

# Can Large Language Models Detect Errors in Long Chain-of-Thought Reasoning?

Yancheng He<sup>1\*</sup>, Shilong Li<sup>1\*</sup>, Jiaheng Liu<sup>1,2\*†</sup>, Weixun Wang<sup>1\*</sup>, Xingyuan Bu<sup>1</sup>, Ge Zhang<sup>3</sup>,  
Zhongyuan Peng<sup>1</sup>, Zhaoxiang Zhang<sup>4</sup>, Zhicheng Zheng<sup>1</sup>, Wenbo Su<sup>1</sup>, Bo Zheng<sup>1</sup>

<sup>1</sup>Alibaba Group, China <sup>2</sup>Nanjing University, China <sup>3</sup>M-A-P, <sup>4</sup>CASIA  
heyancheng.hyc@taobao.com, liujiaheng@nju.edu.cn

## Abstract

Recently, o1-like models have drawn significant attention, where these models produce the long Chain-of-Thought (CoT) reasoning steps to improve the reasoning abilities of existing Large Language Models (LLMs). In this paper, to understand the qualities of these long CoTs and measure the critique abilities of existing LLMs on these long CoTs, we introduce the **DeltaBench**<sup>1</sup> including the generated long CoTs from different o1-like models (e.g., QwQ, DeepSeek-R1) for different reasoning tasks (e.g., Math, Code, General Reasoning), to measure the ability to **Detect Errors in Long CoT Reasoning**. Based on DeltaBench, we first perform fine-grained analysis of the generated long CoTs to discover the effectiveness and efficiency of different o1-like models. Then, we conduct extensive evaluations of existing process reward models (PRMs) and critic models to detect the errors of each annotated process, which aims to investigate the boundaries and limitations of existing PRMs and critic models. Finally, we hope that DeltaBench could guide developers to better understand the long CoT reasoning abilities of their models.

## 1 Introduction

Large Language Models (LLMs) have witnessed remarkable progress in recent years (Liu et al., 2025a; Li et al., 2024a; Liu et al., 2024; Zhang et al., 2024a; Li et al., 2024b; He et al., 2024b; Liu et al., 2025b). Among these advancements, o1-like LLMs have emerged to greatly improve reasoning capabilities by generating long Chain-of-Thought (CoT) reasoning steps (OpenAI, 2024; Guo et al., 2025).

However, despite the growing popularity of o1-like models and their long CoT reasoning approaches, *systematic evaluation of the quality and*

*effectiveness of the generated reasoning chains is not well investigated*, which poses a significant challenge in understanding the capabilities and limitations of these models (Chen et al., 2024b; Yeo et al., 2025; Zhang et al., 2025). Besides, as LLMs continue to evolve, further improving their performance has become increasingly challenging. One of the key areas of focus in this regard is the development of critique ability, which measures the capability to provide detailed analysis, constructive suggestions, and refinement feedback suggestions for solutions generated by other models or even themselves (Lan et al., 2024; Lin et al., 2024; Tan et al., 2024; Song et al., 2025; Zheng et al., 2024a). However, *evaluating the critique abilities of existing LLMs (e.g., critic models or process reward models) on the long CoT reasoning steps has not been explored*.

Therefore, to address the aforementioned challenges, we introduce **DeltaBench**, the first dataset to analyze the qualities of the long CoTs generated by o1-like models and evaluate the critique abilities to **Detect Error in Long CoT Reasoning** of existing critic models and PRMs in Figure 1. Specifically, in DeltaBench, we first collect a diverse collection of long CoTs generated by various o1-like models (i.e., QwQ, DeepSeek-R1, and Gemini-2.0 Flash Thinking) across different reasoning tasks such as **Math**, **Programming**, **PCB** (physics, chemistry and biology), and **General Reasoning**. Then, we divide each long CoT into different sections, where each section denotes an independent subtask. After that, each section is annotated with corresponding labels (e.g., reasoning usefulness, reasoning correctness, and reflection).

Based on DeltaBench, we first conduct a fine-grained analysis on efficiency of the generated long CoTs from different o1-like models, and have the following interesting findings:

- **Fundamental errors (e.g., calculation errors,**

\* First four authors contributed equally.

† Corresponding Author.

<sup>1</sup>The dataset is available at <https://github.com/OpenStellarTeam/DeltaBench>.

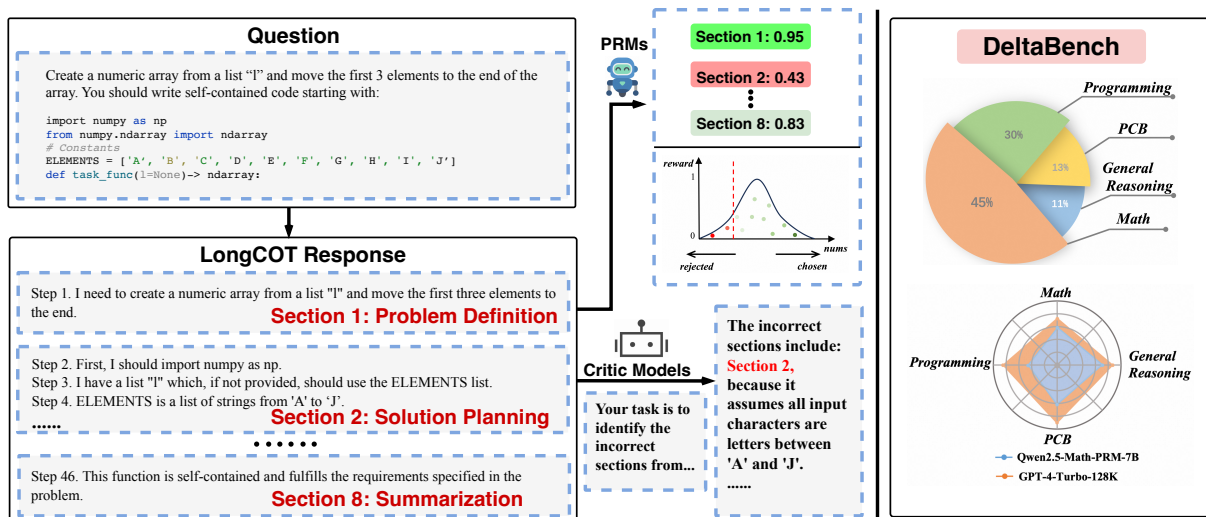


Figure 1: Illustration of the evaluation process for critic models and PRMs for DeltaBench.

**syntax errors and format errors) are usually existed in existing o1-like models.** For example, such errors account for approximately 25% and 23% in QwQ-32B-Preview and Gemini 2.0 Flash Thinking, respectively.

- **The proportion of effective reflection is still very low.** Note that the “effective reflection” denotes this reflection leads to the right answer. For example, on average, approximately 67.8% of the reflections in the collected long CoT responses are useless.
- **The long CoT reasoning process is very redundant for existing o1-like models.** For example, on average, 27% of the reasoning sections in the collected long CoT response are redundant.

After that, we evaluate the critique abilities on LLMs prompted as critic models and PRMs and draw the following insightful observations:

- **For existing LLMs and PRMs, the ability to effectively identify errors in long CoT reasoning is very limited.** For example, the top-performing model in DeltaBench, GPT-4-turbo-128k, achieves an F1-score of only 40.8%.
- **o1-like models (e.g., o1-mini, o1-preview) do not show any advantage over non-o1-like models on the critique abilities.** For example, the results of o1-preview are lower than the results of GPT-4o-mini.
- **The self-critique abilities of existing o1-like models are generally weaker than their critique abilities on other o1-like models.** For

example, DeepSeek-R1 exhibits a 36% reduction in self-critique performance compared to the critique performance of other o1-like models.

## 2 Related Works

**LLM Critic.** As many works have begun to explore the test-time scaling (OpenAI, 2024; Snell et al., 2024), the critique capabilities of LLMs have drawn great interests. For example, CriticBench (Luo et al., 2023; Zheng et al., 2024b) uses LLMs to generate critiques and binary verdicts for input solutions, measuring accuracy by comparing these verdicts to ground truth labels. CriticEval (Lan et al., 2024) evaluates both feedback and correction qualities. Additionally, many works have explored the LLMs’ self-critique for improving reasoning (Tyen et al., 2023; Stechly et al., 2024). For example, Huang et al. (2023) have highlighted the challenges in self-correction without external feedback. ProCo (Wu et al., 2024) facilitates self-correction and iterative verification processes.

## 3 DeltaBench

In this section, we detail the construction of the DeltaBench dataset, developed to assess a model’s capacity to identify and locate errors in long CoT reasoning processes. DeltaBench comprises 1,236 samples across diverse domains, including **Math**, **Programming**, **PCB** (physics, chemistry and biology), and **General Reasoning**. Each sample encompasses a problem, its corresponding long CoT solution, and comprehensive human annotations. Specifically, the long CoT is divided into sections, and each section includes the following tags:

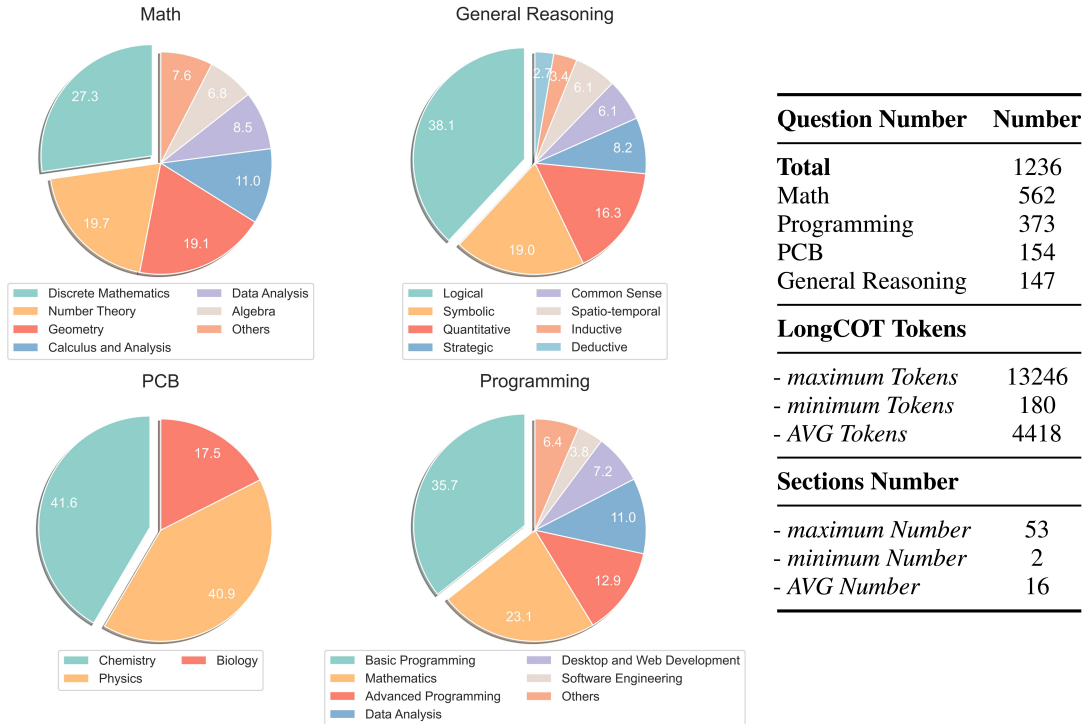


Figure 2: Left: Overview of DeltaBench with pie charts showing the distribution of questions in Math, General Reasoning, PCB (Physics, Chemistry and Biology), and Programming. Right: Statistics of DeltaBench.

- **Strategy Shift:** whether this section introduces a new method or strategy attempt. If a new strategy is introduced, the specific step is annotated.
- **Reasoning Usefulness:** whether the reasoning in this section is useful. If the process of section can help to lead to the right answer, it considered as useful.
- **Reasoning Correctness:** whether this section contains any errors. If an error is present, additional error-related fields are annotated, including the first step number at which the error occurs, explanation and correction.
- **Reflection Efficiency:** whether this section contains reflection and whether the reflection is correct. If reflection is present, the step at which the reflection begins is annotated.
- **Clustering and Deduplication:** Queries are first converted into vector representations using the NV-Embed-v2<sup>2</sup> embedding model. Then, similarity scores are computed between each pair of queries to identify and eliminate duplicates using a predefined threshold. The non-duplicate queries are clustered using DBSCAN (Ester et al., 1996), resulting in 17,510 unique queries.
- **Difficulty Filtering:** For each query, multiple models<sup>3</sup> were employed to generate solutions, and difficulty labels were assigned based on the accuracy of the answers produced by these models. Following this, uniform sampling was carried out according to these difficulty labels to ensure a balanced distribution of difficulties.
- **Subcategory Sampling:** For each query, GPT-4o is used to classify it into a subcategory. The queries are then uniformly sampled based on these subcategories to ensure diversity.

### 3.1 Dataset Construction

**Query collection.** We extract queries from diverse open-source datasets. Detailed data sources are listed in Appendix B.1. The domains include math, programming, physics, chemistry, biology, and general reasoning, which comprise 48 subcategories. To ensure the dataset’s diversity and balance, we employ a multi-step process:

**Data Preprocessing.** We observe low-quality queries exist in open-source datasets. To address

<sup>2</sup><https://huggingface.co/nvidia/NV-Embed-v2>

<sup>3</sup>GPT-4o (OpenAI, 2023), Meta-Llama-3-70B-Instruct (Dubey et al., 2024), Meta-Llama-3-8B-Instruct (Dubey et al., 2024), Qwen2.5-72B-Instruct (Team, 2024), Qwen2.5-32B-Instruct (Team, 2024), and DeepSeek-67B-chat (DeepSeek-AI, 2024).

this issue, we employed GPT-4 and rule-based filtering strategies to identify and remove these low-quality queries. We have recorded encountered issues. Specific details are shown in Appendix B.4.

**Long CoT Generation.** We generate long CoT solutions using several open-source o1-like models, such as QwQ-32B-Preview, DeepSeek-R1, and Gemini 2.0 Flash Thinking, with random sampling to enhance diversity. This method ensures a wide range of reasoning processes and captures potential errors models may produce in real-world scenarios, enabling a robust evaluation of error detection capabilities in long CoT reasoning.

**Section Split.** Previous approaches typically divide solutions into steps. However, long CoT responses often contain numerous steps, which significantly increases the difficulty of human annotation, and many of which are either overly granular or lack meaningful contribution to the overall reasoning process. To address this issue, we segment each long CoT response into multiple sections, each representing an independent sub-task, aligning more closely with human cognitive patterns. Specifically, we use the delimiter "\n\n" to partition the model's response into steps first. Then, we employ GPT-4 to identify the start and end steps of each section and generate a brief summary of the content within each section. This approach not only facilitates manual annotation but also enhances the accuracy of the model's segmentation process. The details are provided in Appendix B.5.

In addition, we visualize the changing distribution of section types in the model's reasoning process. See Appendix D.1 for detailed analysis.

### 3.2 Correctness Assessment

Before manual annotation, we employ automated methods to assess the correctness of the long CoT results. Domain-specific techniques are used to identify potentially erroneous outputs in Appendix B.3 and the evaluation accuracy of each domain and the details are shown in Appendix B.6. This process ensures that the data provided for manual annotation are likely to contain errors, which enhances the annotation efficiency.

### 3.3 Human Annotation

The data annotation process aims to evaluate the reasoning process of each long CoT response systematically. Each section is assessed for **strategy**

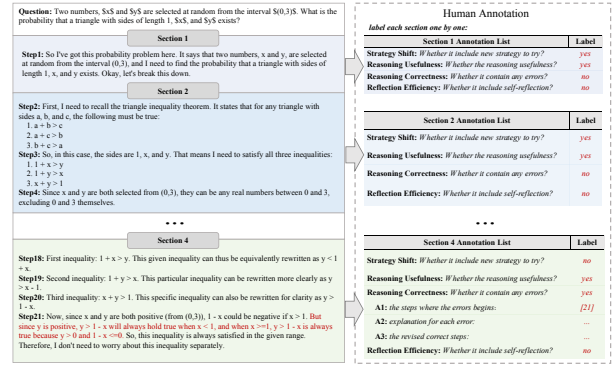


Figure 3: Illustration of the annotation workflow and an example of human annotation.

Benchmark	Source	Long CoT	Granularity
JudgeBench	MMLU-Pro, LiveBench, LiveCodeBench	×	Sample-Level
CriticBench	GSM8K, CSQA, BIGBench, HumanEval, etc	×	Sample-Level
CriticEval	GSM8K, HumanEval, ChatArena, etc.	×	Sample-Level
ProcessBench	GSM8K, MATH, OlympiadBench, Omni-MATH	×	Step-Level
<b>DeltaBench</b>	AIME, BigCodeBench, KOR-bench, GPQA, etc	✓	Section-Level

Table 1: Comparisons between different benchmarks.

**shift, reasoning usefulness, reasoning correctness, and reflection efficiency.** The annotation of whether strategy shift and reflection occurred is to help analyze o1's thinking pattern. The annotation of Reasoning usefulness and error identification is to better analyze and evaluate the performance of system II thinking and further evaluate the critique ability of other models for these problems.

To ensure high-quality annotations, we recruit Master's and Ph.D. graduates from various disciplines and collaborate with specialized suppliers (See Appendix B.2 for more details on the annotation and quality control processes).

### 3.4 Dataset Statistics

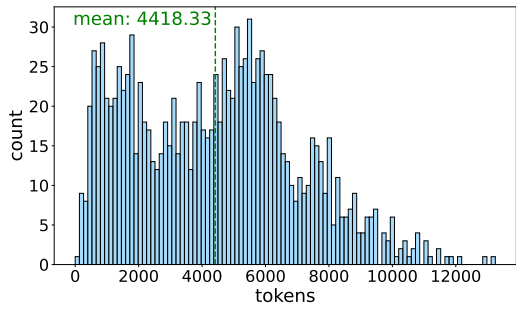
**DeltaBench** contains 1,236 carefully curated samples. These samples are distributed across five major domains and 48 subcategories. The dataset ensures a balance of question types and difficulty levels, incorporating a rich set of long CoT responses.

Figure 4 shows the distribution of both the length and the number of sections of long CoTs. The distribution is relatively balanced overall, enabling a comprehensive evaluation of the performance of PRMs or critic models across a range of different lengths. Additionally, detailed statistics on the category distribution are provided in Appendix B.7.

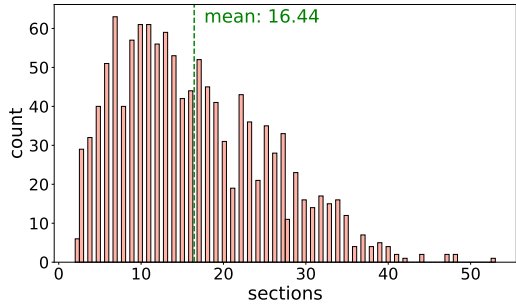
### 3.5 Comparison to other benchmarks

In Table 1, DeltaBench has the following features: (1) We focus on difficult questions, providing a





(a) Distribution of long CoT Length.



(b) Distribution of the number of long CoT sections.

Figure 4: Distribution of long CoT characteristics.

more challenging critical evaluation; (2) We utilize long CoT responses, enabling the assessment of a model’s ability to identify errors within complex reasoning processes; (3) We evaluate the model’s capability to identify all errors in the reasoning process, rather than just the first error or a binary classification of correctness on sample level, which can provide a fine-grained analysis of long CoTs.

## 4 Analysis

### 4.1 Error Analysis of o1-like Models.

**Error Type Lists** We classify the errors that occur during the system II thinking process into 8 major aspects and 23 specific error types based on the manual annotations, including understanding errors, reasoning errors, reflection errors, summary errors, etc. For detailed information about the error categories, see Appendix C.

**What Are the Most Common Errors Across Domains?** To analyze the characteristics of error distribution in different domains, we performed a uniform sampling of the data based on the model, the domain, and the query difficulty. Figure 5 shows the error distribution across different domains, where a detailed analysis is provided in Appendix ??, here are some key findings:

- **Math:** The most frequent error type is *Reasoning Error* (25.3%), followed by *Understanding Error* (15.7%) and *Calculation Error* (15.4%).

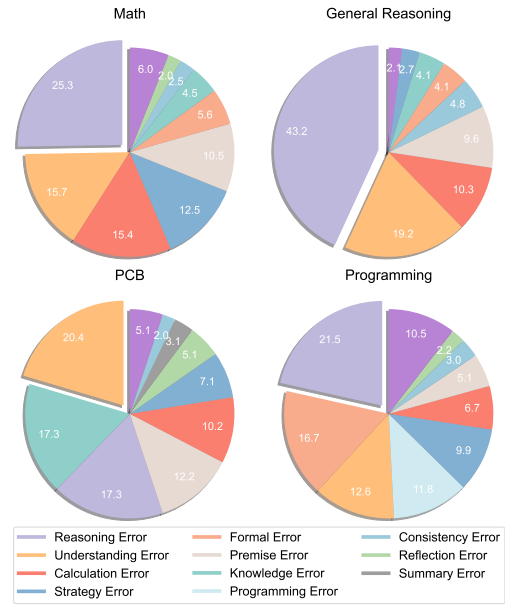


Figure 5: Distribution of Error Types Across Domains.

*Reasoning Error* (25.3%), followed by *Understanding Error* (15.7%) and *Calculation Error* (15.4%). This indicates that while the models often struggle with logical reasoning and problem understanding, low-level computational mistakes also remain a significant issue.

- **Programming:** *Reasoning Error* (21.5%) is the most common, followed by *Formal Error* (16.7%) and *Understanding Error* (12.6%). The high frequency of *Formal Error* and *Programming Error* (11.8%) underscores the models’ struggles with code-specific details and implementation.
- **PCB:** The dominant error types are *Understanding Error* (20.4%) and *Knowledge Error* (17.3%), closely followed by *Reasoning Error* (17.3%). This suggests that the main challenge for current models in the fields of physics, chemistry and biology is to understand field-specific concepts and accurately apply relevant knowledge.
- **General Reasoning:** *Reasoning Error* is the most prevalent, accounting for 43%, followed by comprehension errors, accounting for 19%, showing that logical reasoning is the primary bottleneck.

**What Are the Model-Specific Error Patterns?** We also analyzed errors specific to individual models, providing further insights into model weak-

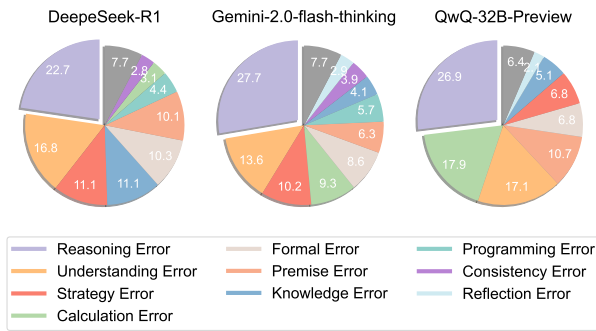


Figure 6: Distribution of Error Types Across Models.

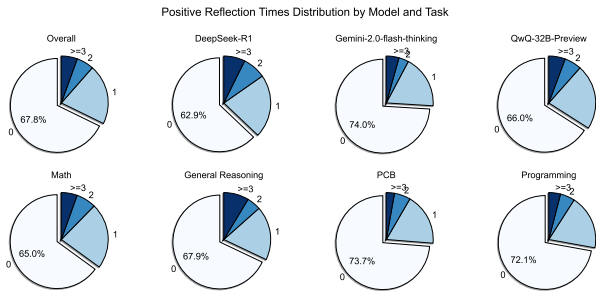


Figure 7: Distribution of Effective Reflection Times by Model and Task on a sample level. The segments within each pie chart represent how many times effective reflection occurs in one sample, with segment '0' indicating there is no effective reflection.

nesses, as illustrated in Figure 6. The error distributions reveal distinct patterns for each model, highlighting their unique strengths and areas for improvement. Due to space constraints, we focus here on the key findings from the most commonly used models, with a comprehensive analysis of all models provided in Appendix ??.

- **DeepSeek-R1** exhibits its most pronounced weakness in *Reasoning Errors* (22.7%), indicating challenges in constructing coherent and accurate logical reasoning paths. However, it demonstrates relative strength in handling fundamental tasks, with minimal *Calculation Errors* (3.1%) and *Programming Errors* (4.4%).
- **QwQ-32B-Preview** excels at identifying correct problem-solving approaches. However, its effectiveness is significantly hindered by deficiencies in handling finer details, particularly in *Calculation Errors* (17.9%).

## 4.2 Reflection Analysis of o1-like Models.

**Statistics.** We also conduct a analysis of the total number of reflections and the proportion of effective reflections in the long CoT output of all ques-

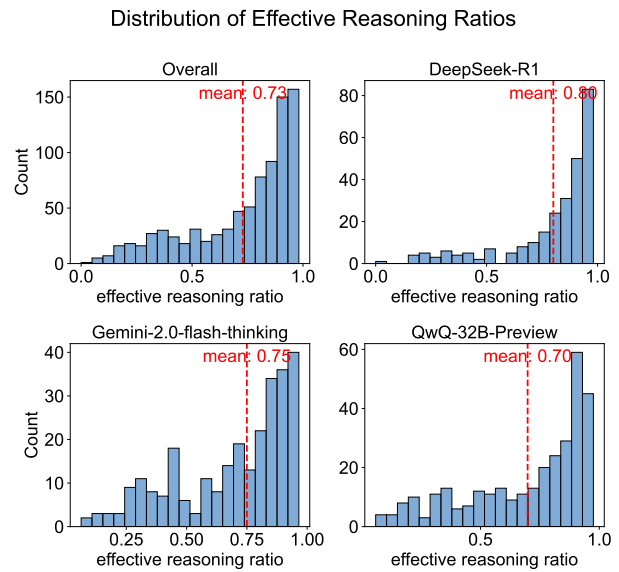


Figure 8: Distribution of Effective Reasoning Ratios.

tions (including questions answered correctly and incorrectly by the model).

## How Effective Are Model Reflections Across Different Models and Domains?

We classify samples with reflections based on the number of valid reflections to evaluate the ability to produce valid reflections. Specifically, we label samples as 0 if no valid reflections occur, and 1, 2, or >=3 for samples with one, two, or three and more valid reflections, respectively(all statistical analyses were performed under strictly controlled conditions, ensuring uniform sampling and balanced tasks for a fair comparison). In Figure 7, DeepSeek-R1 exhibits the highest proportion of effective reflections, and the models show a notably higher rate of effective reflections in the math domain.

### 4.3 Effective Reasoning of o1-like Models.

**Statistics.** Human annotators evaluate the usefulness of the reasoning in each section, enabling us to calculate the proportion of valid reasoning in each response. As illustrated in Figure 8, each graph shows the distribution of effective reasoning ratios for a particular model. The red dashed line in each graph indicates the average effective reasoning ratio.

## What Proportion of Reasoning in Long COT Responses is Effective?

On average, only 73% of the reasoning in the collected long COT responses is useful, highlighting significant redundancy issues. Among the models analyzed, *QwQ-32B-Preview* exhibited the lowest proportion of effective

reasoning at 70%, while *DeepSeek-R1* achieved a notably higher proportion compared to the others, demonstrating superior reasoning efficiency.

#### 4.4 Results on DeltaBench

**Evaluation Metrics.** We employ **recall**, **precision**, and **macro-F1 score** for error sections as evaluation metrics. For the PRMs, we utilize an outlier detection technique based on the Z-Score to make predictions. This method was chosen because threshold-based prediction methods determined from other step-level datasets, such as those used in ProcessBench (Zheng et al., 2024a), may not be reliable due to significant differences in dataset distributions, particularly as DeltaBench focuses on long CoT. Outlier detection helps to avoid this bias. The threshold  $t$  for determining the correctness of a section is defined as:  $t = \mu - \sigma$ , where  $\mu$  is the mean of the rewards distribution across the dataset, and  $\sigma$  is the standard deviation. Sections falling below  $t$  are predicted as error sections. For critic models, all erroneous sections within a long CoT are prompted to be identified. Given that error sections constitute a smaller proportion than correct sections across the dataset, we use macro-F1 to mitigate the potential impact of the imbalance between positive and negative sections. Macro-F1 independently calculates the F1 score for each sample and then takes the average, providing a more balanced evaluation metric when dealing with class imbalance.

**Baseline Models.** For the **PRMs**, we select the following models: Qwen2.5-Math-PRM-7B<sup>4</sup>, Qwen2.5-Math-PRM-72B<sup>5</sup>, Llama3.1-8B-PRM-Deepseek-Data<sup>6</sup>, Llama3.1-8B-PRM-Mistral-Data<sup>7</sup>, Skywork-o1-Open-PRM-Qwen-2.5-1.5B<sup>8</sup>, and Skywork-o1-Open-PRM-Qwen-2.5-7B<sup>9</sup>. We select a group of the most advanced open-source and closed-source LLMs to serve as **critic models** for evaluation, which includes various GPT-4 (OpenAI, 2023) variants (such as GPT-4-turbo-128K, GPT-4o-mini, GPT-4o), the Gemini model (Reid et al., 2024)(Gemini-1.5-pro), several Qwen models (Team, 2024) (such as Qwen2.5-32B-Instruct and Qwen2.5-14B-Instruct), and o1 models (OpenAI, 2024) (o1-preview-0912, o1-mini-0912).

<sup>4</sup>Qwen/Qwen2.5-Math-PRM-7B

<sup>5</sup>Qwen/Qwen2.5-Math-PRM-72B

<sup>6</sup>RLHFlow/Llama3.1-8B-PRM-Deepseek-Data

<sup>7</sup>RLHFlow/Llama3.1-8B-PRM-Mistral-Data

<sup>8</sup>Skywork/Skywork-o1-Open-PRM-Qwen-2.5-1.5B

<sup>9</sup>Skywork/Skywork-o1-Open-PRM-Qwen-2.5-7B

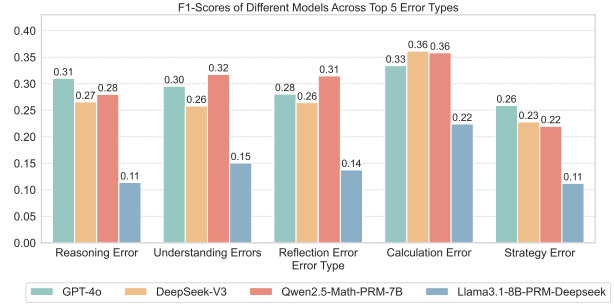


Figure 9: Results of different LLMs on top-5 errors.

#### 4.4.1 Main Results

In Table 2, we provide the results of different LLMs on DeltaBench. For PRMs, we have the following observations: (1). Existing PRMs usually achieve low performance, which indicates that existing PRMs cannot identify the errors in long CoTs effectively and it is necessary to improve the performance of PRMs. (2). Larger PRMs do not lead to better performance. For example, the Qwen2.5-Math-PRM-72B is inferior to Qwen2.5-Math-PRM-7B. For critic models, we have the following findings: (1) GPT-4-turbo-128k archives the best critique results, which is better than other models (e.g., GPT-4o) a lot in DeltaBench. (2) For o1-like models (e.g., DeepSeek-R1, o1-mini, o1-preview), we observe that the results of these models are not superior to non-o1-like models, with the performance of o1-preview is even lower than Qwen2.5-32B-Instruct. A detailed analysis of underperforming models is provided in Appendix E.

#### 4.4.2 Further Analysis

**Performance Analysis Across Different Error Types.** Figure 9 shows the performance of different models on the five most common error types. In terms of error types, most models demonstrate the highest accuracy in recognizing calculation errors. Conversely, the recognition of strategy errors is generally the weakest. In terms of models, there is significant variation in the ability of individual models to recognize different error types. For instance, DeepSeek-V3 achieves an F1 of 36% on calculation errors but only 23% on strategy errors. Meanwhile, Llama3.1-8B-PRM-Deepseek performs poorly, with an F1 score of 22% on calculation errors, and shows a significant decline in performance across the other four error types. This highlights the limited generalization capabilities of most models when recognizing various error types.

Model	Overall			Math	Code	PCB	General
	Recall	Precision	F1	F1	F1	F1	F1
<i>Process Reward Models (PRMs)</i>							
Qwen2.5-Math-PRM-7B	<b>30.30</b>	<b>34.96</b>	<b>29.22</b>	<b>29.64</b>	<b>23.76</b>	31.09	34.19
Qwen2.5-Math-PRM-72B	<u>28.16</u>	<u>29.37</u>	<u>26.38</u>	<u>24.16</u>	<u>22.02</u>	<b>31.14</b>	<b>35.83</b>
Llama3.1-8B-PRM-Deepseek-Data	11.7	15.59	12.02	12.28	10.95	16.76	12.59
Llama3.1-8B-PRM-Mistral-Data	9.64	11.21	9.45	9.40	10.72	13.43	12.40
Skywork-o1-Qwen-2.5-1.5B	3.32	3.84	3.07	1.30	6.66	5.43	7.87
Skywork-o1-Qwen-2.5-7B	2.49	2.22	2.17	0.78	6.28	6.02	3.11
<i>LLM as Critic Models</i>							
GPT-4-turbo-128k	<b>57.19</b>	<b>37.35</b>	<b>40.76</b>	<b>37.56</b>	<b>43.06</b>	45.54	42.17
GPT-4o-mini	<u>49.88</u>	<u>35.37</u>	<u>37.82</u>	<u>33.26</u>	<u>37.95</u>	<b>45.98</b>	<b>46.39</b>
GPT-4o	36.52	32.48	30.85	28.61	28.53	39.25	36.50
Gemini-1.5-pro	35.51	30.32	29.59	26.56	28.20	40.13	33.66
DeepSeek-V3	32.33	28.13	27.33	27.04	27.73	27.35	27.45
Llama-3.1-70B-Instruct	32.22	28.85	27.67	21.49	32.13	28.45	39.18
Qwen2.5-32B-Instruct	30.12	28.63	26.73	22.34	31.37	33.78	24.37
DeepSeek-R1	29.20	32.66	28.43	24.17	29.28	34.78	35.87
o1-preview	27.92	30.59	26.97	22.19	28.09	33.11	35.94
Qwen2.5-14B-Instruct	26.64	27.27	24.73	21.51	29.05	29.98	20.59
Llama-3.1-8B-Instruct	25.71	28.01	24.91	18.12	32.17	27.30	29.93
o1-mini	22.90	22.90	19.89	16.71	21.70	20.37	26.94
Qwen2.5-7B-Instruct	21.99	19.61	18.63	11.61	25.92	29.85	15.18
DeepSeek-R1-Distill-Qwen-32B	17.19	18.65	16.28	13.02	23.55	15.05	11.56
DeepSeek-R1-Distill-Qwen-14B	12.81	14.54	12.55	9.40	18.36	10.44	12.01

Table 2: Experimental results of PRMs and critic models on DeltaBench. Bold indicates the best results within the same group of models, while the underline indicates the second best.

Model	HitRate@k - Avg(%)		
	k = 1	k = 3	k = 5
Qwen2.5-Math-PRM-7B	<b>49.15</b>	<b>69.14</b>	<b>83.14</b>
Qwen2.5-Math-PRM-72B	<u>41.13</u>	<u>62.70</u>	<u>75.73</u>
Llama3.1-8B-PRM-Deepseek-Data	12.63	48.62	69.78
Llama3.1-8B-PRM-Mistral-Data	8.99	42.97	65.33
Skywork-o1-Open-PRM-Qwen-2.5-1.5B	31.90	53.82	69.23
Skywork-o1-Open-PRM-Qwen-2.5-7B	31.58	52.59	69.16

Table 3: Results of HitRate@k. Bold and underlined results indicate the best and the second best.

**Analysis on HitRate evaluation for PRMs.** To better measure the ability of PRMs to identify erroneous sections in long CoTs, we use HitRate@k to evaluate PRMs. Specifically, within a sample, we rank the sections in ascending order based on the rewards given by the PRM, select the smallest k sections, and calculate the recall rate for the erroneous sections among them. Specifically, we define the sorted sections as  $S = \{s_1, s_2, \dots, s_n\}$ , with  $E$  being the set of erroneous sections. We select the top k sections, denoted as  $S_k = \{s_1, s_2, \dots, s_k\}$ . The HitRate@k is calculated as:

$$\text{HitRate@k} = \frac{|S_k \cap E|}{\min(k, |E|)} \quad (1)$$

In this formula,  $|S_k \cap E|$  indicates the number of erroneous sections identified among the top k sections. This metric reflects the ability of PRMs to

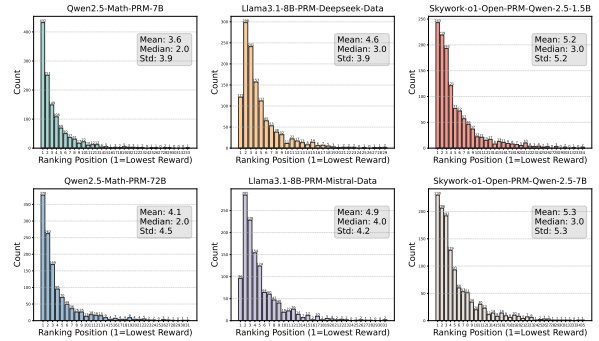


Figure 10: Ranking of rewards for the first incorrect section for different PRMs.

effectively identify erroneous sections within the top k candidate sections. In Table 3, the relative performance rankings among different PRMs are quite similar to the results in Table 2. Additionally, we observe that for  $k = 3$  and  $k = 5$ , the performance differences between various PRMs are not particularly significant. However, when  $k = 1$ , the Qwen2.5-Math-PRM-7B shows a clear performance advantage. Figure 10 illustrates the ranking ability of different PRMs for the first incorrect section within the sample, which is generally consistent with the performance evaluation results of HitRate@k.



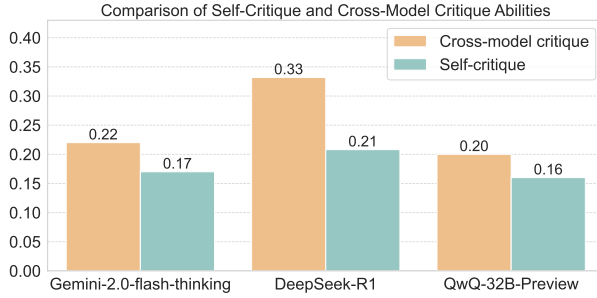


Figure 11: F1 Score Comparison of Self-Critique and Cross-Model Critique Abilities for Different Models.

### Self-Critique Capabilities Lag Behind Cross-Model Critique Abilities.

We randomly sample queries based on domains and models that generate the long CoT output, followed by a statistical analysis of the model’s performance in evaluating its own outputs as well as those of other models. In Figure 11, Gemini-2.0-flash-thinking, DeepSeek-R1, and QwQ-32B-Preview show lower self-critique scores compared to their cross-model critique scores, indicating a prevalent deficiency in self-critic abilities. Notably, DeepSeek-R1 exhibits the largest discrepancy, with a 36% decrease in self-evaluation compared to evaluations of other models. This suggests models’ self-critic abilities remain underdeveloped.

Model	Overall	Math	Code	PCB	General
Qwen2.5-32B-Instruct	26.73	22.34	31.37	33.78	24.37
Qwen2.5-14B-Instruct	24.73	21.51	29.05	29.98	20.59
Qwen2.5-7B-Instruct	18.63	11.61	25.92	29.85	15.18
Qwen2.5-3B-Instruct	16.29	13.81	18.23	25.39	11.31
Qwen2.5-1.5B-Instruct	7.49	5.24	9.51	11.80	6.49
Qwen2.5-0.5B-Instruct	3.79	2.30	5.25	5.24	4.31

Table 4: Experimental results for different Qwen models across various domains.

### Scaling effects on error detection capabilities.

To further investigate the impact of model size on error detection performance, we conduct a comprehensive evaluation across six model variants (ranging from Qwen2.5-0.5B-Instruct to Qwen2.5-32B-Instruct). This systematic examination of how model size influences error detection capabilities. The detailed results, measured by Macro-F1 scores, are presented in Table 4. Our analysis reveals two key findings: (1) Performance Scaling: Error detection capability demonstrates substantial improvement (+22.94 points, from 3.79 to 26.73) with increasing model size, indicating that larger models exhibit significantly stronger overall error detection performance; (2) Capability Leaps: We observe

notable performance leaps between the 3B-to-7B and 7B-to-14B model size transitions, representing critical points where model capabilities undergo significant enhancement.

## 5 Conclusion

In this paper, we have provided a comprehensive evaluation benchmark called DeltaBench to investigate the limitations of existing o1-like models based on the generated long CoTs and measure the critique qualities of existing LLMs. Based on DeltaBench, we discuss the specific error analysis of o1-like models and provide a detailed analysis on critic models and PRMs, where many interesting findings are provided. Finally, we hope our DeltaBench can not only find the limitations of o1-like models, but also provide guidance to further improve these reasoning models.

## 6 Acknowledgement

This work was supported in part by the Jiangsu Science and Technology Major Project (BG2024031) and Nanjing University AI & AI for Science Funding (2024300540).

## 7 Limitations

While DeltaBench offers a comprehensive evaluation for long CoT reasoning and critique abilities, it has some limitations. First, the construction and annotation of the dataset involve high costs, making it challenging to scale to a larger volume of data. Second, although we established a rigorous human annotation mechanism, the process may still introduce subjective biases. Third, as a static benchmark, DeltaBench may not fully capture real-time advancements. Addressing these limitations will be a key focus of our future work.

## References

- Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. 2024a. [Alphamath almost zero: Process supervision without process](#). *Preprint*, arXiv:2405.03553.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuqi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2024b. [Do not think that much for 2+3=? on the overthinking of o1-like llms](#). *ArXiv*, abs/2412.21187.
- Bytedance-Seed-Foundation-Code-Team Yao Cheng, Jianfeng Chen, Jie Chen, Li Chen, Liyu Chen, Wentao Chen, Zhengyu Chen, Shijie Geng, Aoyan Li,

Bowen Li, Bowen Li, Linyi Li, Boyi Liu, Jerry Liu, Kaibo Liu, Qi Liu, Shukai Liu, Si-Han Liu, Tianyi Liu, Tingkai Liu, Yongfei Liu, Rui Long, Jing Mai, Guanghan Ning, Zhongyan Peng, Kai Shen, Jiahao Su, Jing Su, Tao Sun, Yifan Sun, Yu Tao, Guoyin Wang, Siwei Wang, Xuwu Wang, Yite Wang, Zihan Wang, Jinxiang Xia, Liang Xiang, Xianzhong Xiao, Yongsheng Xiao, Chenguang Xi, Shulin Xin, Jingjing Xu, Shi-Bo Xu, Hongxia Yang, Jack Yang, Yingxiang Yang, Jian-Ming Yuan, Jun Zhang, Yufeng Zhang, Yuyu Zhang, Shen Zheng, He Zhu, and Ming Zhu. 2024. [Fullstack bench: Evaluating llms as full stack coders](#). *ArXiv*, abs/2412.00535.

DeepSeek-AI. 2024. [Deepseek llm: Scaling open-source language models with longtermism](#). *arXiv preprint arXiv:2401.02954*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Alonsoius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Milon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Celebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao

Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baeviski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Cvin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard,

- Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhota, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsim-poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiao Cheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. The llama 3 herd of models. *arXiv preprint arXiv: 2407.21783*.
- Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, page 226–231. AAAI Press.
- Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan, Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. 2024. *Omni-math: A universal olympiad level mathematic benchmark for large language models*. *ArXiv*, abs/2410.07985.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. *Reasoning with language model is planning with world model*. *Preprint*, arXiv:2305.14992.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024a. *Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems*. In *Annual Meeting of the Association for Computational Linguistics*.
- Yancheng He, Shilong Li, Jiaheng Liu, Yingshui Tan, Weixun Wang, Hui Huang, Xingyuan Bu, Hangyu Guo, Chengwei Hu, Boren Zheng, Zhuoran Lin, Xuepeng Liu, Dekai Sun, Shirong Lin, Zhicheng Zheng, Xiaoyong Zhu, Wenbo Su, and Bo Zheng. 2024b. *Chinese simpleqa: A chinese factuality evaluation for large language models*. *Preprint*, arXiv:2411.07140.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fan-jia Yan, Tianjun Zhang, Sida I. Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. *Livecodebench: Holistic and contamination free evaluation of large language models for code*. *ArXiv*, abs/2403.07974.
- Tian Lan, Wenwei Zhang, Chen Xu, Heyan Huang, Dahua Lin, Kai Chen, and Xian-Ling Mao. 2024. Criticeval: Evaluating large-scale language model as critic. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Shilong Li, Yancheng He, Hangyu Guo, Xingyuan Bu, Ge Bai, Jie Liu, Jiaheng Liu, Xingwei Qu, Yangguang Li, Wanli Ouyang, et al. 2024a. Graphreader: Building graph-based agent to enhance long-context abilities of large language models. *arXiv preprint arXiv:2406.14550*.
- Shilong Li, Yancheng He, Hui Huang, Xingyuan Bu, Jiaheng Liu, Hangyu Guo, Weixun Wang, Jihao Gu,



- Wenbo Su, and Bo Zheng. 2024b. 2d-dpo: Scaling direct preference optimization with 2-dimensional supervision. *arXiv preprint arXiv:2410.19720*.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. [Let’s verify step by step](#). *ArXiv*, abs/2305.20050.
- Bill Yuchen Lin, Ronan Le Bras, Kyle Richardson, Ashish Sabharwal, Radha Poovendran, Peter Clark, and Yejin Choi. 2025. [Zebralogic: On the scaling limits of llms for logical reasoning](#).
- Zicheng Lin, Zhibin Gou, Tian Liang, Ruilin Luo, Haowei Liu, and Yujiu Yang. 2024. Criticbench: Benchmarking llms for critique-correct reasoning. *arXiv preprint arXiv:2402.14809*.
- Jiaheng Liu, Ken Deng, Congnan Liu, Jian Yang, Shukai Liu, He Zhu, Peng Zhao, Linzheng Chai, Yanan Wu, Ke Jin, et al. 2024. M2rc-eval: Massively multilingual repository-level code completion evaluation. *arXiv preprint arXiv:2410.21157*.
- Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Wang, Chenchen Zhang, Ge Zhang, Jiebin Zhang, et al. 2025a. A comprehensive survey on long context language modeling. *arXiv preprint arXiv:2503.17407*.
- Wei Liu, Yancheng He, Hui Huang, Chengwei Hu, Jiaheng Liu, Shilong Li, Wenbo Su, and Bo Zheng. 2025b. Air: Complex instruction generation via automatic iterative refinement. *arXiv preprint arXiv:2502.17787*.
- Liangchen Luo, Zi Lin, Yinxiao Liu, Lei Shu, Yun Zhu, Jingbo Shang, and Lei Meng. 2023. Critique ability of large language models. *arXiv preprint arXiv:2310.04815*.
- Kaijing Ma, Xinrun Du, Yunran Wang, Haoran Zhang, Zhoufutu Wen, Xingwei Qu, Jian Yang, Jiaheng Liu, Minghao Liu, Xiang Yue, Wenhao Huang, and Ge Zhang. 2024. [Kor-bench: Benchmarking language models on knowledge-orthogonal reasoning tasks](#). *ArXiv*, abs/2410.06526.
- OpenAI. 2023. Gpt-4 technical report. *PREPRINT*.
- OpenAI. 2024. [Learning to reason with llms, september 2024](#).
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem W. Ayoub, Megha Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas, Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal, Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren Sezener, Luke Vilnis, Oscar Chang, Nobuyuki Morioka, George Tucker, Ce Zheng, Oliver Woodman, Nithya Attaluri, Tomás Kociský, Evgenii Eltyshev, Xi Chen, Timothy Chung, Vittorio Selo, Siddhartha Brahma, Petko Georgiev, Ambrose Slone, Zhenkai Zhu, James Lottes, Siyuan Qiao, Ben Caine, Sebastian Riedel, Alex Tomala, Martin Chadwick, J Christopher Love, Peter Choy, Sid Mittal, Neil Houlsby, Yunhao Tang, Matthew Lamm, Libin Bai, Qiao Zhang, Luheng He, Yong Cheng, Peter Humphreys, Yujia Li, Sergey Brin, Albin Cassirer, Ying-Qi Miao, Lukás Zilka, Taylor Tobin, Kelvin Xu, Lev Proleev, Daniel Sohn, Alberto Magni, Lisa Anne Hendricks, Isabel Gao, Santiago Ontañón, Oskar Bunyan, Nathan Byrd, Abhanshu Sharma, Biao Zhang, Mario Pinto, Rishika Sinha, Harsh Mehta, Dawei Jia, Sergi Caelles, Albert Webson, Alex Morris, Becca Roelofs, Yifan Ding, Robin Strudel, Xuehan Xiong, Marvin Ritter, Mostafa Dehghani, Rahma Chaabouni, Abhijit Karmarkar, Guangda Lai, Fabian Mentzer, Bibo Xu, YaGuang Li, Yujing Zhang, Tom Le Paine, Alex Goldin, Behnam Neyshabur, Kate Baumli, Anselm Levskaya, Michael Laskin, Wenhao Jia, Jack W. Rae, Kefan Xiao, Antoine He, Skye Giordano, Lakshman Yagati, Jean-Baptiste Lespiau, Paul Natsev, Sanjay Ganapathy, Fangyu Liu, Danilo Martins, Nanxin Chen, Yunhan Xu, Megan Barnes, Rhys May, Arpi Vezer, Junhyuk Oh, Ken Franko, Sophie Bridgers, Ruizhe Zhao, Boxi Wu, Basil Mustafa, Sean Sechrist, Emilio Parisotto, Thanumalayan Sankaranarayanan Pillai, Chris Larkin, Chenjie Gu, Christina Sorokin, Maxim Krikun, Alexey Guseynov, Jessica Landon, Romina Datta, Alexander Pritzel, Phoebe Thacker, Fan Yang, Kevin Hui, A.E. Hauth, Chih-Kuan Yeh, David Barker, Justin Mao-Jones, Sophia Austin, Hannah Sheahan, Parker Schuh, James Svensson, Rohan Jain, Vinay Venkatesh Ramasesh, Anton Briukhov, Da-Woon Chung, Tamara von Glehn, Christina Butterfield, Priya Jhakra, Matt Wiethoff, Justin Frye, Jordan Grimstad, Beer Changpinyo, Charline Le Lan, Anna Bortsova, Yonghui Wu, Paul Voigtlaender, Tara N. Sainath, Charlotte Smith, Will Hawkins, Kris Cao, James Besley, Srivatsan Srinivasan, Mark Omernick, Colin Gaffney, Gabriela de Castro Surita, Ryan Burnell, Bogdan Damoc, Junwhan Ahn, Andrew Brock, Mantas Pajarskas, Anastasia Petrushkina, Seb Noury, Lorenzo Blanco, Kevin Swersky, Arun Ahuja, Thi Avrahami, Vedant Misra, Raoul de Liedekerke, Mariko Iinuma, Alex Polozov, Sarah York, George van den Driessche, Paul Michel, Justin Chiu, Rory Blevins, Zach Gleicher, Adrià Recasens, Alban Rustemi, Elena Gribovskaya, Aurko Roy, Wiktor Gworek, Sébastien M. R. Arnold, Lisa Lee, James Lee-Thorp, Marcello Maggioni, Enrique Piqueras, Kartikeya Badola, Sharad Vikram, Lucas Gonzalez, Anirudh Baddepudi, Evan Senter, Jacob Devlin, James Qin, Michael Azzam, Maja Trebacz, Martin



Polacek, Kashyap Krishnakumar, Shuo yiin Chang, Matthew Tung, Ivo Penchev, Rishabh Joshi, Kate Olszewska, Carrie Muir, Mateo Wirth, Ale Jakse Hartman, Joshua Newlan, Sheleem Kashem, Vijay Bolina, Elahe Dabir, Joost R. van Amersfoort, Zafarali Ahmed, James Cobon-Kerr, Aishwarya B Kamath, Arnar Mar Hrafnkelsson, Le Hou, Ian Mackinnon, Alexandre Frechette, Eric Noland, Xiance Si, Emanuel Taropa, Dong Li, Phil Crone, Anmol Gulati, S'ebastien Cevey, Jonas Adler, Ada Ma, David Silver, Simon Tokumine, Richard Powell, Stephan Lee, Michael B. Chang, Samer Hassan, Diana Mincu, Antoine Yang, Nir Levine, Jenny Brennan, Mingqiu Wang, Sarah Hodgkinson, Jeffrey Zhao, Josh Lipschultz, Aedan Pope, Michael B. Chang, Cheng Li, Laurent El Shafey, Michela Paganini, Sholto Douglas, Bernd Bohnet, Fabio Pardo, Seth Odoom, Mihaela Rosca, Cicero Nogueira dos Santos, Kedar Soparkar, Arthur Guez, Tom Hudson, Steven Hansen, Chulayuth Asawaroengchai, Ravichandra Addanki, Tianhe Yu, Wojciech Stokowiec, Mina Khan, Justin Gilmer, Jaehoon Lee, Carrie Grimes Bostock, Keran Rong, Jonathan Caton, Pedram Pejman, Filip Pavetic, Geoff Brown, Vivek Sharma, Mario Luvci'c, Rajkumar Samuel, Josip Djolonga, Amol Mandhane, Lars Lowe Sjosund, Elena Buchatskaya, Elspeth White, Natalie Clay, Jiepu Jiang, Hyeontaek Lim, Ross Hemsley, Jane Labanowski, Nicola De Cao, David Steiner, Sayed Hadi Hashemi, Jacob Austin, Anita Gergely, Tim Blyth, Joe Stanton, Kaushik Shivakumar, Aditya Siddhant, Anders Andreassen, Carlos L. Araya, Nikhil Sethi, Rakesh Shivanna, Steven Hand, Ankur Bapna, Ali Khodaei, Antoine Miech, Garrett Tanzer, Andy Swing, Shantanu Thakoor, Zhufeng Pan, Zachary Nado, Stephanie Winkler, Dian Yu, Mohammad Saleh, Lorenzo Maggiore, Iain Barr, Minh Giang, Thais Kagohara, Ivo Danihelka, Amit Marathe, Vladimir Feinberg, Mohamed Elhawaty, Nimesh Ghelani, Dan Horgan, Helen Miller, Lexi Walker, Richard Tanburn, Mukarram Tariq, Disha Shrivastava, Fei Xia, Chung-Cheng Chiu, Zoe Ashwood, Khuslen Baatarsukh, Sina Samangoeei, Fred Alcober, Axel Stjerngren, Paul Komarek, Katerina Tsihlas, Anudhyan Boral, Ramona Comanescu, Jeremy Chen, Ruiibo Liu, Dawn Bloxwich, Charlie Chen, Yanhua Sun, Fangxi aoyu Feng, Matthew Mauger, Xerxes Dotiwalla, Vincent Hellendoorn, Michael Sharman, Ivy Zheng, Krishna Haridasan, Gabriel Barth-Maron, Craig Swanson, Dominika Rogozi'nska, Alek Andreev, Paul Kishan Rubenstein, Ruoxin Sang, Dan Hurt, Gamaleldin Elsayed, Ren shen Wang, Dave Lacey, Anastasija Ili'c, Yao Zhao, Woohyun Han, Lora Aroyo, Chimezie Iwuanyanwu, Vitaly Nikolaev, Balaji Lakshminarayanan, Sadegh Jazayeri, Raphael Lopez Kaufman, Mani Varadarajan, Chetan Tekur, Doug Fritz, Misha Khalman, David Reitter, Kingshuk Dasgupta, Shourya Sarcar, T. Ornduff, Javier Snaider, Fantine Huot, Johnson Jia, Rupert Kemp, Nejc Trdin, Anitha Vijayakumar, Lucy Kim, Christof Angermueller, Li Lao, Tianqi Liu, Haibin Zhang, David Engel, Somer Greene, Anais White, Jessica Austin, Lilly Taylor, Shereen Ashraf, Dangyi Liu, Maria Georgaki, Irene Cai, Yana Kulizhskaya, Sonam Goenka, Bren-

nan Saeta, Kiran Vodrahalli, Christian Frank, Dario de Cesare, Brona Robenek, Harry Richardson, Mahmoud Alnahlawi, Christopher Yew, Priya Ponnappalli, Marco Tagliasacchi, Alex Korchemniy, Yelin Kim, Dinghua Li, Bill Rosgen, Kyle Levin, Jeremy Wiesner, Praseem Banzal, Praveen Srinivasan, Hongkun Yu, cCauglar Unlu, David Reid, Zora Tung, Daniel F. Finchelstein, Ravin Kumar, Andre Elisseeff, Jin Huang, Ming Zhang, Rui Zhu, Ricardo Aguilar, Mai Gim'enez, Jiawei Xia, Olivier Dousse, Willi Gierke, Soheil Hassas Yeganeh, Damion Yates, Komal Jalan, Lu Li, Eri Latorre-Chimoto, Duc Dung Nguyen, Ken Durden, Praveen Kallakuri, Yaxin Liu, Matthew Johnson, Tomy Tsai, Alice Talbert, Jasmine Liu, Alexander Neitz, Chen Elkind, Marco Selvi, Mimi Jasarevic, Livio Baldini Soares, Albert Cui, Pidong Wang, Alek Wenjiao Wang, Xinyu Ye, Krystal Kallarackal, Lucia Loher, Hoi Lam, Josef Broder, Daniel Niels Holtmann-Rice, Nina Martin, Bramandya Ramadhana, Daniel Toyama, Mrinal Shukla, Sujoy Basu, Abhi Mohan, Nicholas Fernando, Noah Fiedel, Kim Paterson, Hui Li, Ankush Garg, Jane Park, Donghyun Choi, Diane Wu, Sankalp Singh, Zhishuai Zhang, Amir Globerson, Lily Yu, John Carpenter, Félix de Chaumont Quitry, Carey Radebaugh, Chu-Cheng Lin, Alex Tudor, Prakash Shroff, Drew Garmon, Dayou Du, Neera Vats, Han Lu, Shariq Iqbal, Alexey Yakubovich, Nilesh Tripurani, James Manyika, Ha roon Qureshi, Nan Hua, Christel Ngani, Maria Abi Raad, Hannah Forbes, Anna Bulanova, Jeff Stanway, Mukund Sundararajan, Victor Ungureanu, Colton Bishop, Yunjie Li, Balaji Venkatraman, Bo Li, Chloe Thornton, Salvatore Scellato, Nishesh Gupta, Yicheng Wang, Ian Tenney, Xihui Wu, Ashish Shenoy, Gabriel Carvajal, Diana Gage Wright, Ben Bariach, Zhuyun Xiao, Peter Hawkins, Sid Dalmia, Clément Farabet, Pedro Valenzuela, Quan Yuan, Christopher A. Welty, Ananth Agarwal, Mianna Chen, Wooyeol Kim, Brice Hulse, Nandita Dukkipati, Adam Paszke, Andrew Bolt, Elnaz Davoodi, Kiam Choo, Jennifer Beattie, Jennifer Prendki, Harsha Vashisht, Rebeca Santamaria-Fernandez, Luis C. Cobo, Jarek Wilkiewicz, David Madras, Ali Elqursh, Grant Uy, Kevin Ramirez, Matt Harvey, Tyler Liechty, Heiga Zen, Jeff Seibert, Clara Huiyi Hu, A. Ya. Khorlin, Maigo Le, Asaf Aharoni, Megan Li, Lily Wang, Sandeep Kumar, Alejandro Lince, Norman Casagrande, Jay Hoover, Dalia El Badawy, David Soergel, Denis Vnukov, Matt Miecnikowski, Jiří ima, Anna Koop, Praveen Kumar, Thibault Sellam, Daniel Vlasic, Samira Daruki, Nir Shabat, John Zhang, Guolong Su, Kalpesh Krishna, Jiageng Zhang, Jeremiah Liu, Yi Sun, Evan Palmer, Alireza Ghaffarkhah, Xi Xiong, Victor Cotruta, Michael Fink, Lucas Dixon, Ashwin Sreevatsa, Adrian Goedeckemeyer, Alek Dimitriev, Mohsen Jafari, Remi Crocker, Nicholas Fitzgerald, Aviral Kumar, Sanjay Ghemawat, Ivan Philips, Frederick Liu, Yannie Liang, Rachel Sterneck, Alena Repina, Marcus Wu, Laura Knight, Marin Georgiev, Hyo Lee, Harry Askham, Abhishek Chakladar, Annie Louis, Carl Crous, Hardie Cate, Dessie Petrova, Michael Quinn, Denese Owusu-Afryie, Achintya Singhal, Nan Wei, Solomon Kim, Damien Vin-

- cent, Milad Nasr, Iliia Shumailov, Christopher A. Choquette-Choo, Reiko Tojo, Shawn Lu, Diego de Las Casas, Yuchung Cheng, Tolga Bolukbasi, Katherine Lee, Saaber Fatehi, Rajagopal Ananthanarayanan, Miteyan Patel, Charbel El Kaed, Jing Li, Jakub Sygnowski, Shreyas Rammohan Belle, Zhe Chen, Jaclyn Konzelmann, Siim Pöder, Roopal Garg, Vinod Koverkathu, Adam Brown, Chris Dyer, Rosanne Liu, Azade Nova, Jun Xu, Junwen Bai, Slav Petrov, Demis Hassabis, Koray Kavukcuoglu, Jeffrey Dean, Oriol Vinyals, and Alexandra Chronopoulou. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *ArXiv*, abs/2403.05530.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *ArXiv*, abs/2311.12022.
- Quan Shi, Michael Tang, Karthik Narasimhan, and Shunyu Yao. 2024. [Can language models solve olympiad programming?](#) *ArXiv*, abs/2404.10952.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- Mingyang Song, Zhao yu Su, Xiaoye Qu, Jiawei Zhou, and Yu Cheng. 2025. [Prmbench: A fine-grained and challenging benchmark for process-level reward models](#).
- Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. 2024. On the self-verification limitations of large language models on reasoning and planning tasks. *arXiv preprint arXiv:2402.08115*.
- Sijun Tan, Siyuan Zhuang, Kyle Montgomery, William Y. Tang, Alejandro Cuadron, Chenguang Wang, Raluca Ada Popa, and Ion Stoica. 2024. [Judgebench: A benchmark for evaluating llm-based judges](#). *ArXiv*, abs/2410.12784.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Gladys Tyen, Hassan Mansoor, Peter Chen, Tony Mak, and Victor Cărbune. 2023. Llms cannot find reasoning errors, but can correct them! *arXiv preprint arXiv:2311.08516*.
- Zhenyu Wu, Qingkai Zeng, Zhihan Zhang, Zhaoxuan Tan, Chao Shen, and Meng Jiang. 2024. Large language models can self-correct with minimal effort. *arXiv preprint arXiv:2405.14092*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). *Preprint*, arXiv:2305.10601.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. 2025. [Demystifying long chain-of-thought reasoning in llms](#).
- Alexander Zhang, Marcus Dong, Jiaheng Liu, Wei Zhang, Yejie Wang, Jian Yang, Ge Zhang, Tianyu Liu, Zhongyuan Peng, Yingshui Tan, et al. 2025. [Codecriticbench: A holistic code critique benchmark for large language models](#). *arXiv preprint arXiv:2502.16614*.
- Ge Zhang, Scott Qu, Jiaheng Liu, Chenchen Zhang, Chenghua Lin, Chou Leuang Yu, Danny Pan, Esther Cheng, Jie Liu, Qunshu Lin, Raven Yuan, Tuney Zheng, Wei Pang, Xinrun Du, Yiming Liang, Yinghao Ma, Yizhi Li, Ziyang Ma, Bill Lin, Emmanouil Benetos, Huan Yang, Junting Zhou, Kaijing Ma, Minghao Liu, Morry Niu, Noah Wang, Quehry Que, RuiBo Liu, Sine Liu, Shawn Guo, Soren Gao, Wangchunshu Zhou, Xinyue Zhang, Yizhi Zhou, Yubo Wang, Yuelin Bai, Yuhan Zhang, Yuxiang Zhang, Zenith Wang, Zhenzhu Yang, Zijian Zhao, Jiajun Zhang, Wanli Ouyang, Wenhao Huang, and Wenhao Chen. 2024a. [Map-neo: Highly capable and transparent bilingual large language model series](#). *arXiv preprint arXiv: 2405.19327*.
- Shudan Zhang, Hanlin Zhao, Xiao Liu, Qinkai Zheng, Zehan Qi, Xiaotao Gu, Xiaohan Zhang, Yuxiao Dong, and Jie Tang. 2024b. [Naturalcodebench: Examining coding performance mismatch on humaneval and natural user prompts](#). *ArXiv*, abs/2405.04520.
- Chujie Zheng, Zhenru Zhang, Beichen Zhang, Runji Lin, Keming Lu, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. 2024a. [Processbench: Identifying process errors in mathematical reasoning](#). *ArXiv*, abs/2412.06559.
- Xin Zheng, Jie Lou, Boxi Cao, Xueru Wen, Yuqiu Ji, Hongyu Lin, Yaojie Lu, Xianpei Han, Debing Zhang, and Le Sun. 2024b. [Critic-cot: Boosting the reasoning abilities of large language model via chain-of-thoughts critic](#). *arXiv preprint arXiv:2408.16326*.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied Sanosi Saied, Weizhu Chen, and Nan Duan. 2023. [Agieval: A human-centric benchmark for evaluating foundation models](#). *ArXiv*, abs/2304.06364.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, et al. 2024. [Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions](#). *arXiv preprint arXiv:2406.15877*.

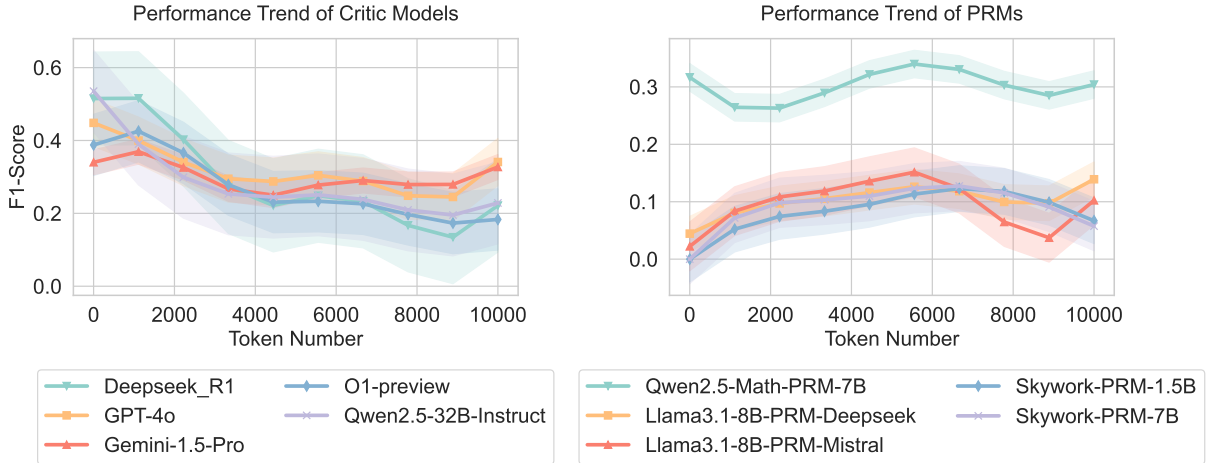


Figure 12: The effect of long CoT length.

## A Effect of Long CoT Length.

In Figure 12, we compare the average F1-Score performance of critic models and PRMs across varying LongCoT token lengths. For critic models, the performance notably declines as token length increases. Initially, models like Deepseek-R1 and GPT-4o exhibit strong performance with shorter sequences (1-3k tokens). However, as token length increases to mid-ranges (4-7k tokens), there is a marked decrease in performance across all models. This trend highlights the growing difficulty for critic models to maintain precision and recall as LongCoT sequences become longer and more complex, likely due to the challenge of evaluating lengthy model outputs. In contrast, PRMs demonstrate greater stability across token lengths, as they evaluate sections sequentially rather than processing the entire output at once. Despite this advantage, PRMs achieve lower overall scores compared to critic models on our evaluation set.

## B Details of Dataset Construction

### B.1 Data Sources

Domain	Source
Math	MATH-500 (Lightman et al., 2023), OlympiadBench (He et al., 2024a), Omni-MATH (Gao et al., 2024), AIME, AMC23, CollegeMath
Programming	CodeForce, BigCodeBench (Zhuo et al., 2024), LiveCodeBench (Jain et al., 2024), FullStackBench (Cheng et al., 2024), NaturalCodeBench (Zhang et al., 2024b), USACO (Shi et al., 2024)
PCB	GPQA (Rein et al., 2023), OlympiadBench (He et al., 2024a), AGIEval (Zhong et al., 2023)
General Reason	ZebraLogicBench (Lin et al., 2025), KOR-Bench (Ma et al., 2024), Geeksforgeeks Puzzles, CS Interview Questions, China Civil Service Exam Questions

Table 5: Query-related data source statistics. Parts without citations come from open internet resources.

Table 5 shows the original data sources from which we extracted high-quality queries and their corresponding solutions, along with additional information (for code data, test cases were extracted).

### B.2 Details on Dataset Annotation

The average cost for annotating each data unit was approximately \$15. The annotation process is organized into three phases, each with specific goals and criteria:

**Initial Assessment.** Annotators initially verify the quality of the question and subsequently evaluate the correctness of the model’s final response.

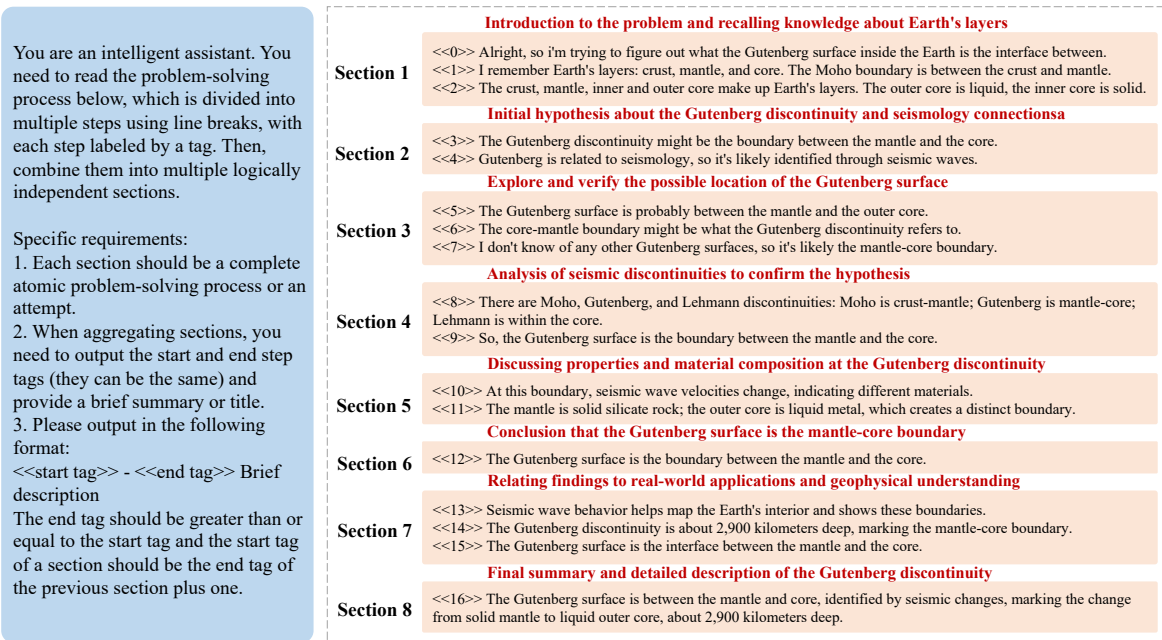


Figure 13: An example of dividing sections.

**Section-Level Evaluation.** The Long COT is divided into sections, each corresponding to a specific sub-task, such as problem analysis, verification of calculation results, and summarization. This phase requires the annotator to check and annotate each section individually. The annotation process and examples are shown in Figure 3.

Label	Fleiss' Kappa Score (Three Annotators)	Cohen's Kappa (Two Reviewers)	Pass Rate (Spot Check)	Consistency Acc. (Cross-Validation)
Strategy Shift	0.767	0.709	96.60%	96.67%
Reasoning Usefulness	0.670	0.608	92.55%	93.33%
Reasoning Correctness	0.744	0.843	97.41%	95.00%
Reflection Efficiency	0.705	0.765	93.60%	91.67%

Table 6: Fleiss' Kappa scores for various annotation domains, indicating the level of agreement among annotators.

**Quality Assurance and Validation.** We have established a strict quality control process to ensure the high quality and consistency of annotations. Each data is assigned three initial annotators, two junior reviewers, and an additional five people who are responsible for overall spot checks. Outsourced personnel and external contractors responsible for annotations receive unified training. We regularly check the consistency and quality between annotators and repeatedly discuss and improve annotation protocols during the annotation process to make the standards more perfect. This process ensures the generation of high-quality annotations and minimizes subjective bias.

We acknowledge the inherent subjectivity in manual annotations. To mitigate this bias, detailed, objective standards and rigorous quality control processes were developed. Meanwhile, to quantify the consistency of our annotations, we employed Fleiss' Kappa, an inter-annotator agreement metric measuring the agreement among multiple annotators. We computed Fleiss' Kappa for each label, assessing the level of agreement among three annotators using binary labels (0/1) for each annotation task. Fleiss' Kappa ranges from -1 to 1, with scores closer to 1 indicating stronger agreement. Scores above 0.60 are regarded as substantial. As shown in Table 6, our annotation process effectively minimizes bias through a voting mechanism and strict cross-validation, while subjectivity cannot be entirely eliminated.

**Profile of Annotation Persons.** In this dataset annotation project, we engaged a diverse group of annotators through three distinct sources. We employed a set of external contractors directly recruited by us and collaborated with two additional suppliers to provide annotation services. The dataset was divided



into three parts, with certain sections deliberately overlapping to facilitate cross-validation. This overlap allows us to compute the annotation consistency rate, and if the results do not meet the required standards, revisions are necessitated.

Our annotator pool is composed of highly qualified individuals: 23 Master’s degree holders and 6 Ph.D. holders in Mathematics, 18 Master’s graduates in Computer Science, 7 Master’s and 2 Ph.D. graduates in Physics, 7 Master’s and 3 Ph.D. graduates in Chemistry, and 6 Master’s and 2 Ph.D. degree holders in Biology. We employ a rotational system to randomly assign two individuals from each academic field to serve as reviewers. Additionally, 5 algorithm specialists are tasked with conducting spot checks on the annotations. This meticulous selection and review process ensures superior data quality and reliability for subsequent analyses.

### B.3 Details on Assessment

**Math** For mathematical queries, we employ a combination of rules and LLMs to evaluate the correctness of the provided solutions. Rule-based systems verify the validity of numerical calculations, while LLMs ensure that reasoning steps adhere to established mathematical principles. This dual approach guarantees high accuracy in error detection within the solutions.

**Programming** For programming tasks, we utilize sandbox testing environments alongside LLM-based evaluations. Specifically, we utilize SandboxFusion (Cheng et al., 2024) as our testing environment. The solution is initially executed in the sandbox environment. Subsequently, the test case, the sandbox environment’s feedback output, and the code are provided to the LLM to determine the correctness of the answer.

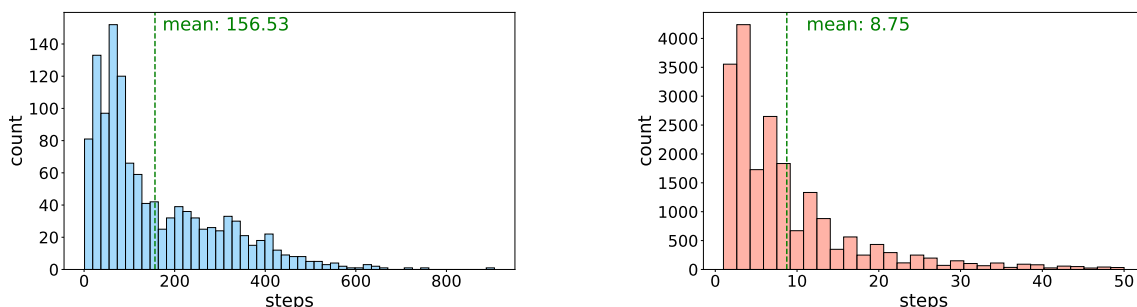
**PCB** Due to the straightforward nature of answers in these domains, we exclusively rely on LLM judgments, which offer high accuracy in assessing correctness.

**General Reasoning** Similarly, for general reasoning questions, LLM judgments are employed to effectively and accurately assess solution validity.

### B.4 Findings in Data Preprocessing

During the data preprocessing stage, we identified several issues with the data collected from open-source datasets. These issues included incomplete queries, incorrect solutions, and excessively high query similarity. To address these problems, we applied a combination of manual review and LLM (Large Language Model) validation to filter out low-quality data. Additionally, for code data specifically, we observed that different sources and types of data sometimes included test cases, while others did not, and the formats of these test cases were inconsistent. To tackle these inconsistencies, we used GPT-4 to filter the data for quality and to extract test cases, standardizing them into executable code for SandboxFusion. This allowed us to conduct uniform sandbox verification to ensure data accuracy.

### B.5 Sections Division



(a) Distribution of the number of steps in long CoT (divided by “\n\n”).

(b) Distribution of the number of steps contained in each divided section.

Figure 14: Statistical distribution of steps in long CoT.

Figure 13 illustrates the prompt for dividing sections along with examples of the resulting divisions. Several steps involved in addressing an atomic problem or exploring an idea are grouped into the same section. The specific outcome of the division is influenced by various factors, such as the task domain. However, compared to a purely long CoT, this approach is more user-friendly for human annotation.

Furthermore, to prevent sections from becoming overloaded with too many steps, which would increase the complexity of the annotation process, we iteratively divide sections that exceed 50 steps. Figure 14 displays the distribution of steps in the original long CoT (subfigure 14a) and the distribution of steps in each divided section (subfigure 14b). Before sectioning, annotators are required to review each individual step, which can be exceedingly challenging for long CoTs with numerous steps. By dividing the sections, annotators can proceed on a section-by-section basis, making the process more comprehensible and significantly reducing the difficulty of annotation.

## B.6 Validation of Long CoT Correctness

Domain	Filtered Total Num	Correct Num	Accuracy(%)
Math	7534	5413	71.84
Programming	2103	1276	60.67
PCB	4517	2626	58.13
General Reasoning	1981	1160	58.56

Table 7: Accuracy statistics for generated Long CoT responses for filtered high-quality queries.

For the filtered high-quality queries, we use a mix of various o1-like models, including QwQ-32B-Preview, DeepSeek-R1, and Gemini 2.0 Flash, to generate the corresponding long CoT. We then use LLM-as-a-judge and a sandbox testing environment to validate the accuracy of the long CoT generated by these o1-like models, obtaining the native erroneous long CoT for subsequent human annotation.

Table 7 shows the accuracy of the generated long CoT. It can be seen that o1-like models enhanced with reinforcement learning in math and programming perform slightly better in these two areas compared to general reasoning and PCB.

## B.7 Statistics on Category Distribution

Table 8 shows the subcategories and corresponding data volumes of DeltaBench across various domains. In obtaining queries and annotations, we strive to ensure balance across categories while also balancing annotation difficulty and accuracy.

## B.8 Analysis of Other Evaluation Metrics

Domain	Subcategory	Num
Math	Discrete Mathematics	144
	Number Theory	104
	Geometry	101
	Others	74
	Calculus and Analysis	58
	Statistics and Other Decision Science	45
	Algebra	36
Programming	Basic Programming	133
	Mathematics	86
	Advanced Programming	48
	Data Analysis	41
	Desktop and Web Development	27
	Others	24
	Software Engineering	14
PCB	Chemistry	64
	Physics	63
	Biology	27
General Reasoning	Logical Reasoning	56
	Symbolic Reasoning	28
	Quantitative Reasoning	24
	Strategic Reasoning	12
	Common Sense Reasoning	9
	Spatio-temporal Reasoning	9
	Others	9
Total		1236

Table 8: Detailed categories of DeltaBench and corresponding data volume statistics.

model	F1-Score	First Error Acc.	Any Error Acc.
GPT-4-turbo-128k	40.76	57.04	69.17
GPT-4o	30.85	36.89	50.89
DeepSeek-V3	27.33	31.72	42.39
Qwen2.5-32B-Instruct	26.73	30.58	42.23
DeepSeek-R1	28.43	29.94	40.78
Qwen2.5-7B-Instruct	18.63	22.25	30.74
GPT-3.5	7.98	6.15	11.65

Table 9: The table compares different accuracy metrics for each model. 'First Error Acc.' is the accuracy in identifying the first error, and 'Any Error Acc.' is the accuracy in detecting any error.

In Table 9, we present the performance of several models across different accuracy metrics: F1-Score, First Error Accuracy, and Any Error Accuracy. These metrics evaluate the models' ability to identify the first error and detect any error within a given sequence. A key observation is that the relative rankings of the models across the First Error Accuracy and Any Error Accuracy metrics closely align with their F1-Score. This consistency across different evaluation measures highlights the robustness of the F1-Score as a comprehensive indicator of model performance and suggests a strong correlation between the ability to detect the first error and the ability to identify any error in the sequence. Additionally, GPT-4-turbo consistently outperforms the other models, regardless of the evaluation metric used. Its Any Error Accuracy reaches 69%, significantly higher than the other models in the comparison. This finding underscores the model's superior performance in error recognition, yet it also points to the limitations that remain in current LLMs.

For PRMs, aside from outlier detection, we also experimented with evaluating using a fixed threshold

Model	Quantile	Threshold	prec	recall	F1
Qwen/Qwen2.5-Math-PRM-7B	5%	0.2168	39.81	73.86	46.48
Qwen/Qwen2.5-Math-PRM-72B	5%	0.2119	33.51	65.11	40.44
RLHFlow/Llama3.1-8B-PRM-Deepseek-Data	5%	0.2021	24.19	56.1	30.88
RLHFlow/Llama3.1-8B-PRM-Mistral-Data	5%	0.2949	23.18	51.46	29.68
Skywork/Skywork-o1-Open-PRM-Qwen-2.5-1.5B	5%	0.0303	19.48	46.76	24.45
Skywork/Skywork-o1-Open-PRM-Qwen-2.5-7B	5%	0.0278	18.68	46.14	23.46

Table 10: Performance of PRMs using the overall reward quantile as the threshold.

based on quantiles. Specifically, we used the ascending 5% quantile of all rewards on DeltaBench as the threshold, considering sections below this value as incorrect. The evaluation results are shown in table 10. However, compared to outlier detection, we found that this approach overestimates the performance of PRMs. This is because using a quantile as a threshold effectively forces PRMs to consider a fixed proportion of sections as incorrect.

## B.9 More Related Works

**Test-time Scaling.** Recently, many works have begun to explore the test-time scaling techniques (OpenAI, 2024; Snell et al., 2024), can greatly enhance performance by increasing the number of generated tokens. Several methods (Hao et al., 2023; Chen et al., 2024a; Yao et al., 2023) have developed the tree search methods to improve reasoning capabilities. In addition, o1-like models (e.g., o1, R1) (OpenAI, 2024; Guo et al., 2025) have investigated the reinforcement learning methods to generate long Chains of Thought (CoT) and enhance the model performance.

## C Error Classification

In Figure 15, we conclude a detailed error classification based on human annotations of the errors contained in the model’s answers.

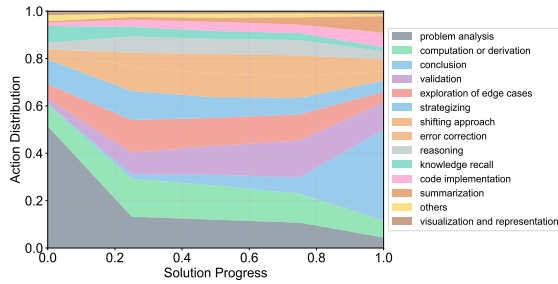
## D Analysis of o1-like Models

### D.1 Distribution of Action Role

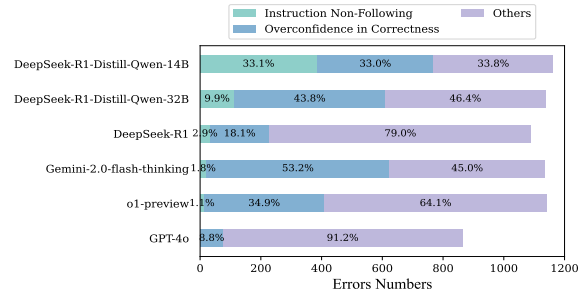


Primary Category	Subcategory	Description
Understanding Error	Problem Misunderstanding	Incorrect interpretation of the problem requirements or context at the initial comprehension stage
	Conceptual Misunderstanding	Misunderstanding of basic principles or theoretical foundations
Knowledge Errors	Factual Error	Incorrect recall or statement of established facts or constants
	Theorem Error	Incorrect recall or application of mathematical theorems or scientific laws
	Definition Error	Incorrect understanding or application of standard definitions or terminology
Logical Errors	Strategy Error	Inappropriate selection of problem-solving approaches or methodologies
	Reasoning Error	Flaws in the logical flow of problem-solving steps
	Premise Error	Invalid or incorrect assumptions in the reasoning process
	Consistency Error	Contradictions or inconsistencies in the logical framework
Calculation Errors	Numerical Error	Inaccuracies in arithmetic operations or numerical computations
	Formula Error	Incorrect application or manipulation of mathematical formulas
	Parameter Error	Mistakes in handling variables or parameters
	Unit Error	Incorrect use or conversion of measurement units
Programming Errors	Syntax Error	Violations of programming language rules and conventions. Including incorrect code structure, missing delimiters, or improper statement construction
	Function Error	Deficiencies in function implementation, usage, or behavior. Including improper function definitions, incorrect function calls, or unexpected function behaviors
	Data Type Error	Incorrect handling of data types and type conversions. Including type mismatches, improper type conversions, or inappropriate data structure usage
Formal Errors	Symbol Error	Incorrect use of mathematical or scientific symbols
	Formatting Error	Improper presentation or organization of solution structure
Completeness Errors	Boundary Omission	Failure to consider edge cases or limiting conditions
Special Cases	Reflection Error	Insufficient analysis of solution validity or limitations
	Summary Error	Inadequate synthesis or conclusion of findings
	Hallucination	Generation of false or unsupported information
	Redundancy	Unnecessary repetition or superfluous information in solutions

Figure 15: Categories of Errors in o1-like Models.



(a) Distribution of Action Roles in Long CoT Process.



(b) Distribution of Error Causes for Each Model.

Figure 16a shows the distribution of action roles in the system 2 thinking process of the o1-like model. Initially, problem analysis dominates, indicating that the model initially focuses on understanding the requirements and constraints of the problem. As the solution progresses, cognitive activities diversify significantly. Verification and error correction gradually increase, indicating that the model begins to think and reflect frequently. In the later stages, verification, conclusion, and summary become more important, reflecting the model’s efforts in verifying the accuracy of the solution and synthesizing the results. In addition, the discrete sections of strategy formulation, knowledge recall, and error correction indicate a continuous adaptation and iteration process. As the model progresses from problem analysis, solution implementation and conclusion, it demonstrates the common reasoning template of slow-thinking models.

### E Analysis of Underperforming Model

We classify the errors produced by underperforming models into three categories: (1) Not following instructions, where models fail to follow the critique instructions; (2) Overconfidence in Correctness, where models incorrectly assume the sample contains no errors; (3) Other errors, such as false negatives, where models incorrectly flag accurate reasoning as erroneous, and error misidentification, where models inaccurately determine the type or location of errors.

As illustrated in Figure 16b, the following observations can be made: (1) DeepSeek-R1-Distill series exhibit significant instruction-following deficiencies. These models tend to directly answer questions rather than critically evaluate the correctness of responses, indicating a potential overfitting problem. (2) o1-like models like DeepSeek-R1 and o1-preview-0912 also demonstrate notable instruction-following challenges, albeit to a lesser extent than GPT-4o. (3) A notable trend among the o1-preview and Gemini-2.0-flash-thinking-exp models was their tendency to assume that samples were correct without thorough evaluation. This overconfidence in correctness was a significant contributor to their error rates. In contrast, GPT-4o demonstrates superior performance, with no instances of "Instructions not followed" errors and a relatively low number of "Overconfidence in Correctness" errors.