

# S1-Bench: A Simple Benchmark for Evaluating System 1 Thinking Capability of Large Reasoning Models

Anonymous ACL submission

## Abstract

We introduce S1-Bench, a novel benchmark designed to evaluate the performance of Large Reasoning Models (LRMs) on simple tasks that favor intuitive *system 1* thinking rather than deliberative *system 2* reasoning. While LRMs have achieved significant breakthroughs in complex reasoning tasks through explicit chains of thought, their heavy reliance on *system 2* thinking may limit their *system 1* thinking capabilities. However, there is a lack of an appropriate benchmark for evaluating LRM’s *system 1* thinking capabilities. To fill this gap, S1-Bench introduces a suite of simple, diverse, and natural questions across multiple domains and languages, specifically designed to assess LRMs’ performance on questions more suitable for *system 1*. We conduct extensive evaluations across 28 LRMs, revealing their inefficiency, inadequate accuracy, and limited robustness when handling simple questions. Additionally, we observe a gap between their difficulty perception and generation length. Overall, this work paves the way toward dual-system compatibility in the development of LRMs<sup>1</sup>.

## 1 Introduction

*“Simplicity is the ultimate sophistication.”*  
— Leonardo da Vinci

Recent advances in Large Reasoning Models (LRMs), notably OpenAI’s o1/o3 (OpenAI, 2024) and the DeepSeek-R1 (Guo et al., 2025) series, have propelled the development of Large Language Models (LLMs). Unlike traditional LLMs that exhibit intuitive, heuristic *system 1* thinking, LRMs demonstrate deliberate and analytical *system 2* reasoning (Qu et al., 2025a; Li et al., 2025b) by explicitly generating external chain-of-thought (COT) (Wei et al., 2022) before producing final answers. Through sophisticated strategies such as self-reflection and multi-path exploration (Li et al.,

<sup>1</sup>The code and benchmark can be found in [Software](#) and [Data](#).

Benchmark	Cross Domain	Realistic Scenarios	Multi-lingual	Acc.
AIME	✗	✓	✗	6.67
GPQA	✓	✓	✗	24.94
Olympiad-Bench	✓	✓	✓	27.94
AMC	✗	✓	✗	31.88
MATH	✗	✓	✗	58.30
MMLU	✓	✓	✗	66.27
GSM8K	✗	✓	✗	87.45
ASDIV	✗	✓	✗	97.51
GSM8K-zero	✗	✗	✗	77.98
RoR-Bench	✗	✗	✗	14.24
<b>S1-Bench (ours)</b>	✓	✓	✓	<b>100.00</b>

Table 1: Characteristics of S1-Bench, with “Acc.” representing the average accuracy of four 7-9B LLMs. See Appendix A.1 for more details.

2025a; Yeo et al., 2025), LRMs can achieve strong performance in tasks that require *system 2* thinking, including advanced mathematical and competition-level problems (Yang et al., 2025a).

However, there remains a lack of appropriate benchmarks for evaluating LRMs’ *system 1* thinking capabilities. Not all real-world problems require *system 2* reasoning. The capacity to dynamically identify simple questions and address them efficiently contributes to both resource optimization and improved user satisfaction for LRMs. Nevertheless, current benchmarks either overemphasize difficulty, are simple yet lack domain diversity, or are only not hard for humans but involve unrealistic adversarial designs. Table 1 presents a collection of recent benchmarks aimed at mitigating the overthinking in LRMs (Sui et al., 2025). The majority of these benchmarks are of high difficulty. For example, AIME and GPQA (MAA Committees; Rein et al., 2024) achieve less than 30% accuracy on conventional small LLMs, which are inherently more suitable for *system 2* reasoning. Although some simple mathematical benchmarks are easy enough, such as GSM8K and ASDIV (Cobbe et al.,

2021; Miao et al., 2021), they often suffer from limited domain variety. Furthermore, some tasks that pose little challenge to humans but incorporate adversarial elements tend to lack relevance to realistic scenarios, such as GSM8K-zero (Chiang and Lee, 2024), which includes the correct answer in the questions. Thus, a benchmark to assess the *system 1* thinking capability of LRMs is still lacking, further hindering our understanding of LRMs’ cognitive flexibility between the two systems (Ziabari et al., 2025; Qu et al., 2025a).

To fill this research gap, we introduce the **System 1 Thinking Capability Benchmark (S1-Bench)**, which measures the performance of LRMs across various simple tasks that commonly encountered in real-world applications. S1-Bench has the following three characteristics: (1) **Simple**. The questions are not hard for humans and can be easily answered by LLMs. LLMs with 7-9B parameters can robustly provide correct answers through direct responses when sampled across multiple temperatures. (2) **Diverse**. S1-Bench is not limited to simple reasoning problems; it encompasses four major categories and 28 subcategories in two languages (English and Chinese), including reasoning problems, commonsense knowledge, instruction following, and analytical problems. (3) **Natural**. The questions are clear, without any misleading elements or ambiguities, ensuring they can be answered intuitively.

We conduct extensive evaluations on S1-Bench across 28 LRMs, yielding the following key findings: (1) Current LRMs exhibit inefficiency and lack *system 1* thinking capabilities across all types of questions, with average output lengths 15.5 times longer than small LLMs on S1-Bench. (2) Despite employing deep reasoning, several LRMs exhibit under-accuracy and limited robustness on simple questions. (3) LRMs exhibit “*gut moment*” at the beginning of some reasoning processes, showing gut feelings about task difficulty. Yet, even when recognizing a question’s simplicity, LRMs often fail to produce shorter responses—revealing a gap between their difficulty awareness and generation behavior. These findings emphasize the significant distance LRMs must traverse to become powerful dual-system compatible models.

Our contributions can be summarized as follows:

- To the best of our knowledge, S1-Bench is the first benchmark to evaluate the *system 1* thinking capabilities of LRMs, which paves

the way for dual-system compatibility.

- We introduce a workflow for constructing a simple dataset for *system 1* evaluation.
- Extensive experiments reveal the inefficiency, under-accuracy, and limited robustness of LRMs on simple questions.
- We find that LRMs exhibit “*gut moment*” on simple problems, and reveal a gap between their difficulty perception and generation length.

## 2 Related work

### 2.1 Large Reasoning Models

Large Reasoning Models (LRMs), characterized by explicitly generating external thinking processes before final answers (Kumar et al., 2025b; Chen et al., 2025), achieve a paradigm shift from intuitive *system 1* thinking to deliberative *system 2* reasoning compared to traditional LLMs (Li et al., 2025b; Qu et al., 2025a), thus achieving superior performance on complex tasks. The development of recent LRMs has largely followed two main approaches: large-scale reinforcement learning (RL) and model distillation. Models trained via large-scale RL (Guo et al., 2025; Team, 2025b; Team et al., 2025b) leverage reward-based optimization to gradually incentivize deliberative reasoning. In contrast, distillation-based LRMs (OpenAI, 2024; Min et al., 2024; Team, 2025a; Ye et al., 2025; Muennighoff et al., 2025) acquire such abilities by transferring structured reasoning patterns from advanced teacher models.

### 2.2 Limitations of LRMs

While LRMs have shown significant performance gains through deliberate reasoning, rigid adherence to this overly cautious thinking can introduce new limitations. On one hand, intermediate reasoning steps can cause excessive token generation and unnecessary solving attempts (Chen et al., 2024b; Hashemi et al., 2025; Kumar et al., 2025a). On the other hand, LRMs’ performance can drop in specific contexts like safety scenarios (Jiang et al., 2025; Zhao et al., 2025) and role-playing (Feng et al., 2025). However, prior studies mainly evaluated LRMs on complex tasks that more suited for deliberative *system 2* thinking. Our work examines how deliberative reasoning impacts extremely simple problems better matched to intuition-driven *system 1* processing.

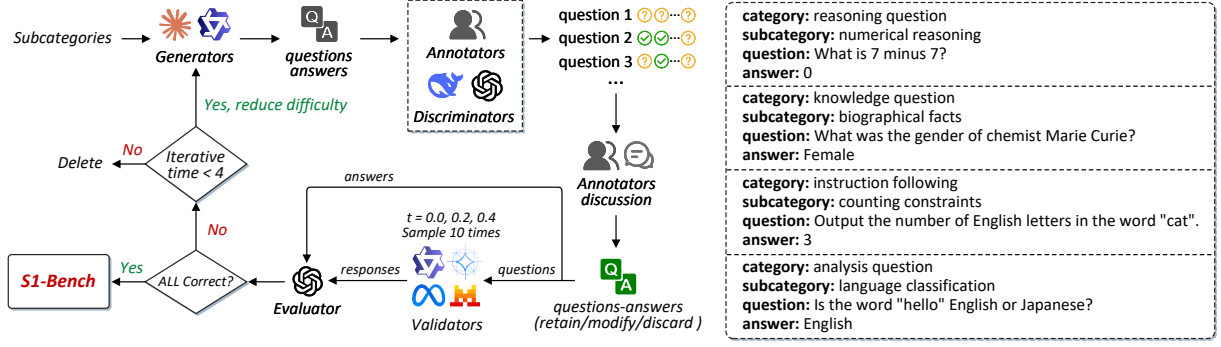


Figure 1: Construction workflow for S1-Bench and an illustrative example from each major category.

### 3 S1-Bench

We introduce S1-Bench, a bilingual, multi-domain benchmark designed to evaluate *system 1* thinking capability of LRM on extremely simple questions. These questions are easily solvable by traditional LLMs and not hard for humans. S1-Bench, which covers both *English* and *Chinese*, is organized into four major categories: *reasoning* (RSN), *knowledge* (KNO), *instruction following* (IF) and *analysis* (ANA), representing major dimensions commonly employed in LLM capability evaluation (Zheng et al., 2023; Chang et al., 2024).

This section begins with how simplicity is ensured, then the detailed construction workflow for S1-Bench, and concludes with an overview of the dataset statistics. Figure 1 shows the construction workflow and an illustrative example per category.

#### 3.1 How to Ensure Simplicity?

We ensure questions are simple and suitable for *system 1* thinking through the following two aspects.

##### 3.1.1 A Priori Simplicity Constraints

We begin by generating question-answer pairs through collaboration between humans and LLMs. Each pair is required to satisfy both the general and the category-specific simplicity criteria.

The general simplicity criteria requires that: (1) Questions must be naturally and clearly expressed, unambiguous, and free of intentional traps. (2) Answers must be unique or easily falsifiable (e.g., providing a three-letter English word).

The category-specific simplicity criteria are as follows. **RSN**: Limited to problems solvable with minimal reasoning or intuition. **KNO**: Restricted to common knowledge with unique, verifiable answers from sources like Wikipedia. **IF**: Involve straightforward instructions without strict formatting requirements. **ANA**: Limited to questions

whose answers can be directly inferred from the prompt, such as binary classification. These constraints ensure all questions remain straightforward for human respondents.

##### 3.1.2 A Posteriori Simplicity Verification

Due to the biases existing between language models and humans (Gallegos et al., 2024), questions that are simple for humans may be difficult for LLMs. Therefore, we introduce additional posteriori verification to ensure that questions are simple enough to be correctly and robustly answered by smaller LLMs from different families.

#### 3.2 Construction Workflow

**Subcategory Preparation.** To ensure diversity, we refer to the subcategories included in existing benchmarks (e.g., MMLU, IFEval, and GSM8K) and evaluation surveys (Chang et al., 2024) to select, merge, or design subcategories for S1-bench, ensuring that each meets the simplicity requirements. The definition and example question for each subcategory can be found in Appendix A.2.

**Implementation of A Priori Simplicity.** First, we use two data *generators*<sup>2</sup> to create 100 initial bilingual question-answer pairs for each candidate subcategory. The data generation prompt (see Appdiix A.3) explicitly incorporates the subcategory definitions, along with both the general and category-specific simplicity criteria, while also aiming to ensure diversity in the generated questions. Second, these question-answer pairs are then independently evaluated by three annotators and two quality *discriminators*<sup>3</sup> according to the general and category-specific simplicity criteria (see Appdiix A.3 for prompt of discriminators), resulting in five evaluation outcomes per pair. The three an-

<sup>2</sup>We select Claude-3.7-Sonnet and Qwen2.5-72B-Instruct.

<sup>3</sup>We select GPT-4o and DeepSeek-V3-241226.

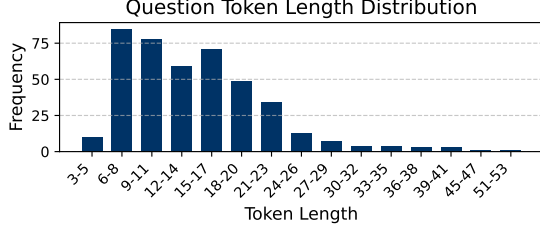


Figure 2: Statistical distribution of token counts for S1-Bench questions.

notators are experienced graduate students familiar with LLMs and well-acquainted with the goals of S1-Bench. **Finally**, based on these evaluation outcomes, three annotators discuss and collectively decide whether to retain, modify, or discard each question.

**Implementation of A Posteriori Simplicity.** First, each question obtained from the previous stage is input into the small LLM *validators*<sup>4</sup> with 7~9 B parameters. For each question, we sample 10 answers at three different temperature settings (0, 0.2, and 0.4), resulting in a total of 30 responses per question. These responses are then individually evaluated for correctness using GPT-4o. **Second**, if all 30 sampled responses are correct, the question is accepted into S1-Bench. Otherwise, the question is returned to the *generators*, where a difficulty-reduction prompt (see Appendix 10) is applied to simplify it. The simplified questions then undergo the same subsequent process. **Finally**, questions fail to meet the full-accuracy criterion (i.e., 30 out of 30 correct) after three rounds of difficulty reduction are excluded from the workflow.

The final S1-Bench comprises questions that satisfy both human-based a priori simplicity constraints and LLM-based a posteriori simplicity verification.

### 3.3 Benchmark Statistics

S1-Bench comprises 422 question-answer pairs across four major categories and 28 subcategories, balanced with 220 English and 202 Chinese questions. Figure 2 shows the token length distribution, with questions averaging 14.46 tokens. To ensure that the a posteriori verification process does not introduce simplicity only tailored to the small LLM validator, we evaluate S1-Bench on five additional LLMs and on Qwen3 Family with reasoning modes disabled. As shown, even the 1.7B model achieves

<sup>4</sup>We select four small LLMs: Qwen2.5-7B, Llama3.1-8B, Mistral-8B, and Gemma2-9B. The full model IDs are detailed in Table B.1.

Model	t=0.0	t=0.2	t=0.4	Tokens
Gemma2-9B	100.00	100.00	100.00	38.77
Llama3.1-8B	100.00	100.00	100.00	42.00
Mistral-8B	100.00	100.00	100.00	44.38
Qwen2.5-7B	100.00	100.00	100.00	42.81
DeepSeek-v3	100.00	100.00	100.00	79.53
Llama3.3-70B	100.00	99.76	99.76	53.71
Qwen2.5-14B	99.74	99.76	99.76	40.00
Qwen2.5-32B	99.98	99.98	99.98	43.17
Qwen2.5-72B	100.00	100.00	100.00	44.61
Qwen3-32B (w/o think)	100.00	100.00	100.00	103.30
Qwen3-14B (w/o think)	100.00	100.00	100.00	86.35
Qwen3-8B (w/o think)	100.00	100.00	99.76	90.54
Qwen3-1.7B (w/o think)	98.10	97.16	95.73	114.32

Table 2: Average accuracy (acc@k) and response token count of different LLMs, each sampled 10 times at three temperature settings.

over 98% accuracy at temperature 0.

## 4 Main Experiment

### 4.1 Baseline Models and Configurations

We evaluated 28 different LRMs, which are explicitly trained to first respond with a *thinking process*, and then generate a *final answer*. These LRMs include open-source families, such as DeepSeek (Guo et al., 2025), Qwen (Yang et al., 2025a), Nemotron (Bercovich et al., 2025), Light-R1 (Wen et al., 2025), s1.1 (Muennighoff et al., 2025), EXAONE (Research et al., 2025), and Sky-T1 (Griggs et al., 2025), as well as closed-source Hunyuan-T1 (Tencent, 2025), spanning from tiny (1.5B) to large (671B) parameter sizes<sup>5</sup>. Notably, OpenAI’s o-series models are not included as they do not disclose thinking processes to users. For each model, we consider two sets of generation configurations: **Greedy** sampling with temperature  $t=0$ ; **Top-p** sampling with temperature  $t=0.6$ ,  $\text{top-p}=0.95$  and sampling size  $k=5$ . Only top-p sampling results are reported in the main text; greedy decoding results are provided in the Appendix C.1.

### 4.2 Evaluation Metrics

**Format Metrics.** To assess the formatting quality of LRM responses, we compute the proportion of responses that satisfy the following two formatting criteria (averaged over 5 runs for top-p sampling). **S-Corr (Strict Format Correctness Rate):** In general, an end thinking marker (e.g., `</think>`) is expected to separate the thinking process from the non-empty final answer. S-Corr measures the proportion of responses that satisfy this criterion.

<sup>5</sup>Model details are presented in Appendix B.1.

Model ID	Size	Loose Format pass@1↑ acc@k↑		Strict Format pass@1↑ acc@k↑		L-Corr ↑	S-Corr ↑	Tokens ↓
Validator LLMs	7-9B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	–	–	42.00
Qwen3-A22B	235B	99.91	99.76	99.91	99.76	<b>100.00</b>	<b>100.00</b>	701.65
Qwen3-A3B	30B	99.95	99.76	99.95	99.76	<b>100.00</b>	<b>100.00</b>	638.40
QwQ-32B	32B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	720.10
Qwen3-32B	32B	99.91	99.53	99.91	99.53	99.91	99.91	668.69
Qwen3-14B	14B	99.95	99.76	99.95	99.76	99.95	99.95	582.99
Qwen3-8B	8B	99.95	99.76	99.95	99.76	99.95	99.95	657.76
Qwen3-1.7B	1.7B	99.34	97.39	99.34	97.39	99.81	99.81	595.90
Hunyuan-T1	–	99.91	99.53	99.91	99.53	<b>100.00</b>	<b>100.00</b>	542.31
DS-R1	671B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	646.40
DS-R1-70B	70B	99.48	97.39	99.38	96.92	<b>100.00</b>	99.91	453.81
DS-R1-32B	32B	99.72	98.82	99.72	98.82	<b>100.00</b>	<b>100.00</b>	429.91
DS-R1-14B	14B	99.57	97.87	99.57	97.87	<b>100.00</b>	<b>100.00</b>	475.46
DS-R1-8B	8B	97.44	97.16	97.39	97.16	99.76	99.53	452.11
DS-R1-7B	7B	95.21	85.78	95.21	85.78	99.24	99.24	454.55
DS-R1-1.5B	1.5B	<b>81.47</b>	<b>54.50</b>	81.47	<b>54.50</b>	97.58	97.58	489.54
Sky-T1-32B	32B	98.82	94.79	94.88	79.62	99.48	95.26	<b>163.00</b>
Nemotron-49B	49B	99.15	97.39	99.15	97.39	<b>100.00</b>	<b>100.00</b>	<b>362.54</b>
Nemotron-8B	8B	86.16	69.91	<b>79.81</b>	59.00	99.43	<b>84.31</b>	372.57
L-R1-32B	32B	97.87	91.00	94.74	79.62	98.91	95.07	<b>1095.36</b>
L-R1-32B-DS	32B	99.57	98.10	99.57	98.10	99.81	99.81	524.12
L-R1-14B-DS	14B	99.05	95.97	99.05	95.97	99.19	99.19	693.19
L-R1-7B-DS	7B	94.64	83.65	94.64	83.65	99.76	99.67	496.47
s1.1-32B	32B	99.53	98.34	99.48	98.10	99.57	99.53	998.00
s1.1-14B	14B	97.63	93.60	97.25	91.94	97.77	97.39	839.86
s1.1-7B	7B	96.68	88.39	88.58	63.98	<b>97.11</b>	<b>88.96</b>	711.49
EXAONE-32B	32B	97.06	94.08	97.06	94.08	99.81	99.81	800.56
EXAONE-7.8B	7.8B	88.15	75.12	87.82	74.41	98.72	98.06	1046.87
EXAONE-2.4B	2.4B	<b>72.42</b>	<b>56.16</b>	<b>72.32</b>	<b>56.16</b>	<b>97.44</b>	97.25	<b>1593.96</b>

Table 3: Main results in the top-p sampling setting on the S1-Bench, sorted by model family. **Bold teal** marks best performance, **teal** second best, **bold burgundy** worst, and **burgundy** second worst.

**L-Corr (Loose Format Correctness Rate):** LRMs may occasionally generate responses with endless thinking. L-Corr quantifies the proportion of responses that do not exhibit this failure mode. Detailed format types are given in Appendix B.4.

**Efficiency Metrics.** We calculate the average token counts for responses (*Tokens*) except for those generate endless thinking. Token counts are obtained using the Qwen2.5 tokenizer.

**Accuracy Metrics.** We calculate accuracy metrics under both strict and loose formatting requirements, respectively. We use GPT-4o as the evaluator to assess the correctness of the responses<sup>6</sup>, with the evaluation prompt in Appendix B.2. For greedy sampling, we directly calculate the accuracy rate. For top-p sampling, we utilize two metrics: *Pass@1* and *Acc@k*. *Pass@1* follows DeepSeek-R1 (Guo et al., 2025), and *Acc@k* is the percentage of questions with all k answers correct. The two metrics use k=5, and their detailed definitions can be found in Appendix B.3. Notably, S-Corr \ L-Corr represents the upper bound for pass@1 and acc@5 under strict \ loose formatting requirements.

<sup>6</sup>If a final answer can be isolated, only the final answer is evaluated; otherwise, the entire response is assessed.

### 4.3 Main Results

Table 3 and Figure 8 presents the main results of LRMs on S1-Bench, revealing two key phenomena.

**LRMs exhibit significantly lower efficiency than LLMs on S1-Bench, and no clear correlation is observed between ART and model size.** We observed the following: **First**, state-of-the-art LRMs, such as DeepSeek-R1 and Qwen3, do not demonstrate a distinct advantage in efficiency. In contrast, Sky-T1-32B, which undergoes specific optimizations to mitigate overthinking using SimPO, achieves the highest efficiency. **Second**, The L-R1-DS 7B/14B/32B models are further post-trained from the DS-R1-7B/14B/32B models. The L-R1-DS models tend to produce longer responses, suggesting that while additional post-training may enhance the model’s capability for complex reasoning, it comes at the cost of response efficiency. **Finally**, the s1.1 models generate considerably longer responses than the DeepSeek-R1-Distilled models. Despite both models being trained solely with SFT to acquire long-COT reasoning ability, the DeepSeek-R1-Distilled models use 800K training samples, while the s1.1 models are trained on only 1K. This discrepancy suggests that the smaller training set may lead to superficial imitation of long reasoning patterns, resulting in verbose thinking on

Model ID	Size	S1-Bench-EN					S1-Bench-ZH					Avg
		RSN	KNO	IF	ANA	Avg	RSN	KNO	IF	ANA	Avg	
Gemma2-9B	9B	74.8	29.4	5.3	52.4	45.9	51.6	19.8	7.5	35.1	31.0	38.8
Llama3.1-8B	8B	91.0	35.4	12.4	61.9	56.0	44.0	28.3	15.2	18.7	26.7	42.0
Qwen2.5-7B	7B	65.5	46.3	6.4	49.6	46.5	50.5	46.6	9.8	36.9	38.8	42.8
Mistral-8B	8B	67.2	55.5	8.6	50.1	49.6	47.3	56.1	14.8	29.7	38.7	44.4
Column Avg	-	74.6	41.6	8.2	53.5	49.5	48.3	37.7	11.8	30.1	33.8	42.0
Sky-T1-32B	32B	<b>215.8</b>	<b>174.1</b>	<b>98.5</b>	<b>233.3</b>	<b>194.3</b>	<b>125.5</b>	<b>125.3</b>	<b>99.4</b>	145.5	<b>128.9</b>	<b>163.0</b>
Nemotron-49B	49B	599.7	587.6	396.5	526.1	540.4	232.9	157.3	235.5	<b>107.8</b>	168.8	362.5
Nemotron-8B	8B	561.0	585.1	458.0	303.1	462.6	369.5	326.0	288.1	166.7	273.5	372.6
DS-R1-32B	32B	421.8	504.4	414.7	521.1	473.7	362.2	385.6	343.1	408.8	382.2	429.9
DS-R1-8B	8B	472.2	528.9	530.7	462.7	491.2	521.9	404.4	266.2	395.5	409.4	452.1
DS-R1-70B	70B	464.1	<b>501.3</b>	378.5	536.1	484.0	450.8	450.2	328.4	416.7	420.9	453.8
DS-R1-7B	7B	447.5	623.9	353.8	510.0	495.5	446.5	463.2	339.5	373.0	409.4	454.5
DS-R1-14B	14B	503.7	674.7	367.3	494.2	519.0	452.0	465.4	375.3	405.8	428.0	475.5
DS-R1-1.5B	1.5B	480.8	584.7	417.4	577.2	529.1	493.0	497.4	329.8	423.1	446.0	489.5
L-R1-7B-DS	7B	568.1	667.1	501.7	566.3	580.3	444.8	454.6	344.1	366.4	405.0	496.5
L-R1-32B-DS	32B	574.5	706.6	647.6	632.8	636.3	431.2	367.0	377.1	418.7	402.2	524.1
Hunyuan-T1	-	561.6	693.8	380.9	435.0	521.2	676.8	553.8	505.1	523.8	565.3	542.3
Qwen3-14B	14B	700.4	639.5	<b>286.2</b>	575.0	579.8	730.4	557.2	403.1	586.0	586.5	583.0
Qwen3-1.7B	1.7B	790.4	720.6	399.9	526.2	624.6	689.8	563.6	406.4	545.9	564.7	595.9
Qwen3-A3B	30B	745.0	729.3	328.1	594.8	625.7	773.7	655.8	453.7	648.6	652.2	638.4
DS-R1	671B	786.1	723.8	711.4	529.2	672.5	727.3	638.5	607.9	533.9	617.9	646.4
Qwen3-8B	8B	853.7	753.1	394.4	629.5	683.2	749.2	623.8	459.3	624.0	630.0	657.8
Qwen3-32B	32B	805.7	774.2	356.9	645.5	674.7	780.2	695.2	446.6	645.3	662.1	668.7
L-R1-14B-DS	14B	951.0	1026.0	829.8	653.5	848.2	594.7	610.1	442.2	451.7	525.7	693.2
Qwen3-A22B	235B	925.3	864.3	487.2	605.7	734.5	803.3	713.4	487.2	611.3	665.9	701.7
s1.1-7B	7B	1039.5	840.8	1923.2	529.4	929.9	489.6	351.3	1034.3	332.4	475.6	711.5
QwQ-32B	32B	873.3	808.1	520.8	634.7	722.4	866.9	707.3	613.3	<b>667.7</b>	717.6	720.1
EXAONE-32B	32B	1323.7	1057.6	1537.0	711.6	1086.4	703.2	348.6	1302.9	<b>125.5</b>	490.3	800.6
s1.1-14B	14B	871.8	746.2	<b>2233.1</b>	708.1	960.2	654.6	546.0	1512.6	579.7	710.7	839.9
s1.1-32B	32B	1077.9	889.7	<b>2055.4</b>	781.7	1081.7	995.6	<b>765.2</b>	<b>1634.6</b>	666.5	<b>906.5</b>	998.0
EXAONE-7.8B	7.8B	1498.3	<b>1398.9</b>	1775.7	<b>882.4</b>	1303.8	<b>1410.3</b>	497.8	1633.1	205.0	767.0	1046.9
L-R1-32B	32B	<b>1614.0</b>	1217.3	1996.9	<b>930.1</b>	<b>1338.3</b>	1035.6	737.7	1240.7	610.2	835.3	<b>1095.4</b>
EXAONE-2.4B	2.4B	<b>1927.3</b>	<b>1426.2</b>	1200.1	825.7	<b>1320.7</b>	<b>2469.7</b>	<b>1622.6</b>	<b>2471.6</b>	<b>1511.2</b>	<b>1898.7</b>	<b>1594.0</b>
Column Avg	-	809.1	766.0	785.1	591.5	718.4	695.8	545.9	677.9	482.0	576.7	650.3
Improvement	-	× <b>10.8</b>	× <b>18.4</b>	× <b>96.0</b>	× <b>11.1</b>	× <b>14.5</b>	× <b>14.4</b>	× <b>14.5</b>	× <b>57.3</b>	× <b>16.0</b>	× <b>17.1</b>	× <b>15.5</b>

Table 4: Average response tokens in the top-p sampling setting on the S1-bench across two languages and four main categories. **Bold teal** marks best performance, **teal** second best, **bold burgundy** worst, and **burgundy** second worst. **Bold** represents the maximum Improvement value for each language.

simple questions.

**Several LRMs exhibit under-accuracy and limited robustness on simple questions.** First, our observations find that, despite employing deep reasoning, most LRMs tend to exhibit lower accuracy on simple questions compared to traditional LLMs. For example, DS-R1-1.5B and EXAONE-2.4B achieve just above 50% acc@k. **Second**, many LRMs struggle with robust correctness in top-p sampling, where acc@k is significantly lower than pass@1. This issue is particularly pronounced in smaller LRMs. For instance, DS-R1-1.5B achieved 81.47% pass@1 but only 54.50% acc@k.

## 5 Efficiency Analysis

### 5.1 Analysis across Question Types

To better understand the efficiency differences of LRMs across question types, we analyze the average response tokens across 4 main categories, 28

subcategories, and two languages. The results are displayed in Table 4 and Appendix C.2.

**LRMs exhibit a substantial increase in response length across all four major categories, 28 subcategories, and two languages.** As shown in Table 4, for each of the four major categories, the average response length of LRMs exceeds that of LLMs by more than a factor of ten. Response lengths also increase significantly across all subcategories (see Appendix C.2). This suggests that while LRMs are primarily trained on reasoning data to produce long CoT style responses, this stylistic pattern generalizes well across a wide range of question types. Moreover, 23 out of the 28 LRMs produce longer responses to questions in English than Chinese.

**LRMs exhibit the most significant increase in ART for instruction following questions and tend to over-explore when the solution space is vast.** As shown in Table 4, although small LLMs provide the most concise responses to instruction

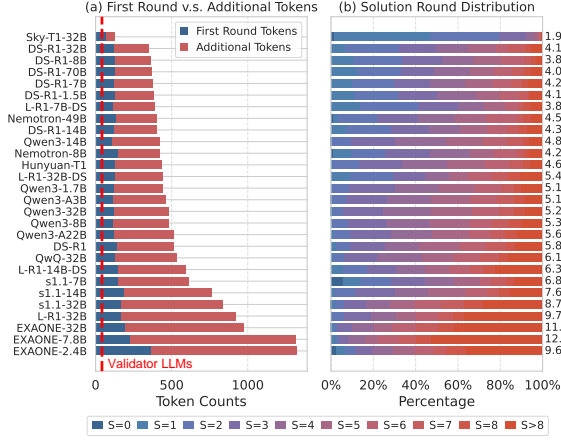


Figure 3: (a) Comparison of first round and additional token costs for each LRM. (b) Distribution of solution rounds for each LRM.

following questions, LRMs generate dramatically longer outputs—96.0 times longer in English and 57.3 times longer in Chinese than small LLMs. To investigate the cause, we further analyze the subcategories of instruction following questions. As shown in Appendix C.2, average tokens is notably longer in the subcategories of *length constraints*, *character constraints*, and *sentence constraints*. These three question types share a similar characteristic: their correctness is easy to verify, but the solution space is vast. We find that, although the model quickly identifies a correct answer, it becomes trapped in the search space, continually exploring alternatives and failing to stop in time. A case can be seen in Table 21. This phenomenon is more pronounced in families with lower efficiency, such as s1.1 and EXAONE.

## 5.2 Thinking Solution Analysis

To better understand the causes of inefficiency in LRMs on S1-Bench, we analyze the solution rounds of their thinking processes<sup>7</sup>. We first use DeepSeek-v3 to segment each thinking process into several solutions, each defined as a point at which LRMs explicitly arrives at a conclusion that matches the correct answer. We then compute the average token counts in the first solution. The detailed experimental setup is provided in Appendix C.3. Our analysis reveals the following:

**The token consumed in the first solution of LRMs significantly exceeds that of validator**

<sup>7</sup>We only analyze well-formatted thinking processes with correct final answers, as incorrect answers make it unclear whether LRMs are over-reasoning or under-reasoning, and malformed thinking processes cannot be precisely extracted.

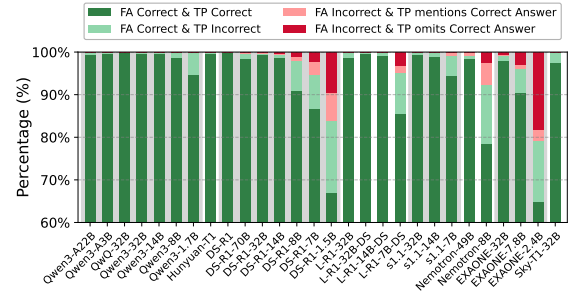


Figure 4: Distribution of the thinking process across four categories. FA and TP refer to Final Answer and Thinking Process, respectively. Green bars indicate cases where the final answer is correct, while red bars indicate cases where it is incorrect.

LRMs, as shown in Figure 3 (a). This suggests that LRMs may involve unnecessary reasoning steps in each solution, which could be one of the reasons for their inefficiency.

**The primary reason for efficiency gaps between LRMs lies in the number of redundant solution rounds they generate, rather than the token cost in the initial round.** As shown in Figure 3 (a), although total thinking token counts vary widely across LRMs, their token counts in the initial round are similar and only account for a small fraction of the total. Figure 3 (b) further shows the distribution of solution rounds on S1-Bench, revealing that LRMs with longer thinking processes tend to generate more solution round, and this redundancy greatly increases computational cost. Furthermore, further experiments reveal that the redundancy in the reasoning process gradually increases over time. Appendix C.4 presents the experimental details.

## 6 Error Analysis

This section analyzes the errors made in the thinking process. Specifically, we utilize DeepSeek-v3 to categorize the responses of LRMs into four cases and compute the corresponding proportions: (1) Final answer correct; thinking process entirely accurate. (2) Final answer correct; thinking process contains intermediate errors. (3) Final answer incorrect; correct answer mentioned in thinking process. (4) Final answer incorrect; correct answer never mentioned in thinking process. The classification details are in Appendix C.5; results are shown in Figure 4. Key findings include:

**Lower-accuracy LRMs tend to produce less reliable reasoning chains; even when they arrive at**

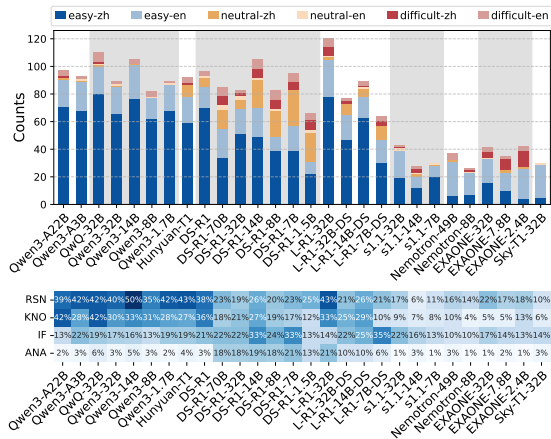


Figure 5: Top: Count of “gut moments” across models. Bottom: Probability of “gut moments” by question type.

the correct final answer, their intermediate steps often contain errors (light green). LRM’s with high accuracy (e.g., DS-R1) show almost no flawed reasoning steps, whereas those with lower accuracy (e.g., DS-R1-1.5B) often generate incorrect intermediate conclusions, further indicating that they lack robust reasoning ability.

Although LRM’s sometimes mention the correct answer during reasoning, they may deviate and ultimately produce incorrect final answers (light red). In one case, the LRM initially arrived at the correct answer but undermined it through excessive verification, a case can be seen in Table 24. In another case, the LRM directly denies the correct answer, a case can be seen in Table 23.

## 7 Gut Moment

We observe an intriguing phenomenon on S1-Bench: LRM’s sometimes show an early sense of question difficulty before solving, which we call the “gut moment”<sup>8</sup>. To explore this phenomenon, we prompt GPT-4o to classify the initial part of model responses (before the first “\n\n”) into four types based on its comment on difficulty: easy, neutral, difficult, and no comment. Figure 5 presents these classifications and their probabilities across four question types. Experimental details and cases are in Appendix C.6. This leads to the following observations:

**First**, all LRM’s show the “gut moment” phenomenon to varying degrees, which is more evident in the Qwen, DeepSeek, and Light-R1 families and

<sup>8</sup>Derived from “gut feeling,” meaning intuition-based judgment without analysis.

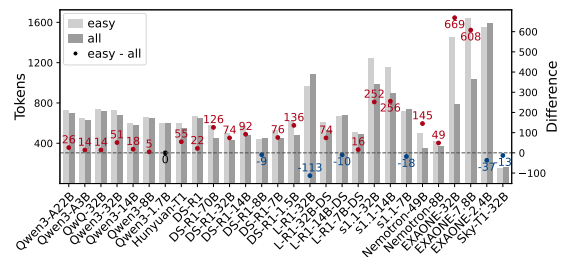


Figure 6: Average response tokens in the easy category vs. all samples. Dots show difference: easy minus all.

Hunyuan-T1. **Second**, LRM’s show stylistic differences in expressing “gut moment.” For example, the Qwen family often views questions as simple, whereas the DeepSeek-distilled models show more diverse difficulty comments. **Third**, some LRM’s show significantly stronger “gut moment” in Chinese than in English, such as the Qwen and DeepSeek families, likely due to a higher proportion of Chinese in their training data. **Finally**, the “gut moment” is most evident in reasoning questions and rarely appears in analytical questions, except in DeepSeek-distilled models.

To investigate whether the early sense of a question as “easy” leads to a corresponding reduction in response length, we compare the average response tokens for questions in the easy category versus all samples. The results are shown in Figure 6. Except for L-R1-32B, other LRM’s do not exhibit a noticeable decrease in response length when questions are viewed as “easy”; in fact, 21 out of 28 LRM’s showed an increase in response length under this condition. **This suggests a discrepancy between the LRM’s initial sense of difficulty and its generative behavior**, the causes and improvements of which warrant further investigation.

## 8 Conclusion

This paper introduces S1-Bench, the first benchmark designed to evaluate *system 1* thinking capabilities in LRM’s. We conduct extensive evaluations across 28 LRM’s, revealing their inefficiency, inadequate accuracy, and limited robustness when handling simple questions. Additionally, we observe “gut moment” and find a gap between their difficulty perception and generation length. Overall, this work paves the way toward dual-system compatibility in the development of LRM’s.

## Limitations

Although S1-Bench pioneers the evaluation of *system 1* thinking in LRMs, it still has several limitations. First, due to our emphasis on ensuring the uniqueness of each sample during dataset construction—for instance, including only one question for basic arithmetic operations such as addition, subtraction, and multiplication—the overall scale of the benchmark remains limited. As a next step, we plan to expand the scale of S1-Bench. Second, while recent months have seen a surge in newly released open-source LRMs, we have only evaluated 28 representative models and have not covered the full spectrum of available models. Lastly, we do not propose methods to improve the efficiency of LRMs on *system 1* tasks in this work; this will be the focus of our future research.

## References

- Pranjal Aggarwal and Sean Welleck. 2025. L1: Controlling how long a reasoning model thinks with reinforcement learning. *arXiv preprint arXiv:2503.04697*.
- AI-MO. 2024. [Amc 2023](#).
- Daman Arora and Andrea Zanette. 2025. Training language models to reason efficiently. *arXiv preprint arXiv:2502.04463*.
- Simon A Aytes, Jinheon Baek, and Sung Ju Hwang. 2025. Sketch-of-thought: Efficient llm reasoning with adaptive cognitive-inspired sketching. *arXiv preprint arXiv:2503.05179*.
- Akhiad Bercovich, Itay Levy, Izik Golan, Mohammad Dabbah, Ran El-Yaniv, Omri Puny, Ido Galil, Zach Moshe, Tomer Ronen, Najeeb Nabwani, Ido Shahaf, Oren Tropp, Ehud Karpas, Ran Zilberstein, Jiaqi Zeng, Soumye Singhal, Alexander Bukharin, Yian Zhang, Tugrul Konuk, and 113 others. 2025. [Llama-nemotron: Efficient reasoning models](#). *Preprint*, arXiv:2505.00949.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, and 1 others. 2024. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wangxiang Che. 2025. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*.
- Xiaoshu Chen, Sihang Zhou, Ke Liang, and Xinwang Liu. 2024a. Distilling reasoning ability from large language models with adaptive thinking. *arXiv preprint arXiv:2404.09170*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, and 1 others. 2024b. Do not think that much for  $2 + 3 = ?$  on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*.
- Cheng-Han Chiang and Hung-yi Lee. 2024. [Over-reasoning and redundant calculation of large language models](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 161–169, St. Julian’s, Malta. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Yingqian Cui, Pengfei He, Jingying Zeng, Hui Liu, Xianfeng Tang, Zhenwei Dai, Yan Han, Chen Luo, Jing Huang, Zhen Li, and 1 others. 2025. Stepwise perplexity-guided refinement for efficient chain-of-thought reasoning in large language models. *arXiv preprint arXiv:2502.13260*.
- Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang, Zengmao Wang, Ziwei Liu, Bo Du, and 1 others. 2025. Dynamic parallel tree search for efficient llm reasoning. *arXiv preprint arXiv:2502.16235*.
- Xiachong Feng, Longxu Dou, and Lingpeng Kong. 2025. Reasoning does not necessarily improve role-playing ability. *arXiv preprint arXiv:2502.16940*.
- Yichao Fu, Junda Chen, Siqi Zhu, Zheyu Fu, Zhongdongming Dai, Aurick Qiao, and Hao Zhang. 2024. Efficiently serving llm reasoning programs with certaintindex. *arXiv preprint arXiv:2412.20993*.
- Yichao Fu, Junda Chen, Yonghao Zhuang, Zheyu Fu, Ion Stoica, and Hao Zhang. 2025. Reasoning without self-doubt: More efficient chain-of-thought through certainty probing. In *ICLR 2025 Workshop on Foundation Models in the Wild*.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3):1097–1179.
- Tyler Griggs, Shiyi Cao, Dacheng Li, Shu Liu, Shishir G. Patil, Matei Zaharia, Joey Gonzalez, and Ion Stoica. 2025. [Think less, achieve more: Cut reasoning costs by 50% without sacrificing accuracy](#).

631	Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. DeepSeek-R1: Incentivizing reasoning capability in llms via reinforcement learning. <i>arXiv preprint arXiv:2501.12948</i> .	686
632		687
633		688
634		689
635		690
636		691
637	Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. 2024. Token-budget-aware llm reasoning. <i>arXiv preprint arXiv:2412.18547</i> .	692
638		693
639		694
640		695
641	Masoud Hashemi, Oluwanifemi Bamgbose, Sathwik Tejaswi Madhusudhan, Jishnu Sethumadhavan Nair, Aman Tiwari, and Vikas Yadav. 2025. Dna bench: When silence is smarter—benchmarking over-reasoning in reasoning llms. <i>arXiv preprint arXiv:2503.15793</i> .	696
642		697
643		698
644		699
645		700
646		701
647	Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. <i>arXiv preprint arXiv:2402.14008</i> .	702
648		703
649		704
650		705
651		706
652		
653		
654	Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. Measuring massive multitask language understanding. In <i>International Conference on Learning Representations</i> .	707
655		708
656		709
657		710
658		711
659	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset. <i>arXiv preprint arXiv:2103.03874</i> .	712
660		
661		
662		
663		
664	Bairu Hou, Yang Zhang, Jiabao Ji, Yujian Liu, Kaizhi Qian, Jacob Andreas, and Shiyu Chang. 2025. Thinkprune: Pruning long chain-of-thought of llms via reinforcement learning. <i>arXiv preprint arXiv:2504.01296</i> .	713
665		714
666		715
667		716
668		717
669	Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. Live-codebench: Holistic and contamination free evaluation of large language models for code. In <i>The Thirteenth International Conference on Learning Representations</i> .	718
670		719
671		720
672		721
673		722
674		723
675		
676	Fengqing Jiang, Zhangchen Xu, Yuetai Li, Luyao Niu, Zhen Xiang, Bo Li, Bill Yuchen Lin, and Radha Poovendran. 2025. Safechain: Safety of language models with long chain-of-thought reasoning capabilities. <i>arXiv preprint arXiv:2502.12025</i> .	724
677		725
678		726
679		727
680		728
681	Abhinav Kumar, Jaechul Roh, Ali Naseh, Marzena Karpinska, Mohit Iyyer, Amir Houmansadr, and Eugene Bagdasarian. 2025a. Overthink: Slowdown attacks on reasoning llms. <i>arXiv e-prints</i> , pages arXiv–2502.	729
682		730
683		731
684		732
685		
	Komal Kumar, Tajamul Ashraf, Omkar Thawakar, Rao Muhammad Anwer, Hisham Cholakkal, Mubarak Shah, Ming-Hsuan Yang, Phillip HS Torr, Salman Khan, and Fahad Shahbaz Khan. 2025b. Llm post-training: A deep dive into reasoning large language models. <i>arXiv preprint arXiv:2502.21321</i> .	733
		734
		735
		736
	Ayeong Lee, Ethan Che, and Tianyi Peng. 2025. How well do llms compress their own chain-of-thought? a token complexity approach. <i>arXiv preprint arXiv:2503.01141</i> .	737
		738
		739
	Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xiangxi Mo, Eric Tang, Sumanth Hegde, Kourosh Hakhmaneshi, Shishir G Patil, Matei Zaharia, and 1 others. 2025a. Llms can easily learn to reason from demonstrations structure, not content, is what matters! <i>arXiv preprint arXiv:2502.07374</i> .	
	Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Xinglin Wang, Bin Sun, Heda Wang, and Kan Li. 2024. Escape sky-high cost: Early-stopping self-consistency for multi-step reasoning. <i>arXiv preprint arXiv:2401.10480</i> .	
	Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, and 1 others. 2025b. From system 1 to system 2: A survey of reasoning large language models. <i>arXiv preprint arXiv:2502.17419</i> .	
	Baohao Liao, Yuhui Xu, Hanze Dong, Junnan Li, Christof Monz, Silvio Savarese, Doyen Sahoo, and Caiming Xiong. 2025. Reward-guided speculative decoding for efficient llm reasoning. <i>arXiv preprint arXiv:2501.19324</i> .	
	Ximing Lu, Seungju Han, David Acuna, Hyunwoo Kim, Jaehun Jung, Shrimai Prabhumoye, Niklas Muenighoff, Mostofa Patwary, Mohammad Shoenybi, Bryan Catanzaro, and 1 others. 2025. Retro-search: Exploring untaken paths for deeper and efficient reasoning. <i>arXiv preprint arXiv:2504.04383</i> .	
	Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. 2025. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning. <i>arXiv preprint arXiv:2501.12570</i> .	
	Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. 2025a. Reasoning models can be effective without thinking. <i>arXiv preprint arXiv:2504.09858</i> .	
	Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. 2025b. Cot-valve: Length-compressible chain-of-thought tuning. <i>arXiv preprint arXiv:2502.09601</i> .	
	MAA Committees. Aime problems and solutions. <a href="https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions">https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions</a> .	

740	Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su.	Jianshu She, Zhuohao Li, Zhemin Huang, Qi Li, Peiran	794
741	2021. A diverse corpus for evaluating and developing	Xu, Haonan Li, and Qirong Ho. 2025. Hawkeye:	795
742	english math word problem solvers. <i>arXiv preprint</i>	Efficient reasoning with model collaboration. <i>arXiv</i>	796
743	<i>arXiv:2106.15772</i> .	<i>preprint arXiv:2504.00424</i> .	797
744	Yingqian Min, Zhipeng Chen, Jinhao Jiang, Jie Chen,	Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wen-	798
745	Jia Deng, Yiwen Hu, Yiru Tang, Jiapeng Wang, Xi-	jing Zhang, Jiangze Yan, Ning Wang, Kai Wang, and	799
746	aoxue Cheng, Huatong Song, and 1 others. 2024.	Shiguo Lian. 2025a. Dast: Difficulty-adaptive slow-	800
747	Imitate, explore, and self-improve: A reproduction	thinking for large reasoning models. <i>arXiv preprint</i>	801
748	report on slow-thinking reasoning systems. <i>arXiv</i>	<i>arXiv:2503.04472</i> .	802
749	<i>preprint arXiv:2412.09413</i> .		
750	Niklas Muennighoff, Zitong Yang, Weijia Shi, Xi-	Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu,	803
751	ang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke	Yali Du, and Yulan He. 2025b. Codi: Compress-	804
752	Zettlemoyer, Percy Liang, Emmanuel Candès, and	ing chain-of-thought into continuous space via self-	805
753	Tatsunori Hashimoto. 2025. s1: Simple test-time	distillation. <i>arXiv preprint arXiv:2502.21074</i> .	806
754	scaling. <i>arXiv preprint arXiv:2501.19393</i> .		
755	Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin	Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu	807
756	Yang, Yujin Kim, and Se-Young Yun. 2025. Self-	Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, An-	808
757	training elicits concise reasoning in large language	drew Wen, Shaochen Zhong, Hanjie Chen, and 1	809
758	models. URL <a href="https://arxiv.org/abs/2502.20122">https://arxiv.org/abs/2502.20122</a> .	others. 2025. Stop overthinking: A survey on ef-	810
		ficient reasoning for large language models. <i>arXiv</i>	811
		<i>preprint arXiv:2503.16419</i> .	812
759	OpenAI. 2024. Learning to reason with LLMs. <a href="https://openai.com/index/learning-to-reason-with-llms/">https://openai.com/index/learning-to-reason-wit</a>	Alon Talmor, Jonathan Herzig, Nicholas Lourie, and	813
760	<a href="https://openai.com/index/learning-to-reason-wit">h-llms/</a> .	Jonathan Berant. 2019. Commonsenseqa: A question	814
761		answering challenge targeting commonsense knowl-	815
762	Rui Pan, Yinwei Dai, Zhihao Zhang, Gabriele Oliaro,	edge. In <i>Proceedings of the 2019 Conference of</i>	816
763	Zhihao Jia, and Ravi Netravali. 2025. Specreason:	<i>the North American Chapter of the Association for</i>	817
764	Fast and accurate inference-time compute via specu-	<i>Computational Linguistics: Human Language Tech-</i>	818
765	lative reasoning. <i>arXiv preprint arXiv:2504.07891</i> .	<i>nologies, Volume 1 (Long and Short Papers)</i> , pages	819
		4149–4158.	820
766	Arkil Patel, Satwik Bhattamishra, and Navin Goyal.	Amir Taubenfeld, Tom Sheffer, Eran Ofek, Amir Feder,	821
767	2021. Are nlp models really able to solve simple	Ariel Goldstein, Zorik Gekhman, and Gal Yona. 2025.	822
768	math word problems? In <i>Proceedings of the 2021</i>	Confidence improves self-consistency in llms. <i>arXiv</i>	823
769	<i>Conference of the North American Chapter of the</i>	<i>preprint arXiv:2502.06233</i> .	824
770	<i>Association for Computational Linguistics: Human</i>		
771	<i>Language Technologies</i> , pages 2080–2094.	Kimi Team, A Du, B Gao, B Xing, C Jiang, C Chen,	825
772	Xiaoye Qu, Yafu Li, Zhaochen Su, Weigao Sun, Jianhao	C Li, C Xiao, C Du, C Liao, and 1 others. 2025a.	826
773	Yan, Dongrui Liu, Ganqu Cui, Daizong Liu, Shuxian	Kimi k1. 5: Scaling reinforcement learning with llms.	827
774	Liang, Junxian He, and 1 others. 2025a. A survey of	URL <a href="https://arxiv.org/abs/2501.12599">https://arxiv.org/abs/2501.12599</a> .	828
775	efficient reasoning for large reasoning models: Lan-		
776	guage, multimodality, and beyond. <i>arXiv preprint</i>	Kimi Team, Angang Du, Bofei Gao, Bowei Xing,	829
777	<i>arXiv:2503.21614</i> .	Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun	830
		Xiao, Chenzhuang Du, Chonghua Liao, and 1 others.	831
778	Yuxiao Qu, Matthew YR Yang, Amrith Setlur,	2025b. Kimi k1. 5: Scaling reinforcement learning	832
779	Lewis Tunstall, Edward Emanuel Beeching, Ruslan	with llms. <i>arXiv preprint arXiv:2501.12599</i> .	833
780	Salakhutdinov, and Aviral Kumar. 2025b. Optimiz-		
781	ing test-time compute via meta reinforcement fine-	NovaSky Team. 2025a. Sky-t1: Train your own o1	834
782	tuning. <i>arXiv preprint arXiv:2503.07572</i> .	preview model within \$450. <a href="https://novasky-ai.github.io/posts/sky-t1">https://novasky-ai.</a>	835
		<a href="https://novasky-ai.github.io/posts/sky-t1">github.io/posts/sky-t1</a> .	836
783	David Rein, Betty Li Hou, Asa Cooper Stickland, Jack-	Qwen Team. 2025b. QwQ-32b: Embracing the power	837
784	son Petty, Richard Yuanzhe Pang, Julien Dirani, Ju-	of reinforcement learning. <a href="https://qwenlm.github.io/blog/qwq-32b/">https://qwenlm.githu</a>	838
785	lian Michael, and Samuel R. Bowman. 2024. GPQA:	<a href="https://qwenlm.github.io/blog/qwq-32b/">b.io/blog/qwq-32b/</a> .	839
786	A graduate-level google-proof q&a benchmark. In		
787	<i>First Conference on Language Modeling</i> .	Tencent. 2025. Reasoning efficiency redefined! meet	840
788	LG Research, Kyunghoon Bae, Eunbi Choi, Ki-	tencent’s ‘hunyuan-t1’—the first mamba-powered	841
789	bong Choi, Stanley Jungkyu Choi, Yemuk Choi,	ultra-large model. <a href="https://llm.hunyuan.tencent.com/#/Blog/hy-t1/">https://llm.hunyuan.tencen</a>	842
790	Seokhee Hong, Junwon Hwang, Hyojin Jeon, Ki-	<a href="https://llm.hunyuan.tencent.com/#/Blog/hy-t1/">t.com/#/Blog/hy-t1/</a> .	843
791	jeong Jeon, and 1 others. 2025. Exaone deep: Rea-		
792	soning enhanced language models. <i>arXiv preprint</i>	Guangya Wan, Yuqi Wu, Jie Chen, and Sheng Li.	844
793	<i>arXiv:2503.12524</i> .	2024. Dynamic self-consistency: Leveraging reason-	845
		ing paths for efficient llm sampling. <i>arXiv preprint</i>	846
		<i>arXiv:2408.17017</i> .	847

848	Junlin Wang, Shang Zhu, Jon Saad-Falcon, Ben Athiwaratkun, Qingyang Wu, Jue Wang, Shuaiwen Leon Song, Ce Zhang, Bhuwan Dhingra, and James Zou. 2025a. Think deep, think fast: Investigating efficiency of verifier-free inference-time-scaling methods. <i>arXiv preprint arXiv:2504.14047</i> .	902
849		903
850		904
851		
852		905
853		906
		907
		908
854	Xinglin Wang, Shaoxiong Feng, Yiwei Li, Peiwen Yuan, Yueqi Zhang, Chuyi Tan, Boyuan Pan, Yao Hu, and Kan Li. 2024. Make every penny count: Difficulty-adaptive self-consistency for cost-efficient reasoning. <i>arXiv preprint arXiv:2408.13457</i> .	909
855		910
856		911
857		912
858		
859	Yiming Wang, Pei Zhang, Siyuan Huang, Baosong Yang, Zhuosheng Zhang, Fei Huang, and Rui Wang. 2025b. Sampling-efficient test-time scaling: Self-estimating the best-of-n sampling in early decoding. <i>arXiv preprint arXiv:2503.01422</i> .	913
860		914
861		915
862		
863		916
864	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	917
865		918
866		919
867		
868		920
869		921
		922
		923
		924
870	Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu Tang, Xiaowei Lv, and 1 others. 2025. Light-R1: Curriculum sft, dpo and rl for long cot from scratch and beyond. <i>arXiv preprint arXiv:2503.10460</i> .	925
871		926
872		927
873		
874		
875	Heming Xia, Yongqi Li, Chak Tou Leong, Wenjie Wang, and Wenjie Li. 2025. Tokenskip: Controllable chain-of-thought compression in llms. <i>arXiv preprint arXiv:2502.12067</i> .	928
876		929
877		930
878		931
		932
879	Jingxian Xu, Mengyu Zhou, Weichang Liu, Hanbing Liu, Shi Han, and Dongmei Zhang. 2025a. Twt: Thinking without tokens by habitual reasoning distillation with multi-teachers’ guidance. <i>arXiv preprint arXiv:2503.24198</i> .	933
880		934
881		935
882		936
883		937
884	Yuhui Xu, Hanze Dong, Lei Wang, Doyen Sahoo, Junnan Li, and Caiming Xiong. 2025b. Scalable chain of thoughts via elastic reasoning. <i>arXiv preprint arXiv:2505.05315</i> .	938
885		939
886		940
887		941
888	Yuchen Yan, Yongliang Shen, Yang Liu, Jin Jiang, Mengdi Zhang, Jian Shao, and Yueting Zhuang. 2025. Infythink: Breaking the length limits of long-context reasoning in large language models. <i>arXiv preprint arXiv:2503.06692</i> .	942
889		943
890		944
891		945
892		
893	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. <i>arXiv preprint arXiv:2505.09388</i> .	946
894		947
895		948
896		949
897		950
898	Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Zheng Lin, Li Cao, and Weiping Wang. 2025b. Dynamic early exit in reasoning models. <i>arXiv preprint arXiv:2504.15895</i> .	
899		
900		
901		
	Junjie Yang, Ke Lin, and Xing Yu. 2025c. Think when you need: Self-adaptive chain-of-thought learning. <i>arXiv preprint arXiv:2504.03234</i> .	
	Wang Yang, Xiang Yue, Vipin Chaudhary, and Xiaotian Han. 2025d. Speculative thinking: Enhancing small-model reasoning with large model guidance at inference time. <i>arXiv preprint arXiv:2504.12329</i> .	
	Wenkai Yang, Shuming Ma, Yankai Lin, and Furu Wei. 2025e. Towards thinking-optimal scaling of test-time compute for llm reasoning. <i>arXiv preprint arXiv:2502.18080</i> .	
	Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. 2025. LIMO: Less is more for reasoning. <i>arXiv preprint arXiv:2502.03387</i> .	
	Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. 2025. Demystifying long chain-of-thought reasoning in llms. <i>arXiv preprint arXiv:2502.03373</i> .	
	Bin Yu, Hang Yuan, Yuliang Wei, Bailing Wang, Weizhen Qi, and Kai Chen. 2025a. Long-short chain-of-thought mixture supervised fine-tuning eliciting efficient reasoning in large language models. <i>arXiv preprint arXiv:2505.03469</i> .	
	Zhaojian Yu, Yinghao Wu, Yilun Zhao, Arman Cohan, and Xiao-Ping Zhang. 2025b. Z1: Efficient test-time scaling with code. <i>arXiv preprint arXiv:2504.00810</i> .	
	Jintian Zhang, Yuqi Zhu, Mengshu Sun, Yujie Luo, Shuofei Qiao, Lun Du, Da Zheng, Huajun Chen, and Ningyu Zhang. 2025. Lightthinker: Thinking step-by-step compression. <i>arXiv preprint arXiv:2502.15589</i> .	
	Weixiang Zhao, Xingyu Sui, Jiahe Guo, Yulin Hu, Yang Deng, Yanyan Zhao, Bing Qin, Wanxiang Che, Tat-Seng Chua, and Ting Liu. 2025. Trade-offs in large reasoning models: An empirical analysis of deliberative and adaptive reasoning over foundational capabilities. <i>arXiv preprint arXiv:2503.17979</i> .	
	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. <a href="#">Judging LLM-as-a-judge with MT-bench and chatbot arena</a> . In <i>Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track</i> .	
	Alireza S Ziabari, Nona Ghazizadeh, Zhivar Sourati, Farzan Karimi-Malekabadi, Payam Piray, and Morteza Dehghani. 2025. Reasoning on a spectrum: Aligning llms to system 1 and system 2 thinking. <i>arXiv preprint arXiv:2502.12470</i> .	

## A More Information of S1-Bench Construction

### A.1 Benchmark Statistics

We survey studies on improving the efficiency of LRMs, as there is potential overlap between these studies and the technical approaches aimed at enhancing *system 1* thinking in LRMs. Table 7 presents the results of our survey. We compile the benchmarks used in these studies for evaluation, that are typically used to verify whether models achieve efficiency improvements. Benchmarks that appear more than four times include: MATH500 (Hendrycks et al., 2021b), GSM8K (Cobbe et al., 2021), AIME24/25 (MAA Committees), GPQA (Rein et al., 2024), AMC23 (AI-MO, 2024), MMLU (Hendrycks et al., 2021a), Olympiad-Bench (He et al., 2024), SVAMP (Patel et al., 2021), LiveCodeBench (Jain et al., 2025), and Common-SenseQA (Talmor et al., 2019).

The accuracy shown in Table 1 is the average result of the four models, Qwen2.5-7B, Llama3.1-8B, Mistral-8B, and Gemma2-9B, at temperature 0, using GPT-4o as the evaluator.

### A.2 Subcategories in S1-Bench

Figure 7 shows the pie chart distribution of 28 subcategories in S1-Bench. For more details on the subcategories, please refer to Table 8,9.

### A.3 Prompt for S1-Bench construction

This section presents the prompts used in the construction of S1-Bench, including the Initial Generation prompt, the Discriminating Generation Quality prompt, and the Reduce Difficulty prompt. See Table 10 for details.

## B Baseline Models and Evaluation Details

### B.1 Baseline Model Details

Table 11 presents the abbreviations, IDs, and URLs of LLMs used in this paper. Table 12 displays the abbreviations, IDs, URLs, organizations, training algorithms, and training data volumes of open-source LRMs evaluated in this study.

### B.2 GPT-4o and Human Evaluation

We use GPT-4o as the evaluator to assess the correctness of the responses. If a final answer can be isolated, only the final answer is evaluated; otherwise, the entire response is assessed. The evaluation prompt is provided in Table 13.



Figure 7: S1-Bench Category Display. The inner circle represents four major categories, and the outer circle includes 28 subcategories.

To evaluate the consistency between the GPT-4 judge’s assessments and human judgments, we conduct a comprehensive human evaluation study involving three of the authors. Specifically, we randomly sample 20 question-answer pairs from each model’s greedy decoding results, resulting in a dataset of 640 pairs derived from 32 models (including 4 verifier LLMs and 28 LRMs). The questions, reference answers, and model responses are then presented to three annotators, who independently judge the correctness of each model response. The final human evaluation results are determined through majority voting. Ultimately, the Cohen’s Kappa between the human evaluators and the GPT-4 judge is calculated to be 0.83, indicating an exceptionally high level of agreement.

### B.3 Accuracy Metrics Details

**Pass@1:** Followed DeepSeek-R1 (Guo et al., 2025), we calculate pass@1 to assess the percentage of correct responses among the k=5 generations. Specifically, it is defined as:

$$\text{pass@1} = \frac{1}{k} \sum_{i=1}^k p_i, \quad (1)$$

where  $p_i$  is the correctness of the  $i$ -th generation.

**Acc@k:** Since S1-Bench is composed of extremely simple questions, we calculate acc@k. Specifically, acc@k=1 if all k responses are correct and acc@k = 0 otherwise. It is defined as:

$$\text{acc@k} = \prod_{i=1}^k p_i, \quad (2)$$

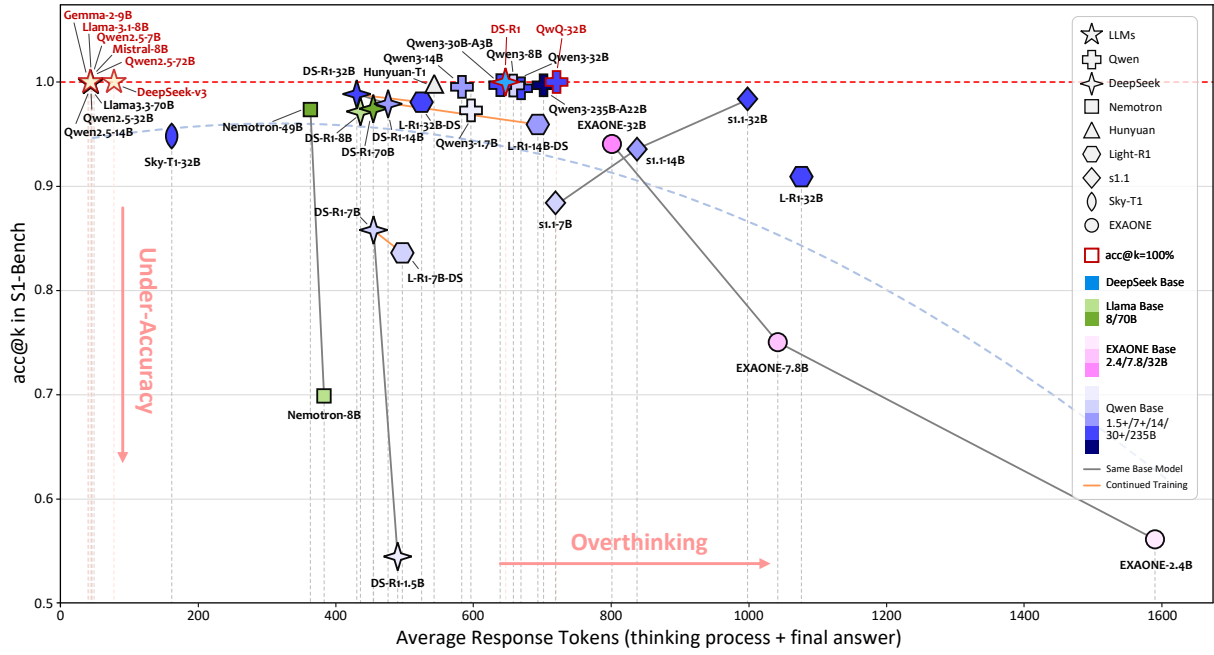


Figure 8: LRMs exhibit under-accuracy and overthinking on simple problems. Shapes represent organizations, colors represent base model families, with darker colors indicating larger models, and connecting lines represent the relationships between model families and training.

#### B.4 Types and Analysis of Format Errors

This section introduces a comprehensive taxonomy of format errors and highlights the importance of addressing these issues in future research. Unlike conventional LLMs, LRMs frequently exhibit format errors. These errors are defined by failing to use a unique end thinking marker (e.g., `</think>`) to separate the thinking process from the final answer. Format errors increase the difficulty of distinguishing the thinking process from the final answer and reveal the vulnerability of LRMs in following pre-defined formats.

To illustrate this phenomenon, we identify 12 distinct types of response formats produced by LRMs, each assigned a unique ID, as shown in Table 5. These 12 types are further grouped into three major categories:

- **Standard-Conforming Responses:** These responses meet the expected format by including exactly one end thinking marker (e.g., `</think>`) to delimit the thinking process from the final answer. Among these, type ID-100 includes a thinking process, while ID-101 omits it. The proportion of such responses is measured using the S-Corr metric.
- **Unreadable Responses:** These refer to generation failures, including cases where LRMs produce endlessly thinking content or solely

produce end thinking markers. The proportion of all other (i.e., readable) responses is measured using the L-Corr metric.

- **Readable but Malformed Responses:** These responses deviate from the standard format yet still contain extractable information. In some cases, the final answer is missing (e.g., ID-200, ID-202, ID-205), and we instead evaluate the correctness of the thinking process. In other cases, multiple (e.g., ID-201, ID-203) or unmatched<sup>9</sup> (e.g., ID-204, ID-206) end thinking markers are generated. In such instances, we treat the content following the last end thinking marker as the final answer for evaluation.

Table 14 and Table 15 present the distributions of 12 format types under top-p sampling and greedy sampling, respectively. we find: (1) The infinite generation phenomenon is widespread across most LRMs, particularly concentrated in LRMs with fewer than 32B parameters. (2) The Nemotron and EXAONE families frequently produce correctly formatted responses without any explicit thinking processes. This behavior can be viewed as a mechanism for mitigating over-thinking. However,

<sup>9</sup>This paper provides a reference collection of unmatched end thinking makers: `</ think>`, `</th think>`, `</ reason>`, `\nanswrn`, `**Final Answer**` and `**答案**`.

the EXAONE family still exhibits substantial overthinking tendencies, suggesting that LRMs’ capability to respond without visible reasoning and their tendency to overthink may be orthogonal characteristics. (3) None of the evaluated LRMs exhibited behaviors classified as ID-205/206.

Format	ID	marker (standard)	marker (unmatched)	marker (number)	thinking process	final answer
Standard	100	✓	–	1	✓	✓
	101	✓	–	1	×	✓
Readable but Malformed	200	✓	–	1	✓	×
	201	✓	–	>1	✓	✓
	202	✓	–	>1	✓	×
	203	✓	–	>1	×	✓
	204	×	✓	≥1	✓	✓
	205	×	✓	≥1	✓	×
	206	×	✓	≥1	×	✓
	207	×	×	0	–	✓
Unreadable	300	✓	✓	≥1	×	×
	301	×	×	0	–	∞

Table 5: Twelve types of response format.

## C More Experimental Setups & Results

### C.1 Greedy Sampling Results

Table 16 presents the performance of LRMs on S1-Bench under greedy sampling. While overall accuracy improves compared to top-p sampling, issues of inefficiency and accuracy degradation on simple questions remain.

### C.2 Efficiency Analysis across Subcategories.

Figure 9 illustrates the average response tokens across the 28 subcategories. In the heatmap, both models (rows) and subcategories (columns) are ordered in descending order according to their average number of response tokens.

### C.3 Solution Analysis Details

For solution analysis, We only use well-formatted thinking processes with correct final answers, as incorrect answers make it unclear whether LRMs are over-reasoning or under-reasoning, and malformed thinking processes cannot be precisely extracted. The segmentation process is performed by DeepSeek-v3, with prompts detailed in Table 17. We compute the average token count in the first solution round; if no solution is found, we use the token count of the entire thinking process.

### C.4 Thinking Redundancy Analysis

We conduct a similarity analysis to analyze how information redundancy in the thinking processes

changes as reasoning sequences increase. Specifically, we first divide the complete thinking process into  $k$  equal-length segments<sup>10</sup>. Then, we encode each segment using the all-MiniLM-L6-v2 model<sup>11</sup>. For each segment, we calculate the cosine similarity with all its preceding segments and use the maximum similarity as a measure of its information redundancy. As shown in Figure 10, information redundancy increases across all four main categories as reasoning sequences increase. Sky-T1-32B shows overall lower similarity, which stems from its shorter thinking process, but still demonstrates an upward trend.

### C.5 Error Analysis Details

In error analysis, we only use well-formatted samples, as malformed thinking processes cannot be precisely extracted. For samples with correct final answers, we categorize them based on whether the thinking process contains explicit incorrect conclusions in intermediate steps. For samples with incorrect final answers, we categorize them based on whether the correct answer is mentioned at least once during reasoning. We use DeepSeek-v3 for categorization, with prompts provided in Table 18.

### C.6 Gut Moment Analysis Details

We prompt GPT-4o to classify the initial part of model responses (before the first ‘\n\n’) into four types based on its comment on difficulty: easy, neutral, difficult, and no comment. The prompts for english question can be seen in Table 19. For Chinese queries, we use the translated version of the prompt in Chinese. In Table 6, we show the most common sentence of all LRMs in each type of “gut moment.”

Type	Sentence	Count
easy-zh	这个问题看起来挺简单的	308
easy-en	that seems straightforward	36
difficult-zh	这个问题看起来有点复杂	308
difficult-en	percentages can sometimes be tricky	7
neutral-zh	这个问题看起来好像不难	24
neutral-en	Hmm, interesting	3

Table 6: The most common sentence in each type of “gut moment.”

## D Error Cases

This section presents several error cases observed in LRMs. See Tables 20, 21, 22, 23, 24, and 25.

<sup>10</sup>We set  $k=15$ , changing its value does not affect the conclusions.

<sup>11</sup><https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

Paper Abbreviation	MATH	GSM8K	AIME	GPQA	AMC	MMLU	Olympiad-Bench	SVAMP	LiveCode-Bench	Common-SenseQA
Codi (Shen et al., 2025b)		✓						✓		
CISC (Taubenfeld et al., 2025)	✓	✓				✓				
CoT-Valve (Ma et al., 2025b)		✓	✓							
Dast (Shen et al., 2025a)	✓		✓							
ATM (Chen et al., 2024a)		✓						✓		✓
DEER (Yang et al., 2025b)	✓		✓	✓	✓					
DPTS (Ding et al., 2025)	✓	✓								
Dynasor (Fu et al., 2024)		✓							✓	
ESC (Li et al., 2024)	✓	✓						✓		✓
Hawkeye (She et al., 2025)	✓	✓	✓							
token complexity (Lee et al., 2025)	✓	✓				✓				
INFTYTHINK (Yan et al., 2025)	✓		✓	✓						
KIMI K1.5 (Team et al., 2025a)	✓		✓			✓			✓	
L1 (Aggarwal and Welleck, 2025)	✓		✓	✓		✓	✓			
LightThinker (Zhang et al., 2025)		✓		✓		✓				
LS-Mixture SFT (Yu et al., 2025a)	✓		✓	✓						
DSC (Wang et al., 2024)	✓	✓								
O1-Pruner (Luo et al., 2025)	✓	✓								
MRT (Qu et al., 2025b)	✓		✓	✓						
Self-Doubt (Fu et al., 2025)	✓		✓		✓					
RASC (Wan et al., 2024)		✓						✓		✓
NoThinking (Ma et al., 2025a)			✓		✓		✓		✓	
Retro-Search (Lu et al., 2025)	✓	✓	✓		✓		✓			
RSD (Liao et al., 2025)	✓	✓		✓		✓	✓			
ST-BoN (Wang et al., 2025b)	✓			✓						
Elastic Reasoning (Xu et al., 2025b)	✓		✓		✓		✓		✓	
FS-BoN (Munkhbat et al., 2025)	✓	✓								
SoT (Aytes et al., 2025)		✓						✓		✓
SpecReason (Pan et al., 2025)	✓		✓	✓						
Speculative Thinking (Yang et al., 2025d)	✓		✓	✓	✓					
SPIRIT (Cui et al., 2025)		✓								
ITC Analysis (Wang et al., 2025a)	✓		✓	✓						
Think when needed (Yang et al., 2025c)			✓							
THINKPRUNE (Hou et al., 2025)	✓		✓		✓		✓			
TALE (Han et al., 2024)		✓								
TokenSkip (Xia et al., 2025)	✓	✓								
TOPS (Yang et al., 2025e)	✓	✓	✓							
efficient reasoning (Arora and Zanette, 2025)	✓	✓	✓							
TWT (Xu et al., 2025a)										✓
Z1 (Yu et al., 2025b)	✓			✓					✓	
Count	28	24	20	11	8	6	6	5	5	5

Table 7: A total of 40 studies on LRM efficiency before May 2025 were included. Benchmarks that appeared more than four times are listed.

cate.	subcatrgories	Explanation and cases
reasoning question	numerical reasoning	Questions that require performing basic mathematical operations or solving simple algebraic equations to arrive at a numerical answer. <b>Case:</b> What's two plus three?
	code reasoning	Questions that require tracing through and executing simple code snippets to determine their output or behavior when run in a specific programming environment. <b>Case:</b> What is the output of the following code when run in Python 3 environment: word = "hello"\nprint(len(word))
	set reasoning	Questions that require applying simple syllogistic reasoning to determine whether elements belong to sets based on clearly stated relationships. <b>Case:</b> All squares are quadrilaterals. A shape is a square, is it a quadrilateral?
	temporal reasoning	Questions that require calculating time durations, ages, or future dates by applying simple arithmetic operations to temporal information. <b>Case:</b> How many minutes equal 120 seconds?
	spatial reasoning	Questions that require determining relative positions, directions, or orientations of objects in space based on simple spatial relationships. <b>Case:</b> If a bird is flying above a tree, where is the tree in relation to the bird?
	causal reasoning	Questions that require determining outcomes by applying simple cause-and-effect relationships based on given conditional statements. <b>Case:</b> If ferromagnetic material is placed in a magnetic field, it will become magnetized. An iron nail was placed next to a strong magnet for some time. Has the nail been magnetized?
	natural law reasoning	Questions that require applying basic knowledge of physical laws and natural phenomena to predict simple observable outcomes in everyday scenarios. <b>Case:</b> Which is faster, an airplane or the propagation of light?
knowledge question	geometry facts	Questions that require recalling simple and fundamental geometric properties about shapes, angles, and basic geometric figures. <b>Case:</b> How many angles does a trapezoid have?
	geographic facts	Questions that require recalling simple factual information about locations, landmarks, political divisions, celestial bodies, and other basic geographic knowledge. <b>Case:</b> Which is the largest continent on Earth?
	historical facts	Questions that require recalling basic facts about historical events. <b>Case:</b> Which country first invented paper?
	biographical facts	Questions that require recalling basic facts about the identities, achievements, and characteristics of historical figures. <b>Case:</b> Who proposed the theory of universal gravitation?
	measurement units	Questions that require recalling simple conversion relationships between standard units of measurement. <b>Case:</b> How many centimeters equal 1 meter?
	scientific notation	Questions that require recalling basic scientific symbols, formulas, and standard units used in scientific communication. <b>Case:</b> What is the chemical symbol for oxygen?
	creative authorship	Questions that require recalling the creators or originators of notable artistic, literary, musical, and cultural works. <b>Case:</b> Who is the author of Hamlet?

Table 8: The subcategory descriptions and cases of reasoning questions and knowledge questions.

cate.	subcatrgories	Explanation and cases
instruction following	repetition constraints	Questions that require outputting specified characters, words, or phrases a specific number of times according to simple formatting instructions. <b>Case:</b> Output the number "7" four times, without using separators.
	length constraints	Questions that require generating outputs of a specific length or with a specific number of components based on simple counting constraints. <b>Case:</b> Output a four-digit number.
	character constraints	Questions that require generating words or numbers that conform to simple specified character patterns or formatting rules. <b>Case:</b> Output a number that begins with 8.
	counting constraints	Questions that require counting specific characters or elements within a given text or sequence. <b>Case:</b> Output the number of letter "y" in the word "yes".
	transformation constraints	Questions that require modifying text or numbers according to simple formatting or character substitution rules to produce a transformed output. <b>Case:</b> Output the word "good" with all letters capitalized directly.
	sentence constraints	Questions that require generating sentences that conform to simple specified content or structural requirements. <b>Case:</b> Give a sentence that contains the phrase "have lunch" directly.
analysis question	sentiment classification	Questions that require determining whether simple statements express positive or negative emotions based on the tone and word choice. <b>Case:</b> Does the sentence "I hate rainy days." express a positive or negative emotion?
	named entity recognition	Questions that require identifying the correct category of named entities (such as people, places, organizations, or time expressions) within simple sentences. <b>Case:</b> In the sentence "Napoleon died in 1821", is "1821" a time or a place name?
	language classification	Questions that require identifying the language of origin for simple words or phrases based on their characteristic writing systems or common vocabulary. <b>Case:</b> Is the word "hello" English or Japanese?
	topic classification	Questions that require identifying the primary subject matter or thematic category of simple sentences based on their content and context clues. <b>Case:</b> Is the topic of the sentence "The stock market rose 2% today" finance or technology?
	intent recognition	Questions that require determining the communicative purpose behind simple utterances or statements based on their phrasing and context. <b>Case:</b> Is the intention of the sentence "I'm sorry I'm late." to apologize or to blame?
	syntax classification	Questions that require identifying the correct grammatical structure or sentence type of simple expressions based on their form, punctuation, and communicative function. <b>Case:</b> Is "Close the door!" an imperative sentence or an interrogative sentence?
	grammar classification	Questions that require identifying simple grammatical properties (like tense, voice, or polarity) of sentences based on their structure and verb forms. <b>Case:</b> Is "The apple was eaten." in active voice or passive voice?
	coreference resolution	Questions that require identifying which entity a pronoun or reference term refers to in simple sentences by tracking relationships between words in the text. <b>Case:</b> In "My computer is broken, and I need to fix it." What does "it" refer to?

Table 9: The subcategory descriptions and cases of instruction following questions and analysis questions.

---

**Prompt for construction workflow for S1-Bench**

---

**Data Generation Prompt**

---

Generate 50 pairs of questions and answers in both Chinese and English based on the category's name, definition, and specific simplicity criteria. The following conditions must be satisfied:

1. Questions must be naturally and clearly expressed, unambiguous, and free of intentional traps.
2. Answers must be unique or easily falsifiable, with no possibility of multiple correct answers.
3. Make the questions as diverse as possible.

# Category Name and Definition:  
{name\_and\_definition}

# Specific Simplicity Criteria:  
{criteria}

# Cases:

## English question:  
{question\_en}

## English Answer:  
{answer\_en}

## Chinese question:  
{question\_zh}

## Chinese Answer:  
{answer\_zh}

Please generate 50 pairs of Chinese and English questions and answers in the following format:

[question]English-question[answer]English-answer[question]Chinese-question[answer]Chinese-answer...

Start generating:

---

**Quality Discrimination Prompt**

---

Given a question, its answer, and its category, please analyze from the following perspectives as comprehensively as possible:

1. Whether the question belongs to the specified category and meet the Specific Simplicity Criteria.
2. Whether the question is easy, clear, unambiguous, and has an absolutely unique answer.
3. Whether the answer is absolutely correct; if not, what the correct answer should be.
4. Whether the question is similar to other given questions, and if similar, whether more diverse questions can be generated.

# Category Name and Definition:  
{name\_and\_definition}

# Specific Simplicity Criteria:  
{criteria}

# Question and Answer:  
{question\_with\_answer}

# Other Questions:  
{questions\_list}

Begin your analysis, aiming to be as detailed and comprehensive as possible:

---

**Difficulty Reduction Prompt**

---

Given a question and answer that are too complex for the model to answer correctly, you need to further reduce their difficulty while trying to:

- Ensure the question aligns with the Category Name and Definition.
- Ensure the question meets the Specific Simplicity Criteria.

# Category Name and Definition:  
{name\_and\_definition}

# Specific Simplicity Criteria:  
{criteria}

# Question and Answer:  
{question\_with\_answer}

The new question and answer:

---

Table 10: "Category Name and Definition" refers to the subcategory name and its definition, while Specific Simplicity Criteria refers to the simplicity requirements specific to the main category.

Model	Model ID	URL
Qwen2.5-7B	Qwen2.5-7B-Instruct	<a href="https://huggingface.co/Qwen/Qwen2.5-7B-Instruct">https://huggingface.co/Qwen/Qwen2.5-7B-Instruct</a>
Llama3.1-8B	Llama-3.1-8B-Instruct	<a href="https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct">https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct</a>
Mistral-8B	Ministral-8B-Instruct-2410	<a href="https://huggingface.co/mistralai/Ministral-8B-Instruct-2410">https://huggingface.co/mistralai/Ministral-8B-Instruct-2410</a>
Gemma2-9B	gemma-2-9b-it	<a href="https://huggingface.co/google/gemma-2-9b-it">https://huggingface.co/google/gemma-2-9b-it</a>
Qwen2.5-14B	Qwen2.5-14B-Instruct	<a href="https://huggingface.co/Qwen/Qwen2.5-14B-Instruct">https://huggingface.co/Qwen/Qwen2.5-14B-Instruct</a>
Qwen2.5-32B	Qwen2.5-32B-Instruct	<a href="https://huggingface.co/Qwen/Qwen2.5-32B-Instruct">https://huggingface.co/Qwen/Qwen2.5-32B-Instruct</a>
Qwen2.5-72B	Qwen2.5-72B-Instruct	<a href="https://huggingface.co/Qwen/Qwen2.5-72B-Instruct">https://huggingface.co/Qwen/Qwen2.5-72B-Instruct</a>
Llama3.3-70B	Llama-3.3-70B-Instruct	<a href="https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct">https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct</a>
DeepSeek-v3	DeepSeek-V3-0324	<a href="https://huggingface.co/deepseek-ai/DeepSeek-V3-0324">https://huggingface.co/deepseek-ai/DeepSeek-V3-0324</a>

Table 11: Mapping of LLM abbreviations and IDs used in this paper, with their open-source URLs.

Model ID	Abbreviation	Base Model	Alg.	Size
<i>DeepSeek</i>				
DeepSeek-R1-Distill-Qwen-1.5B	DS-R1-1.5B	Qwen2.5-Math-1.5B	SFT	800K
DeepSeek-R1-Distill-Qwen-7B	DS-R1-7B	Qwen2.5-Math-7B	SFT	800K
DeepSeek-R1-Distill-Llama-8B	DS-R1-8B	Llama-3.1-8B	SFT	800K
DeepSeek-R1-Distill-Qwen-14B	DS-R1-14B	Qwen2.5-14B	SFT	800K
DeepSeek-R1-Distill-Qwen-32B	DS-R1-32B	Qwen2.5-32B	SFT	800K
DeepSeek-R1-Distill-Llama-70B	DS-R1-70B	Llama-3.3-70B-Instruct	SFT	800K
DeepSeek-R1	DS-R1	DeepSeek-V3-0324	SFT&RL	800K&–
<i>Qwen</i>				
QwQ-32B	QwQ-32B	Qwen2.5-32B	–	–
Qwen3-235B-A22B	Qwen3-A22B	Qwen3-235B-A22B-Base	SFT&RL	–&–
Qwen3-30B-A3B	Qwen3-A3B	Qwen3-30B-A3B-Base	SFT&RL	–&–
Qwen3-32B	Qwen3-32B	Qwen3-32B-Base	SFT&RL	–&–
Qwen3-14B	Qwen3-14B	Qwen3-14B-Base	SFT&RL	–&–
Qwen3-8B	Qwen3-8B	Qwen3-8B-Base	SFT&RL	–&–
Qwen3-1.7B	Qwen3-1.7B	Qwen3-1.7B-Base	SFT&RL	–&–
<i>qihoo360</i>				
Light-R1-7B-DS	L-R1-7B-DS	DeepSeek-R1-Distill-Qwen-7B	SFT	3K
Light-R1-14B-DS	L-R1-14B-DS	DeepSeek-R1-Distill-Qwen-14B	SFT&RL	3K&–
Light-R1-32B-DS	L-R1-32B-DS	DeepSeek-R1-Distill-Qwen-32B	SFT	3K
Light-R1-32B	L-R1-32B	Qwen2.5-32B-Instruct	SFT&DPO	73K&–
<i>simplescaling</i>				
s1.1-7B	s1.1-7B	Qwen2.5-7B-Instruct	SFT	1K
s1.1-14B	s1.1-14B	Qwen2.5-14B-Instruct	SFT	1K
s1.1-32B	s1.1-32B	Qwen2.5-32B-Instruct	SFT	1K
<i>LG AI Research</i>				
EXAONE-Deep-2.4B	EXAONE-2.4B	EXAONE-3.5-2.4B-Instruct	SFT&DPO&RL	1.6M&20K&10K
EXAONE-Deep-7.8B	EXAONE-7.8B	EXAONE-3.5-7.8B-Instruct	SFT&DPO&RL	1.6M&20K&10K
EXAONE-Deep-32B	EXAONE-32B	EXAONE-3.5-32B-Instruct	SFT&DPO&RL	1.6M&20K&10K
<i>NVIDIA</i>				
Llama-3.1-Nemotron-Nano-8B-v1	Nemotron-8B	Llama-3.1-8B-Instruct	SFT&RL	–&–
Llama-3.3-Nemotron-Super-49B-v1	Nemotron-49B	Llama-3.3-70B-Instruct	SFT&RL	–&–
<i>NovaSky</i>				
Sky-T1-32B-Flash	Sky-T1-32B	Qwen2.5-32B-Instruct	SFT&SimPO	17K&10K

Table 12: The open-source LRMs details evaluated for S1-Bench.

---

**Prompt for Correctness Evaluation**

---

***Evaluation on SI-Bench***

---

**\*\*Question:\*\***

{question}

**\*\*Model Answer:\*\***

{model\_answer}

**\*\*Ground Truth:\*\***

{ground\_truth}

Your task is to evaluate whether the model’s answer is correct.

An answer is considered correct as long as it contains the ground truth (regardless of how complex or detailed the description is).

If there are parenthetical notes after the ground truth, then there may be multiple correct answers. In this case, the given answer is just one example, and any answer that meets the requirements specified in the notes can be considered correct.

Additionally, some reasonably uncertain supplementary information is also considered appropriate, including more details, possibilities, and expanded discussion. You should focus more on whether the reply contains the correct answer.

You need to output a standard JSON, providing your explanation of the evaluation in the "explain" field, and giving the evaluation result in the "result" field, where 1 means the answer is correct and 0 means it is incorrect.

Your action should follow the given format: "explain": "", "result": 0/1

---

Table 13: Prompt for Correctness Evaluation.

Model	Standard		Readable but Malformed								Unreadable	
	100	101	200	201	202	203	204	205	206	207	300	301
Qwen3-A22B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-A3B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
QwQ-32B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-32B	99.91	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.09
Qwen3-14B	99.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05
Qwen3-8B	99.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05
Qwen3-1.7B	99.81	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.19
Hunyuan-T1	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-70B	99.91	0.00	0.09	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-32B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-14B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-8B	99.53	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.24	0.00	0.24
DS-R1-7B	99.24	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.76
DS-R1-1.5B	97.58	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.42
Sky-T1-32B	95.26	0.00	0.62	0.09	0.19	0.00	0.28	0.00	0.00	3.03	0.00	0.52
Nemotron-49B	66.07	33.93	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Nemotron-8B	58.06	26.26	0.00	0.00	0.00	0.09	0.00	0.00	0.00	15.02	0.00	0.57
L-R1-32B	95.07	0.00	0.00	0.00	0.00	0.00	0.81	0.00	0.00	3.03	0.00	1.09
L-R1-32B-DS	99.81	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.19
L-R1-14B-DS	99.19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.81
L-R1-7B-DS	99.67	0.00	0.05	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.24
s1.1-32B	99.53	0.00	0.00	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.43
s1.1-14B	97.39	0.00	0.00	0.14	0.00	0.00	0.24	0.00	0.00	0.00	0.00	2.23
s1.1-7B	88.96	0.00	0.00	7.96	0.09	0.00	0.00	0.00	0.00	0.09	0.00	2.89
EXAONE-32B	67.39	32.42	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.19
EXAONE-7.8B	65.83	32.23	0.00	0.00	0.05	0.47	0.00	0.00	0.00	0.14	0.00	1.28
EXAONE-2.4B	81.42	15.83	0.00	0.09	0.00	0.05	0.00	0.00	0.00	0.05	0.00	2.56

Table 14: Format type rates under top-p sampling.

Model	Standard		Readable but Malformed								Unreadable	
	100	101	200	201	202	203	204	205	206	207	300	301
Qwen3-A22B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-A3B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
QwQ-32B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-32B	99.76	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.24
Qwen3-14B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-8B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Qwen3-1.7B	99.76	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.24
Hunyuan-T1	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-70B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-32B	100.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DS-R1-14B	99.76	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.24
DS-R1-8B	99.53	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.24	0.00	0.24
DS-R1-7B	97.87	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.13
DS-R1-1.5B	91.94	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	8.06
Sky-T1-32B	99.29	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.47	0.00	0.24
Nemotron-49B	60.90	39.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Nemotron-8B	55.21	26.78	0.00	0.00	0.00	0.00	0.00	0.00	0.00	16.35	0.00	1.66
L-R1-32B	85.55	0.24	0.00	0.24	0.71	0.24	0.95	0.00	0.00	6.64	2.61	2.84
L-R1-32B-DS	99.29	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.71
L-R1-14B-DS	98.82	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.18
L-R1-7B-DS	98.82	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.18
s1.1-32B	98.82	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.18
s1.1-14B	95.97	0.00	0.00	0.24	0.00	0.00	0.24	0.00	0.00	0.00	0.00	3.55
s1.1-7B	87.91	0.00	0.00	6.64	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.45
EXAONE-32B	65.88	33.89	0.00	0.24	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EXAONE-7.8B	63.51	33.65	0.00	0.00	0.00	0.24	0.00	0.00	0.00	0.24	0.00	2.37
EXAONE-2.4B	78.91	15.88	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.21

Table 15: Format type rates under greedy decoding setting.

Model	Size	acc (Loose)	acc (Strict)	L-Corr ↑	S-Corr ↑	Tokens ↓
Qwen3-235B-A22B	235B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	702.70
Qwen3-30B-A3B	30B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	636.35
QwQ-32B	32B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	750.41
Qwen3-32B	32B	99.76	99.76	99.76	99.76	673.62
Qwen3-14B	14B	99.76	99.76	<b>100.00</b>	<b>100.00</b>	597.06
Qwen3-8B	8B	99.76	99.76	<b>100.00</b>	<b>100.00</b>	649.45
Qwen3-1.7B	1.7B	99.53	99.53	99.76	99.76	579.01
Hunyuan-T1	–	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	541.09
DS-R1	671B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	621.89
DS-R1-70B	70B	99.76	99.76	<b>100.00</b>	<b>100.00</b>	469.78
DS-R1-32B	32B	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	428.46
DS-R1-14B	14B	99.29	99.29	99.76	99.76	463.52
DS-R1-8B	8B	97.63	97.39	99.76	99.53	452.11
DS-R1-7B	7B	94.31	94.31	97.87	97.87	436.87
DS-R1-1.5B	1.5B	76.54	76.54	<b>91.94</b>	91.94	473.67
Sky-T1-32B	32B	99.53	99.05	99.76	99.29	<b>157.12</b>
Nemotron-49B	49B	99.53	99.53	<b>100.00</b>	<b>100.00</b>	<b>337.94</b>
Nemotron-8B	8B	84.60	77.73	98.34	<b>81.99</b>	446.62
L-R1-32B	32B	92.18	85.78	<b>94.55</b>	<b>85.78</b>	<b>996.36</b>
L-R1-32B-DS	32B	99.29	99.29	99.29	99.29	528.45
L-R1-14B-DS	14B	98.82	98.82	98.82	98.82	664.28
L-R1-7B-DS	7B	92.65	92.65	98.82	98.82	514.60
s1.1-32B	32B	98.82	98.82	98.82	98.82	983.38
s1.1-14B	14B	95.97	95.50	96.45	95.97	786.30
s1.1-7B	7B	94.31	87.68	<b>94.55</b>	87.91	630.52
EXAONE-32B	32B	97.63	97.39	<b>100.00</b>	99.76	746.89
EXAONE-7.8B	7.8B	86.73	86.49	97.63	97.16	947.92
EXAONE-2.4B	2.4B	<b>72.99</b>	<b>72.99</b>	94.79	94.79	<b>1394.72</b>

Table 16: Main results in the greedy decoding setting on the S1-Bench, sorted by model family. **Bold teal** marks best performance, **teal** second best, **bold burgundy** worst, and **burgundy** second worst.

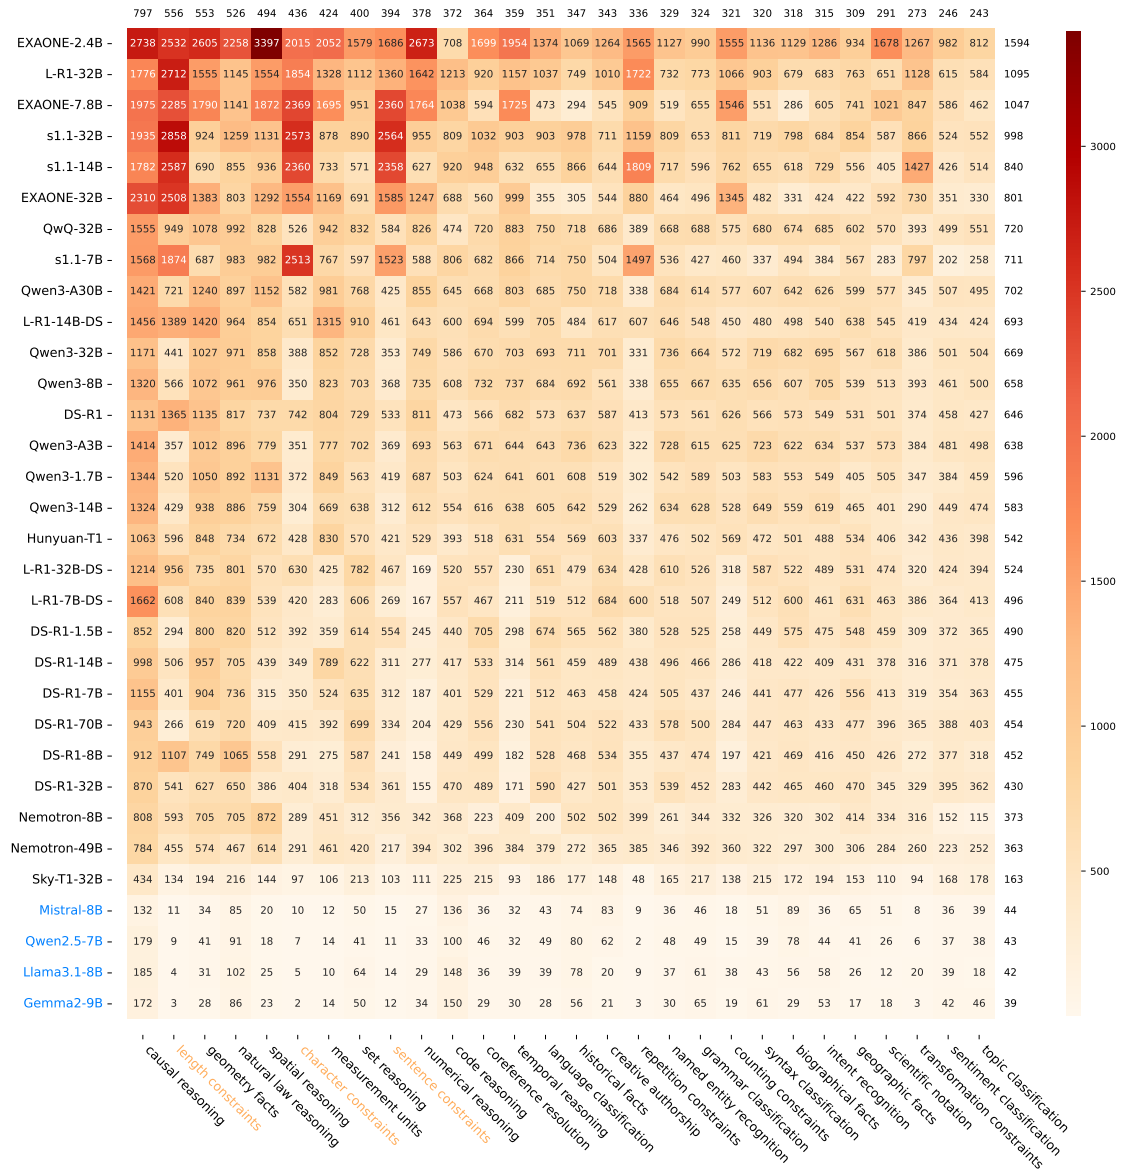


Figure 9: Average response token counts on the 28 subcategories, which is the average result of five generations under top-p sampling.

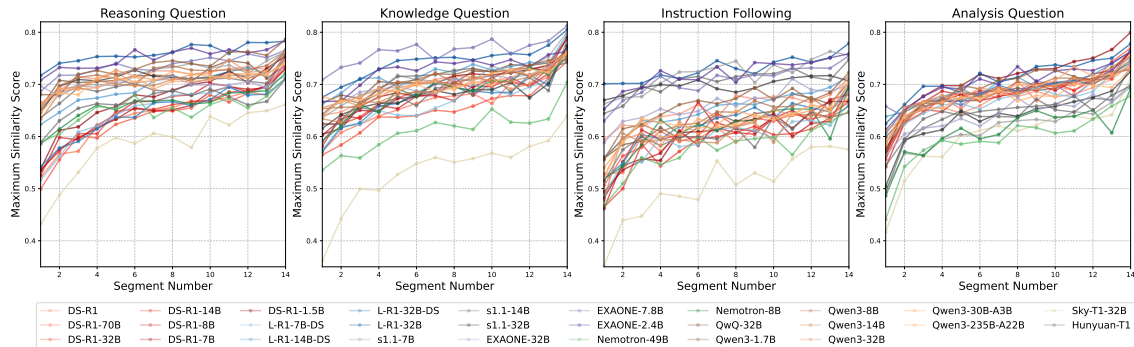


Figure 10: Maximum similarity between each segment and all preceding segments for LRMs across four categories.

---

## Prompts for Solution Segmentation

---

### # Task Description:

Your task is to segment the given Chain of Thought according to the following rules:

#### 1. Segmentation positions:

(1) Please identify and extract all sub-reasoning processes from the Chain of Thought that meet the following condition: **They explicitly arrive at a conclusion (including cases phrased as questions, e.g., "right?") that is directly consistent with the Ground Truth.** Reasoning processes that only indirectly support the Ground Truth or result in partially aligned conclusions should be excluded.

(2) After clearly reaching the conclusion consistent with the Ground Truth, **insert the segmentation marker <split> to distinctly separate each qualifying sub-reasoning process.**

#### 2. Output Restriction:

(1) You should only directly output the segmentation result without adding any additional supplements.

(2) Except for inserting the <split> separator, you must not make any other modifications to the original Chain of Thought, not even minor character-level changes such as punctuation, spacing, or capitalization. In other words, after removing all <split> separators, the output text must match the original Chain of Thought exactly.

### # Some cases:

#### ## Case 1:

##### ### Question:

What is the value of x if x plus 3 equals 8?

##### ### Ground Truth:

5

##### ### Chain of Thought:

Okay, so I need to figure out what x is if x plus 3 equals 8. The problem says "x plus 3 equals 8," so in mathematical terms, that would be:  $x + 3 = 8$

Alright, now I need to solve for x. From what I remember in algebra, the goal is to isolate x on one side of the equation. That means I need to get rid of that "+ 3" that's being added to x. To do that, I should perform the opposite operation. Since it's adding 3, the opposite would be subtracting 3.

So, if I subtract 3 from both sides, that should keep things equal. Let me try that:  $x + 3 - 3 = 8 - 3$

Simplifying both sides, on the left side, the +3 and -3 cancel each other out, leaving just x. And on the right side, 8 minus 3 is 5. So that simplifies to:  $x = 5$

Wait, let me check if that makes sense. If I substitute x back into the original equation, does it hold true? Let's see:

Original equation:  $x + 3 = 8$

Substitute  $x = 5$ :  $5 + 3 = 8$

5 plus 3 is indeed 8, so that works out.

Another way might be to consider what number, when you add 3 to it, gives 8. So starting from 8, if I take away 3, I should get back to the original number. That's the same as subtracting 3 from 8, which again gives 5. So that's consistent with the first method.

Wait a second, what if the original problem was written differently? The problem says "x plus 3 equals 8." In symbols, that's unambiguous:  $x + 3 = 8$ . So no, I think that's right.

\*\*Final Answer\*\*

The value of x is 5.

##### ### Output:

Okay, so I need to figure out what x is if x plus 3 equals 8. The problem says "x plus 3 equals 8," so in mathematical terms, that would be:  $x + 3 = 8$

Alright, now I need to solve for x. From what I remember in algebra, the goal is to isolate x on one side of the equation. That means I need to get rid of that "+ 3" that's being added to x. To do that, I should perform the opposite operation. Since it's adding 3, the opposite would be subtracting 3.

So, if I subtract 3 from both sides, that should keep things equal. Let me try that:  $x + 3 - 3 = 8 - 3$

Simplifying both sides, on the left side, the +3 and -3 cancel each other out, leaving just x. And on the right side, 8 minus 3 is 5. So that simplifies to:  $x = 5$ <split>

Wait, let me check if that makes sense. If I substitute x back into the original equation, does it hold true? Let's see:

Original equation:  $x + 3 = 8$

Substitute  $x = 5$ :  $5 + 3 = 8$

5 plus 3 is indeed 8, so that works out.<split>

Another way might be to consider what number, when you add 3 to it, gives 8. So starting from 8, if I take away 3, I should get back to the original number. That's the same as subtracting 3 from 8, which again gives 5.<split> So that's consistent with the first method.

Wait a second, what if the original problem was written differently? The problem says "x plus 3 equals 8." In symbols, that's unambiguous:  $x + 3 = 8$ . So no, I think that's right.

\*\*Final Answer\*\*

The value of x is 5<split>

...(Other examples are omitted.)

### # Eval Target:

#### ## Question:

{question}

#### ## Ground Truth:

{ground\_truth}

#### ## Chain of Thought:

{thinking\_process}

#### ## Output:

---

Table 17: Prompts for Solution Segmentation.

---

**Prompts for Error Analysis**

---

***Prompts for samples whose final answer is correct***

---

## # Task Description:

You will receive a Question, its corresponding Ground Truth, and a Chain of Thought(CoT) generated by a LLM for that Question. Your task is to carefully analyze the CoT and assign it to one of the two predefined categories listed below.

## # Categories:

- 1: The CoT \*\*\*includes explicit incorrect conclusions\*\*\* in intermediate reasoning steps.
- 2: The CoT \*\*\*doesn't include any explicit incorrect conclusion\*\*\* in intermediate reasoning steps.

Output your evaluation in the following format:

## # TheReason:

[note: Conduct a step-by-step analysis, stating if and where explicit incorrect conclusions occur in the CoT.]

## # ErrorType:

[note: Summarize each incorrect conclusion into a specific error type using a phrase of less than 5 words, such as factual inaccuracies, logical fallacies, comprehension mistakes, calculation errors, formatting issues, and so forth, to better conduct further evaluation and analysis. Directly output a Python list, where each element represents the error type of a specific incorrect conclusion in the CoT. If there are no incorrect conclusions, return an empty list.]

## # TheCategory:

[note: Provide your classification based on your analysis using only the number "1" or "2". Do not add any additional text.]

## # Question:

{question}

## # Ground Truth:

{ground\_truth}

## # CoT:

{thinking\_process}

## # TheReason:

---

***Prompts for samples whose final answer is incorrect***

---

## # Task Description:

You will receive a Question, its corresponding Ground Truth, and a Chain of Thought(CoT) generated by a LLM for that Question. Your task is to carefully analyze the CoT and assign it to one of the two predefined categories listed below.

## # Categories:

- 1: Regardless of whether the CoT ultimately arrives at the correct final answer or not, \*\*\*the correct answer is explicitly mentioned at least once\*\*\* within the reasoning steps (even if it is not ultimately adopted).
- 2: \*\*\*The correct answer is never explicitly mentioned or referenced\*\*\* at any point within the reasoning steps.

Output your evaluation in the following format:

## # TheReason:

[note: Conduct a step-by-step analysis, explicitly stating whether and where a correct answer is mentioned within the reasoning steps.]

## # TheCategory:

[note: Provide your classification based on your analysis using only the number "1" or "2". Do not add any additional text.]

## # Question:

{question}

## # Ground Truth:

{answer}

## # CoT:

{thinking\_part}

## # TheReason:

---

Table 18: Prompts for Error Analysis.

---

**Prompts for classify the “gut moment” for English questions**

---

**# Overall Task Description**

You will be given the beginning portion of a response written by a large language model when answering a question. Your task is to classify the response into one of the following four categories based on the initial comment about **\*\*the difficulty of the question\*\***.

Important: Only consider the **\*\*initial comment\*\*** on difficulty made in the response. If the model later changes its assessment, please ignore those later revisions—focus solely on the first difficulty-related comment.

You must assign the response to exactly one of the four categories below:

Category 1: Initial comment indicates the question is easy. Initial comment includes phrases like “simple,” “basic,” “straightforward,” “common,” etc., clearly stating the question is easy.

Category 2: Initial comment indicates the question is difficult. Initial comment includes phrases like “a bit complex,” “somewhat difficult,” “challenging,” “tricky,” etc., clearly stating the question is hard.

Category 3: Initial comment is neutral or suggests moderate difficulty. Initial comment includes phrases like “interesting,” “not hard,” “not complex,” “fun,” etc., without a clear indication of difficulty level.

Category 4: No comment on difficulty. The response contains no mention of difficulty at all.

Please output your result in the following format:

[Explanation of classification]<sep>[Category (just the number)]<sep>[Short phrase containing the comment]

If it falls into Category 4 (no comment), write "None" for the [Short phrase containing the comment] part.

# Examples for each category will be shown below.

**## Category 1 Examples:**

Question:

What is two plus three?

Response:

Hmm, the user is asking what two plus three equals. This question seems very simple, though it might be a bit confusing for someone just starting to learn math. First, I need to figure out what the user’s intent is. They might be testing my basic calculation ability or genuinely need help solving the problem.

Output:

The initial comment describes the question as "very simple"<sep>1<sep>This question seems very simple

...(Other examples are omitted.)

**## Category 2 Examples:**

Question:

What’s the next floor after basement level two?

Response:

Hmm, this question seems a bit complex, but I’ll try to understand it first. The question is “What’s the next floor after basement level two?” First, I need to clarify what “basement level two” means. Usually, the first basement is the deepest, and then going up you have ground floor, first floor, and so on.

Output:

The initial comment says it’s “a bit complex”<sep>2<sep>This question seems a bit complex

...(Other examples are omitted.)

**## Category 3 Examples:**

Question:

“Doctors recommend drinking at least eight glasses of water a day.” Is this sentence about health or education?

Response:

Hmm, this question seems not complicated, though if you think carefully it might be a bit tricky. First, I need to understand the meaning of the sentence “Doctors recommend drinking at least eight glasses of water a day.”

Output:

The initial comment says it’s “not complicated,” even though it later mentions it might be tricky. Classification should follow the initial comment<sep>3<sep>This question doesn’t seem complicated

...(Other examples are omitted.)

**## Category 4 Examples:**

Question:

All birds have feathers. A sparrow is a bird. Does a sparrow have feathers?

Response:

Hmm, the user is asking a question about birds and feathers—specifically, that all birds have feathers, a sparrow is a bird, so does the sparrow have feathers? This looks like a logic reasoning question. The user might be learning the basics of logic or biology.

Output:

Although the response mentions it’s a “logic reasoning question,” it contains no evaluation of difficulty<sep>4<sep>None

...(Other examples are omitted.)

# Here is the one you need to classify.

Question:

question

Response:

response

Output:

---

Table 19: Prompts for classifying the “gut moment” in English questions.

---

**Categories:** instruction following   **Model:** DeepSeek-R1-Distill-Qwen-7B

---

**Question:**

Output the number "7" four times, without using separators.

**Ground Truth:**

7777

**Thinking process:**

To solve the problem of outputting the number "7" four times without using separators, I need to find a way to repeat the digit 7 multiple times without relying on concatenation or other connecting characters.

One approach is to use mathematical operations that inherently produce the number 7. Addition is a simple and effective operation that can be used to achieve this. By adding 7 to itself multiple times, I can create the desired sequence of numbers.

Specifically, adding 7 four times will give me  $7 + 7 + 7 + 7$ , which equals 28. This method ensures that the number 7 is used exactly four times without the need for any separators.

Therefore, the solution is to use the addition operation to repeat the number 7 four times and output the result.

**Final Answer:**

To output the number "7" four times without using separators, you can utilize mathematical operations that inherently produce the number 7 multiple times. One straightforward way is to use addition:

$[7 + 7 + 7 + 7 = 28]$

This expression shows that the number "7" is used four times in a row without any separators, resulting in the total of 28.

$[7 + 7 + 7 + 7]$

---

Table 20: Error Case for LRM.

**Question:**

输出一个动物的英文单词，并且只有三个字母

Output an English word for an animal, and it must be only three letters long.

**Ground Truth:**

任何一个由三个字母组成的动物名称都可以（例如：cat）

Any animal name consisting of three letters is acceptable (e.g., cat).

**Thinking process:**

嗯，用户让我输出一个动物的英文单词，而且必须是三个字母的。首先，我得确定有哪些常见的动物是三个字母的。可能想到的有猫（Cat）、狗（Dog），不过这两个都是三个字母吗？Cat是C-A-T，确实是三个字母，Dog也是D-O-G，三个字母。那这两个应该都可以。不过用户可能需要更多的例子，或者只需要一个？题目里说的是“输出一个”，所以可能只需要一个。

不过等等，可能还有其他选项，比如Ant（蚂蚁）、Ape（猿）、Bat（蝙蝠）、Bear（熊）是四个字母，所以不行。Fish是四个字母，Fox是三个字母，F-O-X，对，Fox也是三个字母。还有Lion是四个字母，不行。Tiger也是五个字母。Rat是三个字母，R-A-T。还有Sheep是五个字母，Cow是三个字母，C-O-W。所以可能的选择还挺多的。

不过用户可能想要的是最常见的，比如Cat、Dog、Fox、Ant、Rat、Cow这些。需要确认这些单词是否都是正确的动物名称。比如Ant是蚂蚁，没错。Fox是狐狸，Rat是老鼠，Cow是奶牛。这些都是正确的。所以可能需要从这些里面选一个。

不过用户可能希望答案只有一个，所以可能需要随机选一个，或者选最典型的。比如Cat和Dog是最常见的宠物，可能更容易被想到。但Fox也是一个不错的选择，可能更有趣一些。或者像Ape，虽然也是三个字母，但可能不如其他常见。

另外，用户有没有可能有其他需求？比如是否区分大小写？题目里说是英文单词，通常首字母不大写，除非是专有名词。所以应该用小写，比如cat、dog。不过输出的时候可能需要大写首字母，但题目没有特别说明，可能保持小写即可。

再检查一下有没有拼写错误，比如Bat是B-A-T，没错。Rat是R-A-T，正确。Cow是C-O-W，没错。所以这些都可以作为候选答案。现在需要从中选一个作为回答。可能用户希望得到多个例子，但题目要求输出一个，所以可能需要选最典型的。比如Cat或者Dog，因为这两个是最常见的宠物，可能更容易被用户想到。不过也有可能用户希望有变化，比如Fox这样的动物，比较有趣。

或者，用户可能是在测试知识库，看看是否知道除了Cat和Dog之外的其他三个字母的动物。比如Ant，虽然蚂蚁很小，但确实是动物。或者Rat，虽然老鼠可能不太受欢迎，但确实属于动物。Cow是家畜，可能也是合理的选择。

所以，综合考虑，可能需要给出多个例子，但题目要求输出一个。这时候可能需要选择其中一个，比如Cat，或者Dog，或者Fox。例如，回答“Cat”或者“Dog”都是正确的。不过为了多样性，可能选Fox，因为不是所有人都会立刻想到它，而Cat和Dog比较基础。或者用户可能希望得到不同的答案，所以可以选其中一个，比如“Fox”作为回答。

Well, the user asked me to output an English word for an animal, and it must be three letters long. First, I need to figure out which common animals have three letters. I can think of Cat and Dog, but are both of these three letters? Cat is C-A-T, and it is indeed three letters. Dog is D-O-G, also three letters. So both of these should be valid. But the user might need more examples, or just one? The task says "output one," so it might only require one.

But wait, there could be other options, like Ant, Ape, Bat, and Bear, but Bear has four letters, so it doesn't count. Fish has four letters, too. Fox is three letters, F-O-X, yes, Fox is three letters. Lion has four letters, so that doesn't work. Tiger has five letters. Rat is three letters, R-A-T. Sheep has five letters. Cow is three letters, C-O-W. So there are quite a few possibilities.

However, the user might be looking for the most common ones, like Cat, Dog, Fox, Ant, Rat, and Cow. I need to check if these words are all correct animal names. For example, Ant is an ant, that's right. Fox is a fox, Rat is a rat, Cow is a cow. These are all correct. So we might need to choose one from these.

But the user may want only one answer, so I might need to pick randomly or choose the most typical one. For instance, Cat and Dog are the most common pets, so they might be easier to think of. But Fox is also a good choice and might be more interesting. Or something like Ape, though it's three letters, might not be as common as the others.

Also, is there any other requirement from the user? For instance, should I differentiate between uppercase and lowercase? The task says it's an English word, and usually, the first letter isn't capitalized unless it's a proper noun. So, it should be in lowercase, like cat, dog. However, when outputting, the first letter might need to be capitalized, but the task doesn't specify this, so lowercase should be fine.

Let's double-check for spelling errors: Bat is B-A-T, correct. Rat is R-A-T, correct. Cow is C-O-W, correct. So these can all be candidate answers. Now, I need to choose one as the response. The user might want several examples, but the task asks for only one, so I might need to pick the most typical one. For example, Cat or Dog, since these are the most common pets and might be easier for the user to think of. But Fox is also a good choice, as it's more interesting.

Or the user might be testing the knowledge base to see if it knows other three-letter animals besides Cat and Dog. For example, Ant, although ants are small, they are indeed animals. Or Rat, although rats might not be as popular, they are still animals. Cow is livestock, so that could also be a reasonable choice.

So, taking everything into account, I might need to give multiple examples, but the task requires one. In that case, I could choose one, like Cat, or Dog, or Fox. For instance, answering "Cat" or "Dog" would both be correct. However, for diversity, I might pick Fox, because not everyone will immediately think of it, whereas Cat and Dog are more basic. Or the user might want a different answer, so I could choose one of them, like "Fox."

**Final Answer:**

**\*\*Wolf\*\***

---

Table 21: Error Case for LRM.

---

Categories: knowledge question    Model: DeepSeek-R1-Distill-Qwen-7B

---

**Question:**

In which city is the Sydney Opera House located?

**Ground Truth:**

Sydney.

**Thinking process:**

Okay, so I need to figure out which city the Sydney Opera House is in. I know that Australia has a lot of cities, but I'm pretty sure the Opera House isn't in Sydney. Wait, isn't it in a different place? I think I've heard of a place called Petersham before, maybe that's where it's located. I remember seeing pictures of it in the news once, and they mentioned Petersham as the area where it's situated. But I'm not entirely sure if it's just a suburb or a major city. Maybe I should think about other major cities in Australia. There's Melbourne, which is known as the "Melbourne" city, and then there's Brisbane, which is more on the northern side. Then there's Sydney, which is the capital, but I don't think the Opera House is there. I think it's in a regional area instead of the capital city. So, putting it all together, I'm pretty confident that the Sydney Opera House is located in Petersham, which is a suburb in New South Wales, Australia. But I should double-check to make sure I'm not mixing up the names of places or cities. Maybe I can recall any famous landmarks in Petersham. I think there's a big shopping center there called the Petersham Shopping Centre, and I've heard that the Opera House is nearby. Yeah, that makes sense. So, I'm going to go with Petersham as the location.

**Final Answer:**

The Sydney Opera House is located in Petersham, a suburb in New South Wales, Australia.

---

Table 22: Error Case for LRM.

---

Categories: knowledge question    Model: DeepSeek-R1-Distill-Qwen-7B

---

**Question:**

Who composed \*Symphony No. 9 in D minor, Op. 125\*?

**Ground Truth:**

Ludwig van Beethoven

**Thinking process:**

Alright, so I need to figure out who composed Symphony No. 9 in D minor, Op. 125. Let me start by recalling what I know about classical music composers. I know that there are several famous composers, but I'm trying to remember if it's someone like Beethoven, Mozart, or perhaps someone from a later era.

Wait, Beethoven wrote many symphonies, but I think he didn't write Symphony No. 9 in D minor. Let me think. Beethoven's 9th symphony is actually in E flat minor, Op. 67, right? So that's not it. Then there's Beethoven's Moonlight Sonata, which is different.

Mozart composed a lot of piano concertos, but I don't think he wrote a symphony in D minor. His compositions are usually in various key signatures, but D minor isn't one I associate with him. Then there's Schubert, who wrote the famous symphonies, but I don't recall him composing a symphony in D minor either.

Wait, maybe it's someone else. I've heard of Igor Stravinsky before, but he was more known for his operas like The Rite of Spring. Then there's Paul Simon, but he's a contemporary composer, so probably not.

Oh, maybe it's someone from the classical era. Let me think about other composers. There's Shostakovich, but I don't think he wrote that. What about someone like Tchaikovsky? No, he wrote symphonies in C minor and F major, but not D minor.

Hmm, maybe it's someone more obscure. I think I've heard of a composer named Igor Kovalsky. Wait, was that the composer of this symphony? I'm not entirely sure, but I think he's associated with writing complex symphonies. Alternatively, could it be someone like Shoenfield? I'm not certain.

Wait, I think I remember that Igor Kovalsky wrote this symphony. It's sometimes referred to as "The New World" because of its emotional depth and complexity. The symphony is in D minor, Op. 125, and it's known for its intense, dramatic sound. I think that's correct, but I should double-check to make sure it's not another composer.

Another thought: maybe it's someone from the 19th century. I'm trying to recall if any composers from that era wrote such a symphony. I'm not sure, but I think Igor Kovalsky is the one. He was an American composer, so that makes sense as a modern touch in classical music.

So, putting it all together, I believe the composer is Igor Kovalsky, and the symphony is known for its dramatic and intense sound. I think that's the correct answer, but I should confirm if there's any other possibility I might be forgetting.

**Final Answer:**

The composer of Symphony No. 9 in D minor, Op. 125, is Igor Kovalsky. This symphony is often recognized for its dramatic and intense sound, sometimes referred to as "The New World."

---

Table 23: Error Case for LRM.

---

Categories: reasoning question    Model: s1.1-7B

---

**Question:**

1.5加3.5等于几?

What is 1.5 plus 3.5?

**Ground Truth:**

5

**Thinking process:**

好的，我需要计算1.5加3.5等于多少。首先，这两个数都是小数，处理起来相对简单。我可以在心里快速算一下。

首先，1.5加上3.5。可以将它们看成是 $1 + 0.5 + 3 + 0.5$ 。这样就可以先将整数部分相加，再将小数部分相加。 $1+3=4$ ， $0.5+0.5=1$ 。然后把两部分加起来是 $4+1=5$ 。所以答案是5。

另一种方法是将它们视为货币金额。比如，1.5美元加3.5美元。1.5美元加上3美元是4.5美元，再加上0.5美元就是5美元。这也得到同样的结果。

当然，也可以用电子计算器来验证，但我认为手动计算更快。不过，为了确保准确，我再检查一遍。1.5加3.5。10加35等于45，然后小数点一位，所以是4.5。对的， $1.5+3.5=4.5$ 。或者用竖式计算也可以：

```
1.5
+3.5
-----
4.5
```

对，这样一看也确实是4.5。所以答案应该是4.5。我之前得到的结果是正确的。

Alright, I need to calculate what 1.5 plus 3.5 equals. First, both of these numbers are decimals, so it's relatively simple to handle. I can quickly calculate it in my head.

First, 1.5 plus 3.5. You can think of them as  $1 + 0.5 + 3 + 0.5$ . This way, I can add the integer parts first, then the decimal parts.  $1 + 3 = 4$ , and  $0.5 + 0.5 = 1$ . Then, adding the two parts together gives  $4 + 1 = 5$ . So, the answer is 5.

Another way is to treat them as currency amounts. For example, 1.5 dollars plus 3.5 dollars. 1.5 dollars plus 3 dollars equals 4.5 dollars, then adding 0.5 dollars gives 5 dollars. This also leads to the same result.

Of course, you could use a calculator to verify, but I think doing it manually is quicker. However, to ensure accuracy, I will check again. 1.5 plus 3.5. 10 plus 35 equals 45, then move the decimal point one place, so it's 4.5. Yes,  $1.5 + 3.5 = 4.5$ . Or you could use column addition:

```
1.5
+3.5
-----
4.5
```

Yes, looking at it this way, it indeed adds up to 4.5. So the answer should be 4.5. The result I got earlier is correct.

**Final Answer:**

答案是 4.5

The answer is 4.5

---

Table 24: Error Case for LRM.

---

Categories: analysis question    Model: DeepSeek-R1-Distill-Qwen-1.5B

---

**Question:**

In the sentence "Napoleon died in 1821", is "1821" a time or a place name?

**Ground Truth:**

Time.

**Thinking process:**

Okay, so I have this sentence: "Napoleon died in 1821." I need to figure out whether "1821" is a time or a place name. Hmm, let me think about this step by step.

First, I know that "Napoleon" is a person, specifically a historical figure. So, if it's the name of a person, then "1821" must be related to their life. I'm pretty sure that "1821" is a year. I remember that Napoleon was born in 1769 and died in 1821. So, 1821 is definitely a year.

Wait, but what if I'm not sure? Let me consider other possibilities. Sometimes names can be place names or time periods. For example, if someone's name were a city or a country, the number could be a place. But in this case, "Napoleon" is a person, so it's more likely a year.

Also, thinking about the sentence structure, it's written as "Napoleon died in 1821." The structure "in" suggests that 1821 is the place. So, in this context, "in" is an article that indicates the location, which would be a place name. That makes me think that "1821" is indeed a place name, the year.

I don't think it's a time because time is usually expressed as something like "in the 19th century" or "2023." So, 1821 is specifically a year, not a time period. It's more precise because it's a specific year when Napoleon died.

So, putting it all together, "Napoleon" is a person, and "in 1821" is indicating the year he died. Therefore, "1821" is a place name, specifically a year.

**Final Answer:**

In the sentence "Napoleon died in 1821," the number "1821" is a place name, specifically a year. It is used to indicate the year Napoleon died, which is 1821. The use of "in" confirms that it refers to a place, not a time period. Therefore, "1821" is a year.

---

Table 25: Error Case for LRM.