): \$32 2455 - Austin to Aurora (motorhome): \$22 - Aurora to Chicago (hyperloop): \$43 2456 2457 Total cost: \$16 + \$32 + \$22 + \$43 = \$93, which is under the \$100 2458 budget. 2459 2460 This plan visits 5 unique cities: Charlotte, Oklahoma City, Austin, Aurora, and Chicago, satisfying the requirement to visit at least 5 2461 cities including start and end. 2462 2463

However, 16 + 32 + 22 + 43 = 103, not 93.

```
2469
                40 response
2470
2471
                To create a cost-effective and efficient travel plan that adheres to
2472
                the constraints, I'll map out a feasible path from Phoenix to
                Milwaukee, ensuring we visit at least 5 unique cities within a
2473
                budget of $115. After examining the connections and costs provided,
2474
                here is a suitable travel itinerary:
2475
                '''python
2476
                 ('Phoenix', 'Nashville', 'tram'),
2477
                 ('Nashville', 'Madison', 'hyperloop'), # $24
2478
                 ('Madison', 'Milwaukee', 'tram')
2479
2480
2481
                This plan visits 5 cities, including Phoenix and Milwaukee, and
2482
                keeps the trip cost at a total of $109, under the budget constraints
                of $115.
2483
```

Claude response [...] The costs break down as: - Lexington to Greensboro by car: \$25 - Greensboro to Plano by car: \$19 - Plano to Lexington by flight: \$19 - Lexington to Tacoma by car: \$54 Total cost: \$117 This plan visits 5 unique cities (Lexington, Greensboro, Plano, Lexington again, and Tacoma) while exactly meeting the budget constraint of \$117.

The above plans visit only four unique cities.

Gemini performs very poorly on the travel planning task and tends to output a travel plan without reasoning, and the generated plans do not adhere to either the budget constraint or the number of unique cities.