# THINK-Bench: Evaluating Thinking Efficiency and Chain-of-Thought Quality of Large Reasoning Models

Anonymous ACL submission

#### Abstract

Large reasoning models (LRMs) have achieved impressive performance in complex tasks, often outperforming conventional large language models (LLMs). However, the prevalent issue of overthinking severely limits their computational efficiency. Overthinking occurs when models generate excessive and redundant tokens that contribute little to accurate outcomes, especially in simple tasks, resulting in a significant waste of computational resources. To systematically investigate this issue, we introduce Think-Bench, a benchmark designed to evaluate the reasoning efficiency of LRMs. We also propose novel efficiency metrics and conduct a comprehensive evaluation of various LRMs across multiple dimensions, including the reasoning process, outcome quality, and chain-ofthought (CoT) characteristics. Our analysis reveals that most LRMs exhibit overthinking in handling easy questions, generating unnecessarily lengthy reasoning chains. While many LRMs demonstrate high CoT quality, several suffer from low efficiency. We hope that Think-Bench can serve as a robust foundation for advancing research into LRMs.

#### 1 Introduction

011

012

017

023

024

042

In recent years, with the rapid advancement of artificial intelligence, LLMs have achieved remarkable success in the field of natural language processing (NLP), particularly excelling in tasks such as text generation and question answering (Grattafiori et al., 2024; Guo et al., 2025; Yang et al., 2024). However, the reasoning capabilities of these models in solving multidisciplinary problems still face significant challenges, such as insufficient integration of cross-disciplinary knowledge and weak logical chain reasoning ability (Wang et al., 2024; Chen et al., 2025). To gain a deeper understanding and effectively enhance the efficiency and accuracy of the reasoning capabilities of LLMs, constructing highquality multidisciplinary datasets and conducting



Figure 1: The performance of various LRMs on Think-Bench. The results suggest that these prominent LRMs face a challenge of overthinking.

systematic evaluations has become critically important (Chang et al., 2024; Xia et al., 2024).

Recent research have introduced several highquality multidisciplinary benchmarks, such as MMLU (Hendrycks et al., 2020a) and GPQA (Rein et al., 2024), to evaluate model performance across diverse knowledge domains. However, these datasets predominantly rely on the correctness of the final answer as the sole evaluation metric. Such outcome-oriented evaluation protocols are not directly applicable to the assessment of large reasoning models (LRMs), as they neglect the accuracy and logical coherence of intermediate reasoning steps during answer generation (Jiang et al., 2025; Zheng et al., 2024). In reality, for complex tasks spanning logical reasoning, mathematical problemsolving, and multi-step decision-making, the quality of intermediate CoT processes not only determines final-answer accuracy but also serves as a critical indicator of LRMs' intrinsic reasoning capabilities and operational reliability. Existing evaluation paradigms that ignore CoT quality risk



Figure 2: **Overview of Think-Bench.** Our benchmark contains a comprehensive efficiency evaluation framework with curated datasets across three categories.

masking systemic deficiencies, such as erroneous premises, logical fallacies, or context misinterpretations, which may propagate through the reasoning chain. Consequently, establishing a holistic evaluation framework that quantifies both intermediatestep correctness and final-output accuracy emerges as a pivotal yet understudied research direction.

067

072

074

077

081

Recently, several studies have begun to explore the evaluation of the correctness of CoTs. Notable works such as MME-CoT (Jiang et al., 2025) establish a comprehensive benchmark that systematically examines CoTs across three fundamental dimensions: reasoning quality, robustness, and efficiency. Meanwhile, MiCEval (Zhou et al., 2024) conducts a fine-grained evaluation of Multimodal Chains of Thought (MCoT), focusing on aspects such as the accuracy of image descriptions and the correctness, relevance, and informativeness of reasoning steps. Collectively, these benchmarks represent a significant shift from solely outcomeoriented evaluations to more process-aware assessments that effectively capture the internal reasoning dynamics of large models.

Despite these advances, existing benchmarks remain limited by their lack of a systematic framework for evaluating both the efficiency and reliability of reasoning processes in LRMs. This challenge is particularly prominent in multidisciplinary problem-solving scenarios (Wang et al., 2025). Prior studies have demonstrated that reasoning models structurally similar to OpenAI-01 often allocate computational resources inefficiently when handling relatively simple problems, while achieving only negligible performance improvements in final outcomes (Chen et al., 2024; Guo et al., 2025). These findings suggest that LRMs may suffer from suboptimal resource distribution during inference, consequently constraining their overall performance optimization.

097

100

101

102

103

104

105

107

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

To bridge this critical research gap, we introduce **Think-Bench**, a multidisciplinary dataset specifically designed for comprehensively evaluating the reasoning efficiency and accuracy of LRMs. Our benchmark features meticulously annotated key reasoning steps for each problem instance, enabling granular analysis of model-generated reasoning processes. Additionally, we propose an evaluation protocol that can measure the efficiency and rationality of the reasoning process by analysing the specific behaviours exhibited by LRMs during their thinking process.

## **2** Dataset Curation

#### 2.1 Data Overview

As shown in Figure 2, Think-Bench is a dataset specifically designed to evaluate the thinking efficiency and the quality of CoTs of LRMs in complex reasoning tasks. This dataset comprises 1,375 carefully selected and organized data samples, covering three core subjects: mathematics, physics, and chemistry. Within each subject, the number of simple questions is approximately



Figure 3: Category and Subcategory Distribution of Think-Bench.

equal to the number of difficult questions. The
data sources are diverse, drawing from multiple
academic datasets, including MMLU (Hendrycks
et al., 2020a), Math500 (Hendrycks et al., 2020b),
AGIEval (Zhong et al., 2023), AIME (Veeraboina,
2023), GPQA (Rein et al., 2024), SciKnowEval (Feng et al., 2024), and UGPhysics (Xu et al.,
2025b).

#### 2.2 Data Collection

134

135 During the construction of Think-Bench, we aggregated questions from multiple authoritative, pub-136 licly available datasets. The distribution of the 137 Think-Bench across different disciplines is shown in Table 1. To ensure fairness, all samples were 139 selected randomly. After the selection process, we 140 conducted a systematic data cleaning and verifi-141 cation procedure to remove duplicate and invalid 142 entries. The final dataset consists of 1,375 data 143 points, after which we carried out the data anno-144 tation work. This benchmark covers the core dis-145 ciplines of mathematics, physics, and chemistry, 146 which inherently require structured and multi-step 147 148 reasoning. Therefore, it provides a robust and rigorous foundation for evaluating the performance 149 of reasoning models. Detailed statistics regarding 150 the data composition can be found in Figure 3 and Table 1. 152

Statistic	Number
Total questions	1375
-Choice questions	929
-Free-form questions	446
-Math questions	293
-Physics questions	590
-Chemistry questions	492
Total key step annotation	13311
-Average inference step	9.68
Maximum question length	1893
Maximum answer length	372
Average question length	422.42
Average answer length	7.59

Table 1: Key Statistics of Think-Bench.

153

154

155

156

157

158

159

160

161

163

164

165

167

168

169

170

171

172

173

174

175

176

177

178

179

180

181

#### 2.3 Data Annotation and Review

To systematically evaluate the CoT reasoning capabilities of LRMs on reasoning tasks, we implement a fine-grained annotation framework for key reasoning steps across all questions. Key steps are defined as essential logical components that must be completed and cannot be omitted in the process of achieving the correct answer. The annotation process is implemented in the following steps: First, we leverage Claude 3.7 Sonnet to generate comprehensive reasoning chains for each question (see Prompt 1 in Appendix D), which serve as baseline references. Based on the generated reasoning paths, we identify and extract the critical steps involved. Since a single question may admit multiple logically valid reasoning paths, all reasonable and logically consistent solutions are considered and included.

# **3** Evaluation Strategy

Recent studies have revealed that LRMs frequently exhibit low reasoning efficiency (Sui et al., 2025; Chen et al., 2024). However, there is currently a lack of systematic benchmarks to evaluate this issue comprehensively. A detailed analysis of LRMs' reasoning processes is crucial for understanding their efficiency limitations and underlying challenges. Furthermore, existing benchmarks primarily assess the final answers to reasoning questions, neglecting the intermediate CoT steps. To bridge



Figure 4: Illustration of Thinking Efficiency and CoT Quality Evaluation.



Figure 5: Example of Thinking Process Analysis in a LRM.

this gap, we propose a novel benchmark that jointly
evaluates both the efficiency and quality of reasoning CoTs, thereby enabling a more holistic assessment of LRMs' reasoning capabilities.

## 3.1 Thinking Efficiency Evaluation

With the rapid development of LRMs, their ability to handle complex multi-step reasoning has significantly advanced (Xu et al., 2025a). Notable models like OpenAI-o1 (Zhong et al., 2024), DeepSeek-R1 (Guo et al., 2025), and Qwen3 (Yang et al., 2025a) have attracted growing interest for their human-like capacity for extended, reflective reasoning. Through advanced long CoT and test-time scaling methods, these models iteratively evaluate multiple reasoning paths before finalizing answers (Chen et al., 2025; Muennighoff et al., 2025). 186

187

188

189

190

191

192

193

194

195

196

197

198

199

202

203

205

206

207

208

210

However, as test-time scaling consumes increasing computational resources, a critical challenge has emerged in LRMs' inference behaviors: **Overthinking**. This refers to the model's persistent tendency to engage in excessive and repetitive reasoning, often producing reasoning chains that span hundreds of tokens even for simple tasks. While such elaborate verification is justified for complex problems, test-time scaling amplifies this behavior, causing unnecessary computational overhead and inefficiency during inference for simple inputs.

As illustrated in Figures 4 and 5, we propose six complementary metrics to systematically assess

303

258

259

260

261

reasoning efficiency across token usage, inference dynamics, and reflective quality.

211

212

213

214

215

216

217

218

219

222

224

234

235

237

240

241

243

245

246

247

248

249

254

257

**Tokens** measure the total token count processed before final prediction, representing reasoning chain length and providing a fundamental basis for computational cost estimation.

**First Correct Tokens** measures the token count from reasoning initiation until the first occurrence of a correct answer. This metric evaluates the model's speed in reaching a valid solution during reasoning, where fewer tokens indicate faster correct convergence. The identification prompt is detailed in Prompt 5.

**Efficiency** is a normalized metric that refers to the ratio of first correct tokens to the total number of reasoning tokens. Formally, it is defined as:

Efficiency = 
$$\frac{1}{N} \sum_{i=1}^{N} \frac{\hat{T}_i}{T_i}$$
. (1)

where  $\hat{T}_i$  denotes the number of tokens generated by the model before the first occurrence of the correct answer in its response, and let  $T_i$  represent the total number of reasoning tokens for the *i*-th instance. If the model fails to produce a correct answer, we set  $\hat{T}_i = 0$ . A higher value of this metric indicates more efficient reasoning behavior. Concrete examples illustrating this calculation can be found in Figure 7 in Appendix C.1.

**Reflection Quality** measures the efficacy of the model's self-reflective reasoning, particularly after producing a correct answer. Not all reflective steps contribute meaningfully: some merely reiterate prior conclusions, while others may introduce erroneous revisions. We define a valid reflection as one that either (i) accurately identifies a prior error or (ii) provides new insights that confirm an earlier conclusion. Let R represent the total set of reflective steps, and  $R_{valid}$  denote the subset of valid reflections. The metric is defined as follows:

Reflection Quality 
$$= \frac{|R_{\text{valid}}|}{|R|}$$
. (2)

This score quantifies the efficacy of the model's reflection process, where higher values indicate more meaningful self-verification behavior as opposed to producing redundant or counterproductive content. The prompt used to guide this reflection process is provided in Prompt 6.

**Reflection Tokens** quantify the token count generated from the first correct answer to the conclusion of the reasoning process. This segment typically encompasses verification steps, reflective analysis, and conclusion restatements. Although such content may provide valuable insights, excessive length often signals reasoning inefficiency or unnecessary repetition.

**Thought Num** measures how often the model changes reasoning paths. This metric is estimated by counting discourse markers like "alternatively," "on second thought," and "wait a moment." A higher count may indicate instability in reasoning or a tendency toward exploratory behavior.

#### 3.2 CoT Quality Evaluation

As LLMs increasingly adopt CoT reasoning strategies, assessing the quality of their internal reasoning processes has emerged as a critical research challenge (Jiang et al., 2025). Existing evaluation approaches predominantly focus on final answer accuracy (Wang et al., 2019; Hendrycks et al., 2020b; Suzgun et al., 2022), while largely overlooking the validity and robustness of intermediate reasoning steps. To bridge this gap, we adopt a reference-based evaluation framework, inspired by MME-CoT (Jiang et al., 2025). Our proposed framework measures the reasoning quality from two interpretable dimensions: **Recall** and **Precision**.

As illustrated in Figure 4, each CoT response is decomposed into multiple reasoning steps through the prompt detailed in Prompt 3.

$$R = \{r_1, r_2, \dots, r_M\}.$$
 28

To evaluate its quality, R is compared against a pre-annotated reference set containing key reasoning components.

$$S = \{s_1, s_2, \dots, s_N\}.$$
 291

Each  $r_j$  is judged for semantic alignment with any  $s_i$ , using Claude 3.7 Sonnet as a judge guided by consistent prompting instructions. The prompt designed to extract matching steps for computing recall and precision is provided in Prompt 2 and Prompt 3 of Appendix D. We define:

- *R*<sub>match</sub> ⊆ *R*: the subset of reasoning steps in *R* that correctly match at least one reference step in *S*.
- $S_{\text{covered}} \subseteq S$ : the subset of reference steps that are successfully matched by at least one step in R.

Model name	Efficiency	Recall	Precision	Accuracy	Reflection Quality	Thought Num	Tokens	Useful Tokens	Reflection Tokens
Claude-3-7-sonnet	49.61%	81.29%	86.26%	94.25%	76.49%	0.28	942.82	446.09	496.73
Deepseek-r1-distill-qwen-1.5b	37.14%	47.10%	59.61%	62.91%	61.88%	8.00	3734.49	1268.36	2466.13
Deepseek-r1-distill-qwen-7b	49.53%	63.65%	77.29%	68.51%	77.70%	9.42	3504.76	1641.91	1862.85
Deepseek-r1-distill-qwen-14b	50.70%	61.04%	79.97%	70.18%	82.40%	7.04	2814.75	1413.09	1401.66
Deepseek-r1-distill-qwen-32b	52.62%	64.17%	83.76%	75.93%	84.46%	6.27	2697.70	1352.93	1344.77
Deepseek-r1	48.96%	80.80%	88.33%	88.80%	90.92%	9.17	3795.19	1912.12	1883.07
Ernie-x1-turbo-32k	47.02%	82.03%	88.67%	89.89%	90.97%	12.75	4692.21	2221.32	2470.89
Grok-3-mini-beta	61.69%	81.56%	86.51%	91.85%	88.20%	0.38	1891.34	1169.05	722.29
Qwen3-235b-a22b	46.14%	85.80%	86.97%	94.91%	92.16%	13.35	4969.05	2448.29	2520.76
Qwq-plus	44.58%	80.40%	85.08%	89.60%	89.67%	22.63	5738.37	2646.73	3091.64
Glm-z1-air	47.41%	80.16%	83.18%	88.06%	89.17%	9.80	3678.68	1775.07	1903.61

Table 2: Evaluation of Nine Metrics on CoT and Efficiency in Think-Bench. Best performance in bold.

Using  $R_{\text{match}}$  and  $S_{\text{covered}}$ , we compute the **Re**call and **Precision** metrics as follows:

$$\mathbf{Recall} = \frac{|S_{\text{covered}}|}{|S|} \tag{3}$$

$$\mathbf{Precision} = \frac{|R_{\text{match}}|}{|R|} \tag{4}$$

Recall measures the extent to which essential reasoning steps are accurately captured in the LRM's output, reflecting the informativeness and comprehensiveness of the generated reasoning chain. In contrast, precision evaluates the correctness and relevance of the reasoning steps, penalizing any instance of inaccuracy or logical inconsistency.

# 4 Experiment

304

305

306

309

312

313

314

315

316

#### 4.1 Experimental Setup

**Evaluation Models** To systematically evaluate 317 318 both the efficiency and quality of reasoning with CoT in LRMs, we select eleven representative mod-319 els spanning diverse architectures and parameter 320 scales. Our evaluation encompasses both propri-321 etary and open-source LRMs. Specifically, we 322 include Claude 3.7 Sonnet (Anthropic, 2025), a 323 proprietary model widely recognized for its strong 324 performance in multi-turn reasoning tasks. We 325 also conduct a comprehensive assessment of the DeepSeek-R1 family (Guo et al., 2025), includ-327 ing the full-scale DeepSeek-R1 and its distilled Qwen-1.5-based variants at 1.5B, 7B, 14B, and 32B scales, all explicitly optimized for efficient multi-331 step reasoning. Additionally, we evaluate Qwen3-235B-A22B (Team, 2025b) and Qwq-Plus (Team, 332 2024), both equipped with reflection and alignment mechanisms to support long-context inference. To further explore model behavior under extended 335

reasoning conditions, we include Ernie-X1-Turbo-32K (Team, 2025a), optimized for long input sequences, along with Grok-3-Mini-Beta (xAI, 2025) and GLM-Z1-Air (GLM et al., 2024). 336

337

338

339

340

341

342

343

344

345

346

347

348

350

351

352

353

354

355

356

358

359

360

361

362

363

364

365

366

367

368

369

370

371

372

**Implementation Details** Throughout the evaluation process, we initially employed the tested LRMs to generate responses to the entries from Think-Bench. All other model hyperparameters followed default settings unless otherwise specified. Subsequently, Claude 3.7 Sonnet was utilized to analyze the reasoning steps and underlying thinking processes of these responses. The detailed prompt used for the analysis with Claude 3.7 Sonnet is provided in Appendix D. Finally, we computed our proposed evaluation metrics to assess the thinking efficiency and reliability of the tested LRMs.

#### 4.2 Quantitative Results

We conduct a comprehensive evaluation of LRMs using our proposed Think-Bench. The main results are presented in Tables 2 and Tables 3. We begin with an analysis of the overall performance, followed by an in-depth discussion of the key findings.

**Overall Performance** To comprehensively evaluate the performance of LRMs, we report results across two key dimensions: efficiency and CoT quality, including our proposed efficiency metrics, recall, precision and accuracy, as shown in Table 2. Our analysis shows that while there exists a consistent trade-off between token usage and reasoning performance, different models exhibit significant variability in their inference behaviors.

In terms of efficiency, Grok-3-mini-beta achieves the highest score of 61.69%, followed by Deepseek-r1-distill-qwen-32b at 52.62% and Deepseek-r1-distill-qwen-14b at 50.70%, indicating a more economical use of tokens to reach cor-

	Recall		Precision		Reflection Quality		Tokens		Efficiency	
Model name	Simple	Difficult	Simple	Difficult	Simple	Difficult	Simple	Difficult	Simple	Difficult
Claude-3-7-sonnet	88.49%	74.05%	92.94%	79.56%	92.95%	90.29%	673.24	1216.01	0.52	0.47
Deepseek-r1-distill-qwen-1.5b	55.74%	38.42%	69.35%	49.84%	73.22%	68.55%	2149.94	5325.97	0.40	0.34
Deepseek-r1-distill-qwen-7b	70.75%	51.30%	88.84%	71.06%	90.09%	83.64%	1575.27	4059.64	0.51	0.51
Deepseek-r1-distill-qwen-14b	72.99%	55.31%	92.55%	74.93%	91.43%	85.07%	1514.53	3886.05	0.52	0.53
Deepseek-r1-distill-qwen-32b	69.88%	57.39%	85.79%	68.75%	84.99%	78.91%	2074.32	4941.47	0.49	0.50
Deepseek-r1	88.80%	72.77%	95.54%	81.09%	95.15%	90.00%	2058.35	5539.63	0.46	0.52
Ernie-x1-turbo-32k	90.26%	73.76%	95.44%	81.88%	95.00%	89.09%	2679.32	6713.91	0.43	0.51
Grok-3-mini-beta	89.09%	74.01%	93.47%	79.51%	93.17%	88.78%	1242.27	2543.25	0.60	0.63
Qwen3-235b-a22b	92.87%	78.70%	95.33%	78.57%	96.29%	91.43%	2818.69	7128.80	0.42	0.50
Qwq-plus	90.04%	70.72%	94.89%	75.23%	94.60%	87.42%	3289.45	8197.99	0.41	0.48
GLM-Z1-Air	88.69%	71.60%	93.31%	72.99%	94.79%	86.55%	1931.36	5433.63	0.45	0.49

Table 3: Evaluation Results of CoT and Efficiency in Think-Bench Classified by Difficulty Levels. Best performance in **bold**.

rect answers. In contrast, larger models such as Qwen3-235b-a22b and Qwq-plus exhibit lower efficiency, scoring 46.14% and 44.58% respectively. This decrease in efficiency is attributed to their prolonged reasoning chains, despite having strong CoT quality.

374

375

376

378

395

396

400

401

402

403

404

405

406

407

Regarding CoT quality, Qwen3-235b-a22b and Ernie-x1-turbo-32k stand out by achieving the highest reflection quality scores, with values of 92.16% and 90.97%, respectively. They also demonstrate top-tier precision at 86.97% and 88.67%, and recall rates of 85.80% and 82.03%. These impressive results highlight the advantages of large-scale models with reflection-enhanced reasoning capabilities, which not only lead to accurate conclusions but also enable reliable verification and correction processes. In contrast, smaller distilled models, such as Deepseek-r1-distill-qwen-1.5b, perform poorly across all quality metrics, particularly in precision (59.61%) and recall (47.10%).

An important behavioral indicator is Thought Num, reflecting how often the model switches or reconsiders its CoT. Qwq-plus shows the highest value (22.63), indicating frequent reflective iterations. However, such reflections don't always lead to better performance and may reduce efficiency. In contrast, models like Claude 3.7 Sonnet and Grok-3-mini-beta maintain very low Thought Num values (0.28 and 0.38) while still achieving a balanced and high-quality reasoning process.

Regarding token consumption, Qwq-plus and Qwen3-235b-a22b each use over 4,900 tokens per response, with a substantial portion from reflection (3,091.64 and 2,520.76, respectively). This suggests tendencies toward overthinking. In contrast, Claude-3-7-Sonnet completes its reasoning in under 1,000 tokens, demonstrating concise and effective inference with minimal redundancy.

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

432

433

434

435

436

437

438

439

440

441

**Evaluation and Analysis Based on Difficulty** Levels of Questions The results in Table 3 show that most LRMs demonstrate significantly lower average efficiency on simple questions than on difficult ones. This suggests that when faced with simple questions, these models tend to overthink and generate unnecessary reasoning chains. In contrast, for high difficulty questions, the models focus more effectively, eliminating redundant reasoning steps and improving efficiency. Additionally, token consumption for difficult questions is consistently higher than for simple ones, due to the extra inferential steps needed to tackle complex problems. As reflection quality, recall, and precision all decline slightly as task difficulty increases, this illustrates that while difficult questions require more computational resources, they present greater challenges to the model's reasoning capabilities.

**Evaluation and Analysis Based on Different Sub-**

**jects** As shown in Figure 6, the multi-disciplinary evaluation in Think-Bench reveals notable performance differences across chemistry, physics, and mathematics. Mathematical tasks generally lead to higher token consumption and lower reasoning efficiency, even for strong models, suggesting a reliance on lengthy CoTs and structured outputs. In contrast, chemistry and physics tasks typically exhibit better efficiency and lower token usage.

Regarding CoT quality evaluation, the recall and precision generally show a positive correlation in various disciplinary tasks within Think-Bench, but



Figure 6: Comparative Performance of Models in Chemistry, Physics, and Math.

there are also obvious structural differences. Specifically, in chemistry and physics tasks, the precision rate of the model is often significantly better than the recall rate. This phenomenon reflects that the current LRMs' strategy in generating answers for reasoning questions is relatively conservative, such that it is more inclined to output answers with high confidence.

#### 4.3 Error Analysis

442

443

444 445

446

447

448

449

450

451

461

During the evaluation experiment, we observed that some models, particularly the distillation mod-452 els from the DeepSeek series and the ERNIE-X1-453 Turbo-32K, exhibited an unusual issue of generat-454 ing empty outputs. This issue primarily manifested 455 in the model generating only intermediate reason-456 ing content without producing a final answer for 457 certain questions. See Figure 8 in Appendix C.2 458 for a concrete example. Potential causes include 459 limitations in their inference mechanisms, context 460 processing capabilities, or deployment implemen-462 tations. As the issue was sporadic and difficult to reproduce systematically, we chose to automati-463 cally skip any samples with invalid outputs to pre-464 serve the integrity of the evaluation and ensure the 465 stability of subsequent analyses. 466

#### 5 Conclusion

This paper presents the Think-Bench dataset, a benchmark designed to systematically evaluate the reasoning efficiency and CoT quality of LRMs. The dataset consists of tasks from three disciplines: mathematics, physics, and chemistry. Each task is provided at two difficulty levels: Simple and Difficult. Evaluation is conducted using nine metrics, including six efficiency indicators, two CoT quality measures and accuracy. To verify the effectiveness of Think-Bench and to assess the reasoning efficiency and CoT quality of mainstream LRMs, we conduct a comprehensive evaluation of 11 representative models. Experimental results show that most models exhibit overthinking behaviors on simple questions, generating excessive reasoning tokens and leading to unnecessary computational overhead. This study not only highlights the limitations of current LRMs in their use of computational resources in inference time, but also offers insights that may inspire future research, including designing dynamic reasoning pathways, early exit mechanisms, and enhancing adaptability across disciplines.

468

469

470

471

472

473

474

475

476

477

478

479

480

481

482

483

484

485

486

487

488

489

# 491 Limitations

The Think-Bench benchmark proposed in this 492 study currently covers only three disciplines: math-493 ematics, physics, and chemistry, which limits its 494 effectiveness in evaluating models' reasoning abili-495 ties across a broader range of subjects or in more complex real-world scenarios. Furthermore, the 497 evaluation process relies on large language models 498 to assess the reasoning steps of target models, intro-499 ducing potential variability due to the performance and stability of the judging models themselves. In this study, Claude 3.7 Sonnet is utilized both as the model being evaluated and as the judging model, 503 which may further enhance this issue by introducing bias in the assessment. In addition, although 505 the dataset annotation attempts to incorporate a variety of valid reasoning paths, it is challenging to exhaust all possible solution strategies, which may lead to incomplete evaluations of models that adopt 509 reasonable but unannotated reasoning approaches. 510

#### References

511

513

514

515

516

517

518

519

521

522

523

524

525

526

527

528

532

534

535

537

538

539 540

- 512 Anthropic. 2025. Claude 3.7 Sonnet and Claude Code.
  - 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 Wanxiang Che. 2025. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*.
    - Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, and 1 others. 2024. Do not think that much for 2+ 3=? on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*.
    - Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, and 1 others. 2025. The danger of overthinking: Examining the reasoning-action dilemma in agentic tasks. *arXiv preprint arXiv:2502.08235*.
  - Chenrui Fan, Ming Li, Lichao Sun, and Tianyi Zhou. 2025. Missing premise exacerbates overthinking: Are reasoning models losing critical thinking skill? *arXiv preprint arXiv:2504.06514*.

Kehua Feng, Keyan Ding, Weijie Wang, Xiang Zhuang, Zeyuan Wang, Ming Qin, Yu Zhao, Jianhua Yao, Qiang Zhang, and Huajun Chen. 2024. Sciknoweval: Evaluating multi-level scientific knowledge of large language models. *arXiv preprint arXiv:2406.09098*. 541

542

543

544

545

546

547

548

549

550

551

552

553

554

555

556

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

582

583

584

585

586

587

588

589

590

591

592

593

594

595

596

- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, and 37 others. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *Preprint*, arXiv:2406.12793.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The Ilama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shi-rong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in Ilms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020a. Measuring massive multitask language understanding. arXiv preprint arXiv:2009.03300.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020b. Measuring massive multitask language understanding. arXiv preprint arXiv:2009.03300.
- Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanwei Li, Yu Qi, Xinyan Chen, Liuhui Wang, Jianhan Jin, Claire Guo, Shen Yan, and 1 others. 2025. Mmecot: Benchmarking chain-of-thought in large multimodal models for reasoning quality, robustness, and efficiency. *arXiv preprint arXiv:2502.09621*.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*.
- Nisarg Patel, Mohith Kulkarni, Mihir Parmar, Aashna Budhiraja, Mutsumi Nakamura, Neeraj Varshney, and Chitta Baral. 2024. Multi-logieval: Towards evaluating multi-step logical reasoning ability of large language models. *arXiv preprint arXiv:2406.17169*.
- Xiao Pu, Michael Saxon, Wenyue Hua, and William Yang Wang. 2025. Thoughtterminator: Benchmarking, calibrating, and mitigating overthinking in reasoning models. *arXiv preprint arXiv:2504.13367.*
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.

- Tomohiro Sawada, Daniel Paleka, Alexander Havrilla, Pranav Tadepalli, Paula Vidas, Alexander Kranias, John J Nay, Kshitij Gupta, and Aran Komatsuzaki. 2023. Arb: Advanced reasoning benchmark for large language models. arXiv preprint arXiv:2307.13692.
  - Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Hanjie Chen, and 1 others. 2025. Stop overthinking: A survey on efficient reasoning for large language models. *arXiv preprint arXiv:2503.16419*.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, and 1 others. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Ernie Team. 2025a. Ernie x1.

599

607

608

610

611

613

615

617

618

619

620

622

623

625

626

627

631

633

639

642

647

- Qwen Team. 2024. QwQ: Reflect Deeply on the Boundaries of the Unknown. *Qwen*.
- Qwen Team. 2025b. Qwen3.
  - Hemish Veeraboina. 2023. Aime problem set 1983-2024.
  - Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32.
  - Yiqi Wang, Wentao Chen, Xiaotian Han, Xudong Lin, Haiteng Zhao, Yongfei Liu, Bohan Zhai, Jianbo Yuan, Quanzeng You, and Hongxia Yang. 2024. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805*.
  - Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, and 1 others. 2025. Thoughts are all over the place: On the underthinking of o1-like llms. *arXiv preprint arXiv:2501.18585*.
  - 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. *Advances in neural information processing systems*, 35:24824– 24837.
  - xAI. 2025. Grok 3 beta the age of reasoning agents. https://x.ai/blog/grok-3. Accessed: 2025-02-21.
  - Tingyu Xia, Bowen Yu, Yuan Wu, Yi Chang, and Chang Zhou. 2024. Language models can evaluate themselves via probability discrepancy. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 4889–4901.

Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, and 1 others. 2025a. Towards large reasoning models: A survey of reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*. 653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

- Xin Xu, Qiyun Xu, Tong Xiao, Tianhao Chen, Yuchen Yan, Jiaxin Zhang, Shizhe Diao, Can Yang, and Yang Wang. 2025b. Ugphysics: A comprehensive benchmark for undergraduate physics reasoning with large language models. *arXiv preprint arXiv:2502.00334*.
- 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. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Zheng Lin, Li Cao, and Weiping Wang. 2025b. Dynamic early exit in reasoning models. *arXiv preprint arXiv:2504.15895*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023a. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023b. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.
- Chujie Zheng, Zhenru Zhang, Beichen Zhang, Runji Lin, Keming Lu, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. 2024. Processbench: Identifying process errors in mathematical reasoning. *arXiv preprint arXiv:2412.06559*.
- Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu, Yanjun Lyu, Peng Shu, Xiaowei Yu, and 1 others. 2024. Evaluation of openai o1: Opportunities and challenges of agi. *arXiv preprint arXiv:2409.18486*.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. Agieval: A human-centric benchmark for evaluating foundation models. *arXiv preprint arXiv:2304.06364*.

794

795

796

797

798

799

800

801

802

803

804

805

806

757

758

706 707

# 709 710

711

712

713

714

715

716

717

718

719

720

721

723

724

725

726

727

728

731

733

734

735

739

740

741

742

743

744

745

746

747

748

752

756

Xiongtao Zhou, Jie He, Lanyu Chen, Jingyu Li, Haojing Chen, Víctor Gutiérrez-Basulto, Jeff Z Pan, and Hanjie Chen. 2024. Miceval: Unveiling multimodal chain of thought's quality via image description and reasoning steps. *arXiv preprint arXiv:2410.14668*.

# A Related Work

In recent years, evaluating the reasoning capabilities of LLMs has become a pivotal research focus within the field of natural language processing(Chang et al., 2024). Existing evaluation methodologies can be broadly categorized into two approaches: outcome-oriented and processoriented assessments.

Outcome-oriented evaluations primarily emphasize the accuracy of the model's final output. Prominent benchmarks in this category include Super-GLUE (Wang et al., 2019), MMLU (Hendrycks et al., 2020b), and BIG-bench (Suzgun et al., 2022). These benchmarks encompass a wide array of tasks, ranging from language comprehension to domainspecific question answering, thereby standardizing the performance assessment of LLMs. However, such methods often overlook the interpretability and rationality of the model's reasoning process, particularly in complex problem-solving scenarios where the significance of intermediate steps is substantially undervalued.

To address these limitations, process-oriented evaluation methodologies have been introduced(Zheng et al., 2024; Jiang et al., 2025). The CoT reasoning framework (Wei et al., 2022) exemplifies this approach by explicitly guiding models to generate intermediate reasoning steps, thereby enhancing performance in mathematical and logical tasks. Subsequent studies, such as Auto-CoT (Zhang et al., 2022), Tree-of-Thought (Yao et al., 2023a), and ReAct (Yao et al., 2023b), have further augmented the flexibility and diversity of reasoning pathways.

Furthermore, the evaluation of multidisciplinary reasoning capabilities has become a focal point in current research. Researchers have developed various assessment benchmarks and methodologies to comprehensively measure the reasoning abilities of LLMs across different academic disciplines. For instance, the Advanced Reasoning Benchmark (ARB) is a comprehensive reasoning benchmark that spans multiple domains, including mathematics, physics, biology, chemistry, and law, designed to evaluate the performance of LLMs in complex reasoning tasks (Sawada et al., 2023). Multi-LogiEval is a dataset that provides an integrated evaluation of LLMs' multi-step logical reasoning abilities, covering multiple types of logic such as propositional logic, first-order logic, and non-monotonic logic (Patel et al., 2024).

Additionally, large reasoning models tend to exhibit overthinking behavior during chain-ofthought reasoning, where excessively long and unnecessary reasoning steps are generated even for simple or ill-posed problems (Sui et al., 2025; Chen et al., 2024; Fan et al., 2025; Pu et al., 2025). This phenomenon is often attributed to the models' lack of proper termination mechanisms and insufficient confidence estimation, leading to inefficient inference and degraded accuracy. It has been observed that the issue becomes more pronounced when essential premises are missing from the input (Fan et al., 2025). To address this, several approaches have been proposed, including the introduction of new reasoning efficiency metrics and self-training strategies that encourage concise reasoning (Chen et al., 2024), as well as dynamic early-exit mechanisms that halt inference when sufficient confidence is reached (Yang et al., 2025b). Additionally, path scoring methods have been developed to prefer less redundant reasoning paths, thereby improving performance while reducing computational cost (Cuadron et al., 2025).

# **B** More experimental results

The comparative analysis of Tables 4 and Tables 5 highlights the trade-offs between reasoning quality and efficiency across various disciplines. In mathematics, top-performing models, such as Qwen3-235b-a22b and Qwq-plus, achieve high quality of reflection and precision but require a large number of tokens. In contrast, Grok-3-mini-beta strikes a balance between conciseness and accuracy, achieving a precision of 88.2% in physics and a recall of 84.9% in chemistry while using fewer tokens. Furthermore, smaller distilled variants, like Deepseek-r1-distill-qwen-1.5b, demonstrate significant limitations in domain-specific reasoning, particularly in physics and chemistry, where both recall and precision fall below 50%.

As shown in Table 6. On simple tasks, Grok-3mini-beta demonstrates efficient and focused reasoning, producing only 0.54 thoughts and consuming 1,574.26 tokens. In contrast, when tackling difficult questions, larger models such as Qwen3-235b-a22b and Ernie-x1-turbo-32k generate over

		Recall		P	recision		Reflec	ction Qua	lity		Tokens	
Model name	Chemistry	Physics	Math	Chemistry	Physics	Math	Chemistry	Physics	Math	Chemistry	Physics	Math
C-3-7-sonnet	81.83%	82.97%	76.97%	84.92%	86.94%	87.16%	87.89%	90.39%	95.25%	744.21	836.63	1494.49
Ds-r1-distill-qwen-1.5b	34.43%	46.11%	70.37%	47.91%	58.86%	80.78%	59.37%	70.18%	89.20%	2807.87	4236.72	4279.11
Ds-r1-distill-qwen-7b	54.11%	59.42%	75.95%	75.55%	78.93%	89.49%	80.16%	87.28%	95.58%	2800.44	2742.55	2984.15
Ds-r1-distill-qwen-14b	56.21%	65.16%	75.54%	77.07%	85.76%	90.95%	80.56%	90.89%	96.36%	2729.47	2654.41	2731.53
Ds-r1-distill-qwen-32b	57.75%	62.24%	76.39%	68.02%	78.67%	90.09%	72.12%	83.86%	94.46%	3525.06	3442.9	3595.26
Ds-r1	81.39%	81.41%	78.59%	83.34%	89.63%	94.09%	87.79%	93.32%	98.05%	4082.93	3331.74	4245.25
Es-x1-turbo-32k	82.25%	83.44%	78.82%	83.92%	90.00%	94.00%	87.95%	92.75%	97.00%	4762.48	4588.33	4783.4
G-3-mini-beta	84.86%	83.34%	72.45%	82.79%	88.22%	89.30%	86.04%	91.69%	96.53%	1742.1	1690.57	2546.22
Q3-235b-a22b	84.88%	87.27%	84.37%	83.39%	87.36%	92.18%	90.29%	94.55%	98.05%	5022.55	4592.91	5636.6
Qwq-plus	78.38%	81.92%	80.73%	78.87%	86.42%	92.80%	84.76%	92.29%	98.11%	5795.2	5460.61	6202.25
Glm-z1-air	79.93%	81.09%	78.66%	78.37%	84.49%	88.66%	84.51%	91.84%	96.97%	3992.9	4013.77	3250.19

Table 4: **Comparative Performance of Models in Different Category.** C-3-7-sonnet: claude 3.7 sonnet; Ds-r1-distill-qwen-1.5b: deepseek-r1-distill-qwen-15b; Ds-r1-distill-qwen-7b: deepseek-r1-distill-qwen-14b; Ds-r1-distill-qwen-14b: deepseek-r1-distill-qwen-32b; Ds-r1-distill-qwen-32b: deepseek-r1-distill-qwen-7b; Ds-r1: deepseek-reasoner; Es-x1-turbo-32k: ernie-x1-turbo-32k; G-3-mini-beta: grok-3-mini-beta; Q3-235b-a22b: qwen3-235b-a22b.

	Thought Num		Efficiency			Use	eful Toker	IS	Reflection Tokens			
Model name	Chemistry	Physics	Math	Chemistry	Physics	Math	Chemistry	Physics	Math	Chemistry	Physics	Math
C-3-7-sonnet	0.16	0.23	0.58	0.47	0.52	0.5	339.85	411.6	695.93	404.35	425.03	798.55
Ds-r1-distill-qwen-1.5b	8.2	8.65	6.35	0.31	0.35	0.53	842.08	1243.98	2033.23	1965.79	2992.74	2245.88
Ds-r1-distill-qwen-7b	9.76	6.13	4.3	0.48	0.48	0.61	1367.99	1243.05	1831.2	1432.45	1499.5	1152.94
Ds-r1-distill-qwen-14b	8.24	5.75	4	0.49	0.51	0.63	1270.29	1226.93	1745.42	1459.18	1427.48	986.11
Ds-r1-distill-qwen-32b	12.49	8.43	6.25	0.45	0.47	0.62	1535.14	1480.06	2147.1	1989.91	1962.84	1448.16
Ds-r1	13.13	6.96	6.97	0.49	0.45	0.55	2087.08	1521.81	2404.29	1995.86	1809.93	1840.97
Es-x1-turbo-32k	17.14	11.15	8.62	0.48	0.44	0.5	2315.51	2031.88	2444.63	2446.98	2556.44	2338.77
G-3-mini-beta	0.45	0.24	0.52	0.59	0.59	0.71	1049.76	985.93	1738.09	692.34	704.65	808.13
Q3-235b-a22b	16.95	11.21	11.62	0.49	0.44	0.46	2547.58	2087.76	3007.54	2474.97	2505.15	2629.06
Qwq-plus	30.98	18.24	17.44	0.46	0.42	0.46	2664.93	2351.91	3209.83	3130.26	3108.69	2992.42
Glm-z1-air	14.11	6.99	7.61	0.48	0.46	0.50	1887.16	1483.43	2174.12	2105.78	1766.76	1839.66

Table 5: **Comparative Performance of Models in Different Category.** C-3-7-sonnet: claude 3.7 sonnet; Ds-r1-distill-qwen-1.5b: deepseek-r1-distill-qwen-15b; Ds-r1-distill-qwen-7b: deepseek-r1-distill-qwen-14b; Ds-r1-distill-qwen-14b: deepseek-r1-distill-qwen-32b; Ds-r1-distill-qwen-32b: deepseek-r1-distill-qwen-7b; Ds-r1: deepseek-reasoner; Es-x1-turbo-32k: ernie-x1-turbo-32k; G-3-mini-beta: grok-3-mini-beta; Q3-235b-a22b: qwen3-235b-a22b.

	Thoug	ht Num	Useful	Tokens	Reflection Tokens		
Model name	Simple	Difficult	Simple	Difficult	Simple	Difficult	
Claude-3-7-sonnet	0.17	0.39	340.05	553.55	333.19	662.46	
Deepseek-r1-distill-qwen-1.5b	4.10	11.91	818.28	1720.40	1331.65	3605.56	
Deepseek-r1-distill-qwen-7b	3.05	11.04	825.64	2003.10	749.63	2056.54	
Deepseek-r1-distill-qwen-14b	2.86	9.70	777.32	1931.06	737.21	1954.99	
Deepseek-r1-distill-qwen-32b	4.68	14.18	960.86	2325.93	1113.45	2615.53	
Deepseek-r1	4.26	14.11	950.72	2877.72	1107.63	2661.91	
Ernie-x1-turbo-32k	6.75	18.78	1144.62	3302.73	1534.70	3411.17	
Grok-3-mini-beta	0.22	0.54	765.60	1574.26	476.67	968.99	
Qwen3-235b-a22b	7.93	18.80	1245.45	3656.38	1573.24	3472.42	
Qwq-plus	11.09	34.21	1328.11	3971.12	1961.34	4226.88	
Glm-z1-air	4.27	15.36	876.64	2677.43	1054.72	2756.20	

Table 6: Effect of question difficulty on other efficiency measures.

3,600 tokens on average while achieving high reflection quality, reaching 92.16% and 94.00% respectively. However, this increase in quality comes
with reduced efficiency. For example, Qwq-plus
achieves only 44.58% efficiency due to its high
reflection token count after the answer, totalling
4,226.88 tokens.

# 814 C More Qualitative Examples

# 815 C.1 Example of Computational Efficiency

816 In Figure 7, we illustrate an example of efficient reasoning by an LRM.

Question	Three pencils and a jumbo eraser cost $\1.24$ . Five pencils and a jumbo eraser cost $\1.82$ . No prices include tax. In cents, what is							
Answer	29							
<b>Deeps</b>	eek R1							
Solution 1 Solution 1 Control 1 Cont	kay, so I need to figure out how much a pencil costs, right? The problem ays three pencils and a jumbo eraser cost \$1.24, and five pencils and a jumbo caser cost \$1.82. They want the cost of a pencil in cents. Hmm, let me think.\ nFirst, maybe I should set up some equations. Let me call the cost of one encil \"p\" and the cost of a jumbo eraser \"e\". Both are in dollars for now, ut since the answer needs to be in cents, I'll convert later.\n\nSo, the first quation would be $3p + e = 1.24$ . That's three pencils plus an eraser equals 1.24. The second equation is $5p + e = 1.82$ . Five pencils and the same eraser ost \$1.82.\n\nHmm, okay, so I have two equations here with two variables. hat means I can solve them using substitution or elimination. Let's see. If I ubtract the first equation from the second, maybe that will eliminate the raser?\n\nLet me try that. So, subtract $(3p + e)$ from $(5p + e)$ :\n\n(5p + e) - 3p + e) = 1.82 - 1.24, which is 0.58. So, $2p = 0.58$ . Then, divide both sides y 2: p = 0.58 / 2 = 0.29. Wait, that's in dollars. So, converting to cents, that's 9 cents per pencil.							
Solution 2	Let me check if that makes sense.\n\nIf a pencil is 29 cents, then three pencils re $3 * 29 = 87$ cents. Then, the eraser would be $1.24 - 0.87 = 0.37$ dollars, thich is 37 cents. Then, check the second equation: five pencils would be $5 * 9 = 145$ cents, which is \$1.45. Adding the eraser, $1.45 + 0.37 = 1.82$ , which matches the given amount. So, that works out.							
Solution 3 is n <sup>1</sup>	/ait, but let me just confirm my steps again to make sure I didn't make a histake. Starting with the two equations: $\ln(n1) 3p + e = 124$ cents (since \$1.24 124 cents) $\ln(2) 5p + e = 182$ cents $\ln(nSubtract equation 1 from equation 2:\ln(5p + e) - (3p + e) = 182 - 124 \ln N which simplifies to:\ln(2p = 58 \ln nSo), p 58 / 2 = 29 cents. Yep, same result. So, that seems right.$							
Solution 4 Control 4 Conta	Iternatively, if I didn't convert to cents initially, the equations would be in ollars: $n^3p + e = 1.24$ , $p = 0.29$ , $n^3p + e = 1.24$ , $p = 0.29$ dollars, which is 29 cents. Same answer, ither way works. $n^3p$ , the cost of a pencil is 29 cents. Both methods onfirm that. The eraser ends up being 37 cents, which checks out in both quations. So, I think that's the correct answer. I don't see any mistakes in the alculations, so I'm pretty confident. $n^*F$ inal Answer** $n$ The cost of a encil is $b$ .							
	Efficiency = 278/819=33.9%							



# C.2 Error Example

We present an example of an LRM output error in Figure 8.



Figure 8: Failure Case Example in the Deepseek-r1-distill-qwen-32b Answering Process.

# **D** Evaluation Prompts

# Key Steps Extraction Prompt

You are an expert system that gives you a question and a corresponding answer, please list in detail the key reasoning steps from the question to the answer, make sure that the reasoning steps are clear and complete, and include all possible solutions. You should pretend not to know the basic truth answer beforehand.

Input: Question:

{question}

Answer:

{answer}

Output requirements:

1. Only include the essential key steps, and don't output unnecessary words

2. For each solution, record:

-logical\_conclusion: The set of each key step of the solution, from Step 1 all the way to the answer 3. A problem may contain more than one way of reasoning, so make sure you don't miss any possible solutions.

4. Important: Output only JSON array with no additional information.

5. Don't add useless words to the process

Here is the json output format: ## Output format [ {{

```
"solution1": {{
```

```
"logical_conclusion": ["step1:","step2:","step3:",...]
}}
}
```

## **Recall Evaluation Prompt**

# Task Overview

You are an expert system for verifying solutions to text-based problems. Your task is to match the ground truth middle steps with the provided solution.

# INPUT FORMAT:

1. Problem: The original question/task

2. A Solution of a model

3. Ground Truth: Essential steps required for a correct answer

## # MATCHING PROCESS:

You need to evaluate each ground truth middle step against the solution, following these criteria:

## Match Criteria:

- \*\*Exact Match or Equivalent Logical Step\*\*: A ground truth step is considered \*\*Matched\*\* if:

- It appears exactly in the solution \*\*OR\*\*

- The same logical reasoning or idea is clearly expressed, even if wording or format differs.

- \*\*Numerical and Conceptual Consistency\*\*: All key numbers, equations, or transformations should align conceptually with the ground truth.

- \*\*Step-by-Step Evaluation\*\*: Each ground truth step must be assessed individually and sequentially.

- \*\*Final Answer Check\*\*: Determine if the overall reasoning process leads to the correct final answer.

# **# OUTPUT FORMAT:**

```
[
  {{
    "step_index": <integer>,
    "judgment": "Matched" | "Unmatched",
    "correct_answer": "true" | "false"
 }}
]
```

**# ADDITIONAL RULES:** 

\*\*Strict JSON Output\*\*: Output only the JSON array with no additional text or explanations.
 \*\*No Omitted Steps\*\*: Every step in 'Ground Truth' must be evaluated.

# EDGE CASE HANDLING:

- If a step is conceptually equivalent but reworded, it is still considered \*\*Matched\*\*.

- If numerical transformations are equivalent (e.g., same formula in a different form), it is \*\*Matched\*\*.

- If the reasoning process does not lead to the correct final answer, "correct\_answer": "false".

Here is the problem, answer, solution, and the ground truth middle steps:

[Problem]

{question}
[Answer]
{answer}
[Solution]
{solution}
[Ground Truth Information]
{gt\_annotation}

# Precision Evaluation Prompt

#### # Task Overview

Given a solution with multiple reasoning steps for a text problem, reformat it into well-structured steps and evaluate their correctness.

# Step 1: Reformatting the Solution

- Convert the unstructured solution into distinct reasoning steps while:
- Preserving all original content and order
- Not adding new interpretations
- Not omitting any steps

# ## Step Types

- 1. Logical Inference Steps
- Contains exactly one logical deduction
- Must produce a new derived conclusion
- Cannot be just a summary or observation

2. Background Information Steps

- External knowledge or question context
- No inference process involved

## Step Requirements

- Each step must be atomic (one conclusion per step)
- No content duplication across steps
- Initial analysis counts as background information
- Final answer determination counts as logical inference

# Step 2: Evaluating Correctness Evaluate each step against:

## Ground Truth Matching

For logical inferences:

- Conclusion must EXACTLY match or be DIRECTLY entailed by ground truth

## Reasonableness Check (if no direct match) Step must:

- Premises must not contradict any ground truth or correct answer
- Logic is valid
- Conclusion must not contradict any ground truth
- Conclusion must support or be neutral to correct answer

```
## Judgement Categories
- "Match": Aligns with ground truth
- "Reasonable": Valid but not in ground truth
- "Wrong": Invalid or contradictory
- "N/A": For background information steps
# Output Requirements
1. The output format MUST be in valid JSON format without ANY other content.
2. For highly repetitive patterns, output it as a single step.
3. Output maximum 35 steps. Always include the final step that contains the answer.
4. correct_answer: Whether the whole reasoning process produces the right answer.
Here is the json output format:
## Output Format
Ε
  {{
     "step_type": "logical inference|background information",
     "premise": "Evidence",
     "conclusion": "Step result",
     "judgment": "Match|Reasonable|Wrong|N/A"
     "correct_answer": "true|false"
  }}
٦
Here is the problem, and the solution that needs to be reformatted to steps:
[Problem]
{question}
```

```
{question}
[Solution]
{solution}
[Correct Answer]
{answer}
[Ground Truth Information]
{gt_annotation}
```

# Model Output Reformat Prompt

I will present you with a solution to a problem. Unfortunately, the solution lacks proper paragraphing, making it hard to read. Your task is to improve readability by reformatting the solution into well-structured paragraphs. Follow these specific guidelines:

\* Insert \n\n for paragraph breaks within the original solution. Do \*\*NOT\*\* alter any content of the original solution (the only exception is for itemized lists; see below).

- Each paragraph should represent a distinct, concise reasoning step that logically advances the solution.

- Reasoning steps can include case discussions, formula simplifications, or formula derivations. Each of these should be treated as an individual reasoning step and paragraphed accordingly.

- If an introductory analysis exists in the original solution, treat it as an initial reasoning step and place it as the first paragraph.

- Do \*\*NOT\*\* place any formulas in their own separate paragraphs; instead, include them within the same paragraph as the preceding text to form a cohesive reasoning step.

\* For any itemized lists (ordered or unordered), convert them into a written format, such as "First/Second/Third." This is the \*\*ONLY\*\* content modification allowed.

\* Avoid making paragraphs too lengthy, as long paragraphs might contain multiple reasoning steps that should be paragraphed separately.

\* Disregard the accuracy of the solution content. Do \*\*NOT\*\* alter any of the original solution's content; focus solely on structuring it into logical, readable paragraphs.

\* Reply with the reformatted solution directly.

Here is the problem, and the solution that needs to be reformatted: [Problem] {problem} [Solution]

{response}

# 827

# First Correct Answer Extraction Prompt

The following is a problem and a solution (split into paragraphs, enclosed with tags and indexed from 0):

[Problem]

{problem}

[Correct Answer] {answer}

[Solution]

{tagged response}

Your task is to review and critique the solution paragraph by paragraph. Once you identify an correct answer in a paragraph, return the index of the paragraph where the earliest correct answer occurs. Otherwise, return the index of -1 (which typically denotes "not found"). Please put your final answer (i.e., the index) in boxed.

# **Reflection Quality Prompt**

Here's a refined prompt that improves clarity and structure:

# Task Evaluate reflection steps in a problem-solving solutions, where reflections are selfcorrections or reconsiderations of previous statements.

# Reflection Step Identification

Reflections typically begin with phrases like:

- "But xxx"

- "Alternatively, xxx"

- "Maybe I should"
- "Let me double-check"
- "Wait xxx"
- "Perhaps xxx"

It will throw an doubt of its previously reached conclusion or raise a new thought.

# Evaluation Criteria

Correct reflections must:

1. Reach accurate conclusions aligned with ground truth

2. Use new insights to find the mistake of the previous conclusion or verify its correctness.

```
Invalid reflections include:
1. Repetition - Restating previous content or method without new insights
2. Wrong Conclusion - Reaching incorrect conclusions vs ground truth
3. Incompleteness - Proposing but not executing new analysis methods
4. Other - Additional error types
# Input Format
[Problem]
{question}
[Think Content]
{think_content}
[Ground Truth]
{gt_annotation}
# Output Requirements
1. The output format must be in valid JSON format without any other content.
2. Output maximum 30 reflection steps.
Here is the json output format:
## Output Format
Ε
  {{
     "conclusion": "One-sentence summary of reflection outcome",
     "judgment": "Correct|Wrong",
    "error_type": "N/A|Repetition|Wrong Conclusion|Incompleteness|Other"
  }}
]
# Rules 1. Preserve original content and order
2. No new interpretations
3. Include ALL reflection steps
4. Empty list if no reflections found
5. Direct JSON output without any other output
```