

# An Efficient and Precise Training Data Construction Framework for Process-supervised Reward Model in Mathematical Reasoning

Wei Sun<sup>1,2</sup>, Qianlong Du<sup>1</sup>, Fuwei Cui<sup>1</sup> and Jiajun Zhang<sup>†1,2,3</sup>

<sup>1</sup>Institute of Automation, Chinese Academy of Sciences

<sup>2</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>3</sup>Wuhan AI Research

{sunwei2023, fuwei.cui}@ia.ac.cn {qianlong.du, jjzhang}@nlpr.ia.ac.cn

## Abstract

Enhancing the mathematical reasoning capabilities of Large Language Models (LLMs) is of great scientific and practical significance. Researchers typically employ process-supervised reward models (PRMs) to guide the reasoning process, effectively improving the models' reasoning abilities. However, existing methods for constructing process supervision training data, such as manual annotation and per-step Monte Carlo estimation, are often costly or suffer from poor quality. To address these challenges, this paper introduces a framework called *EpicPRM* (Efficient, Precise, Cheap), which annotates each intermediate reasoning step based on its quantified contribution and uses an adaptive binary search algorithm to enhance both annotation precision and efficiency. Using this approach, we efficiently construct a high-quality process supervision training dataset named *Epic50k*, consisting of 50k annotated intermediate steps. Compared to other publicly available datasets, the PRM trained on *Epic50k* demonstrates significantly superior performance. Getting *Epic50k* at <https://github.com/xiaolizh1/EpicPRM>.

## 1 Introduction

Improving the mathematical reasoning ability of Large Language Models (LLMs) can significantly accelerate advancements in artificial intelligence across various scientific domains. To enhance the mathematical reasoning capability of LLMs, several studies (Yu et al., 2023a; Cobbe et al., 2021; Shao et al., 2024; Chen et al., 2025a) employ Outcome-supervised Reward Models (ORMs) to guide Chain-of-Thought (CoT) generation and leverage reinforcement learning to optimize generative performance. Among these approaches, process supervision, leveraging Process-supervised Reward Models (PRMs), has emerged as a more effective approach (Li et al., 2022; Uesato et al.,

2022; Lightman et al., 2023; Chen et al., 2024b), providing substantial improvements through step-by-step supervision.

In process supervision, researchers typically construct an annotated dataset of reasoning processes and then train a PRM to guide the reasoning process. The quality of this annotated data is crucial for PRM effectiveness. Two primary approaches exist for dataset construction: manual annotation (Lightman et al., 2023) and automatic annotation (Wang et al., 2024a,b; Luo et al., 2024). While manual annotation ensures high data quality, it is labor-intensive and difficult to scale to diverse domains. Conversely, automatic annotation reduces labor costs but often results in lower annotation quality. This necessitates the use of larger datasets to compensate, which in turn demands significant computational resources.

To address these challenges, this paper introduces *EpicPRM*, a fully open-source framework designed for the precise and efficient construction of high-quality annotated process supervision training datasets. Our approach leverages multiple LLMs with varying capabilities and from diverse sources to generate a more comprehensive set of CoT reasoning chains. We annotate intermediate steps by quantifying their contribution and identifying the first erroneous step within each chain. This identification is achieved through an adaptive binary search process, which adjusts the starting position and the number of Monte Carlo estimation samples based on problem difficulty. These optimizations result in a 64.39% reduction in annotation costs compared to traditional sequential search algorithms. Using this framework, we constructed *Epic50k*, a dataset containing 50k annotated intermediate steps. The PRM trained on *Epic50k* achieves supervision performance comparable to, and in some cases surpassing, PRMs trained on the PRM800k and Math-Shepherd datasets, despite being less than 10% the size.

<sup>†</sup>Corresponding author.

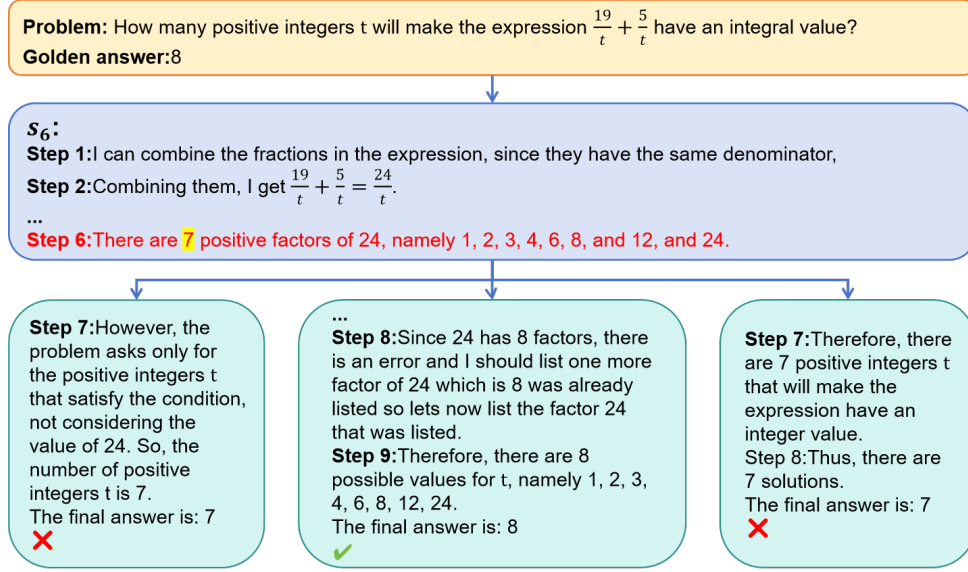


Figure 1: An examples of getting the correct answer based on the wrong steps. On the basis of  $s_6$ , completer continues to sample a number of rollouts. We can see that since the completer has a certain error correction capability, even if there is an obvious error in step 6.

Our main contributions are as follows:

1. We propose a more efficient and precise framework for constructing process supervision data by quantifying the contribution of intermediate reasoning steps and optimizing the binary search algorithm.
2. We publicly release *Epic50k*, a high-quality process-supervised training dataset containing 50k intermediate reasoning steps.
3. Our experimental results demonstrate the importance of data quality over quantity for training PRMs, showing that even a smaller, high-quality dataset can yield superior results.

## 2 Background

Existing methods for constructing process-supervision data first use LLMs to generate complete CoT chains, and then use Monte Carlo estimation to label the correctness of each step.

**Monte Carlo estimation.** Several related studies, including Math-Shepherd (Wang et al., 2024a), MiPS (Wang et al., 2024b), and OmegaPRM (Luo et al., 2024), have introduced a Monte Carlo (MC) estimation method to approximate the probability of obtaining a correct solution from a given state. This method designates a large language model (LLM) as a "completer," tasked with generating the subsequent reasoning steps from the current state  $s_t$  ( $s_t = [q; a_1; a_2; \dots; a_t]$ , where  $q$  denotes the question and  $a_1$  through  $a_t$  represent the first  $t$  selected intermediate steps) until a final answer is reached. This generative process is termed a "roll-

out." After performing  $N$  independent rollouts, the number of successful rollouts  $M$  that arrive at the correct final answer is counted, and the probability of achieving a correct solution from the current state is estimated as  $M/N$ .

**Identification of the first erroneous step.** In automated data construction, evaluating the correctness of intermediate steps within a complete erroneous CoT reasoning chain using MC estimation is computationally expensive. Prior work (Lightman et al., 2023) has shown that effective PRM training only requires the identification of the first erroneous step in a complete CoT reasoning chain. Once the first erroneous step is identified, all preceding steps are annotated as correct, while subsequent steps are annotated as incorrect. To optimize this identification process, OmegaPRM (Luo et al., 2024) employs a binary search algorithm, reducing the number of MC estimations needed and thereby improving computational efficiency.

## 3 Method

In this section, we introduce how to make MC estimation more precise and how to find the first erroneous step more efficiently.

### 3.1 Make MC Estimation More Precise

#### 3.1.1 Perplexity Instead of Count

Using MC estimation to represent the probability of obtaining the correct answer from a given state has inherent limitations. When  $N$  is not large enough,

counting will be accidental, we argue that using  $M/N$  as a direct representation of the probability of sampling correct rollouts is imprecise. According to the Law of Large Numbers, estimating probabilities using counting methods necessitates a large number of samples to achieve reliable accuracy. For instance, observing two heads in two coin tosses does not imply a 100% probability of obtaining heads. Similarly, if we sample  $N$  rollouts and find  $M$  correct ones, estimating the probability of correctness as  $M/N$  inherently carries uncertainty due to sampling variability. By using perplexity, we can directly calculate the probability of the model generating each rollout, thereby mitigating errors introduced by sampling. Furthermore, if a correct rollout has an exceptionally low generation probability but is included in the sample, a counting-based method would overestimate the likelihood of obtaining correct rollouts. Therefore, we propose using perplexity (PPL) calculated by logarithmic probability to represent the sampling probability of each rollout:

$$PPL(j; s_t, \theta_k) = \exp \left\{ -\frac{1}{L} \sum_{i=1}^L \log p_{\theta_k}(x_i | x_{<i}) \right\} \quad (1)$$

where  $j$  represents a rollout,  $s_t$  represents the input prefix,  $\theta_k$  represents the  $k$ -th completer;  $p_{\theta_k}(x_i | x_{<i})$  is the probability that completer  $k$  generates the token  $x_i$  given the preceding tokens  $x_{<i}$ ; and  $L$  is the total number of tokens in rollout  $j$ .

Consequently, each of  $K$  completers which independently samples  $N$  rollouts, with  $M$  of these rollouts yielding the correct answer, our MC estimate based on perplexity (PPL), denoted as  $MC_{PPL}$ :

$$MC_{PPL}(s_t, \theta_{1:K}) = \frac{\frac{1}{K} \sum_{k=1}^K \sum_{m=1}^M \log PPL(j; s_t, \theta_k)}{\sum_{n=1}^N \log PPL(j; s_t, \theta_k)} \quad (2)$$

Compared with the counting, directly calculating the probability of obtaining correct answer eliminates the influence of accidental, that's why we use  $MC_{PPL}$  instead of  $M/N$ .

### 3.1.2 Quantify Contribution of Steps

**Limitations of MC Estimation.** The prevailing automatic annotation methods operate under the assumption that if at least one of  $N$  sampled rollouts originating from state  $s_t$  produces the correct final answer (equivalent to  $MC_{PPL}$  value greater than 0), then all  $t$  intermediate steps within state  $s_t$  are considered correct. However, the inherent error correction capabilities of LLMs can obscure underlying inaccuracies in these intermediate steps.

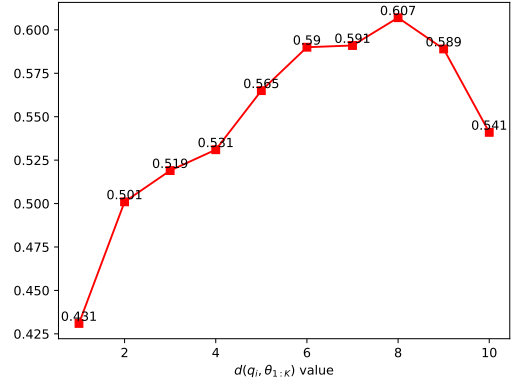


Figure 2: Relationship between the position of the first erroneous step and the difficulty of the problem. The horizontal axis represents the difficulty of the problem, and the vertical axis represents the ratio of the position of the first error step to the total number of steps.

As demonstrated in Figure 1, a completer may rectify an error at step  $a_t$  and still arrive at the correct final answer, despite  $a_t$  itself being flawed. This phenomenon results in false positives for  $a_t$ , consequently diminishing the quality of the process supervision data and adversely affecting the supervisory efficacy of the trained PRM.

**Contribution-Based Annotation Method.** To annotate intermediate reasoning steps more precisely, we quantify the contribution of each step  $a_t$  to problem solving as the ratio of  $MC_{PPL}(a_t)$  to  $MC_{PPL}(a_0)$ . This represents the ratio of the probability that the completer will continue to answer the question and ultimately arrive at the correct answer, given the first  $t$  steps (with the first  $t - 1$  steps have been verified correct), to the probability that the completer will arrive at the correct answer without any reasoning steps. In Section 4.1, we specifically analyze the relationship between the quantified contribution  $C(a_t; \theta_{1:K})$  and the correctness of the steps annotated by human expert through experiments and establish a **threshold**  $\alpha$  to be used as the criterion for annotating the steps:

$$C(a_t; \theta_{1:K}) = \frac{MC_{PPL}(s_t, \theta_{1:K})}{MC_{PPL}(s_0, \theta_{1:K})} \quad (3)$$

$$Label(a_t; \theta_{1:K}) = \begin{cases} 0 & \text{if } C(a_t; \theta_{1:K}) \leq \alpha \\ 1 & \text{if } C(a_t; \theta_{1:K}) > \alpha \end{cases}$$

### 3.2 Adaptive Binary Search Algorithm

In this subsection, we detail how to use the adaptive binary search algorithm to improve the efficiency of finding the first error step.

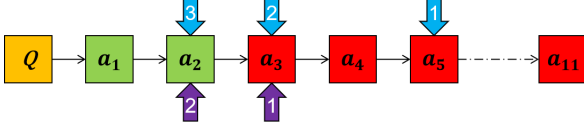


Figure 3: An example of comparing search processes between our adaptive binary search and traditional binary search. The blue arrow is traditional binary search, requiring 3 searches (starting from  $a_5$ ) to find the first erroneous step  $a_3$ . Conversely, our adaptive binary search (purple arrow) sets the initial search position to  $a_3$  based on question difficulty, requiring only 2 searches.

### 3.2.1 Dynamic Starting Position

Human problem-solving behavior suggests that errors typically occur early in more complex problems, while simpler problems often involve calculation mistakes later in the solution. We hypothesize that LLMs follow similar patterns. To test this, we sampled 22k erroneous rollouts and identified the position of the first erroneous step. We used the MC perplexity value of the initial problem state,  $MC_{PPL}(s_0)$ , to represent problem difficulty—lower indicates greater difficulty. We classified problem difficulty  $d(q_i, \theta_{1:K})$  into 11 levels:

$$d(q_i, \theta_{1:K}) = \text{round}(10 * MC_{PPL}(s_0, \theta_{1:K})) \quad (4)$$

Our results, shown in Figure 2, reveal a strong correlation between the first erroneous step’s position and problem difficulty. Based on these insights, we propose an enhanced binary search algorithm with a dynamic starting position, an example is shown in Figure 3. Instead of starting at  $T/2$ , we adapt the starting point based on the problem’s difficulty  $d_{q_i}$  using Eq.(5). This adjustment allows faster narrowing of the search range, reducing the number of MC estimations. The complete data construction process is shown in Algorithm 3.2.2.

$$mid = \begin{cases} \left\lfloor \frac{l+r}{2} \right\rfloor - \left\lfloor \frac{Len}{4} \right\rfloor & \text{if } d(q_i, \theta_{1:K}) < 2 \\ \left\lfloor \frac{l+r}{2} \right\rfloor & \text{if } 2 \leq d(q_i, \theta_{1:K}) < 6 \\ \left\lfloor \frac{l+r}{2} \right\rfloor + \left\lfloor \frac{Len}{4} \right\rfloor & \text{if } d(q_i, \theta_{1:K}) \geq 6 \end{cases} \quad (5)$$

Where  $d(q_i, \theta_{1:K})$  represents the difficulty of  $q_i$  with K completers,  $r$  and  $l$  represent the upper and lower bounds of the binary search, and  $Len$  represents the total number of steps.

### 3.2.2 Dynamic Sampling Count

We observe that the computational overhead in automatic annotation is primarily attributable to the completer models’ sampling of  $N$  rollouts. Existing methods employ a fixed value for  $N$  irrespective of problem difficulty, resulting in wasted computation on simpler problems and potentially insufficient sampling for complex problems, where the

number of valid rollouts  $M$  at both  $s_0$  and  $s_t$  may be zero. To address this, we propose dynamically adjusting  $N$  based on problem difficulty. Specifically, we determine  $N$  at the initial state  $s_0$  (as defined in Eq.(2)) by initializing with a value of 16 and iteratively incrementing until  $M$  exceeds 10, thereby ensuring a minimum of 10 correct rollouts from each completer. This adaptive sampling strategy mitigates computational overhead while maintaining annotation quality.

---

#### Algorithm 1 Adaptive Binary Search for Annotating Process Supervision Data

---

**Input:** Question  $q_i$ , response  $A_T = (a_1, \dots, a_T)$ , completers  $\theta_{1:K}$ , threshold coefficient  $\alpha$

**Output:** The subscript of the first error in  $A_T$

- 1: Initialize the lower bound  $Left = 0$ , upper bound  $Right = T - 1$ , flag for the first search  $f = True$ ;
  - 2: Calculate the average  $MC_{PPL}$  value and the threshold  $H$  for error steps,  $V = MC_{PPL}(s_0, \theta_{1:k})$ ,  $H = \alpha * V$ ;
  - 3: **while**  $Left \leq Right$  **do**
  - 4:   **if**  $f = True$  and  $T \geq 4$  **then**
  - 5:     Compute  $mid$  according to Eq.(5);
  - 6:   **else**
  - 7:      $mid = \left\lfloor \frac{Left+Right}{2} \right\rfloor$ ;
  - 8:   **end if**
  - 9:    $s_{mid} = (q_i; a_1, \dots, a_{mid})$ ;
  - 10:    $f = False$ ;
  - 11:   **if**  $MC_{PPL}(s_{mid}, \theta_{1:K}) \leq H$  **then**
  - 12:      $Right = mid - 1$ ;
  - 13:   **else**
  - 14:      $Left = mid + 1$ ;
  - 15:   **end if**
  - 16: **end while**
  - 17: **return**  $Left$
- 

## 4 Experiments

### 4.1 Annotation Threshold Selection

In this section, we experimentally determined the threshold in Section 3.1.2. We sampled 1,000 steps from the manually annotated PRM800k(Lightman et al., 2023) and analyzed the ratio of their  $MC_{PPL}$  values at  $s_0$  and  $s_t$  to assess  $a_t$ ’s contribution. Figure 4 shows that some manually annotated incorrect steps still had a positive ratio, indicating at least one correct rollout. This suggests that certain errors have minor negative effects, while some correct steps may introduce slight disruptions due to



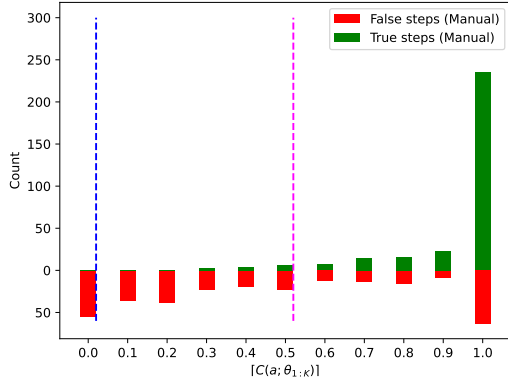


Figure 4: Relationship between manually annotated labels and step contribution. The x-axis depicts the quantified contribution per step, where the step values have been rounded up. All values greater than 1 are counted as 1. The vertical axis represents the number of manually annotated steps. The blue and pink dotted lines represent the existing method and our method respectively. The left side of dotted lines is annotated as the incorrect step, and the right side is annotated as the correct step. Our method yields more precise results in identifying correct and incorrect steps.

unconventional reasoning. As the ratio increases, correct steps become more prevalent while incorrect ones decrease. So our human experts have selected 0.5 as the threshold ( $\alpha$ ) after random sampling validation. As depicted in the bar chart of Figure 4, indicate that a randomly chosen threshold within the range of 0 to 1 surpasses the discriminative capabilities of existing approaches, suggesting that human expert validation may not be necessary.

## 4.2 PRM Training Data Construction

### 4.2.1 CoT Reasoning Chain Generation

We use the MATH dataset (Hendrycks et al., 2021) to construct Epic50k. We adopted the same training and test set splits as PRM800k (Lightman et al., 2023), creating the MATH500 test set by selecting 500 representative questions from the original 5k-question MATH test set, and expanding the training set to 12k questions. Initially, we randomly selected 3,500 questions from the training set, maintaining a level distribution ratio of 1:1:1:2:2 across levels 1 through 5. We then employed the LLaMA3-8B-Instruct (Touvron et al., 2023; AI@Meta, 2024), LLaMA3.1-8B-Instruct, and Qwen2-7B-Instruct (Yang et al., 2024a) models to generate multiple complete solutions for these questions. LLaMA3-8B-Instruct was used for levels 1 and 2, while LLaMA3.1-8B-Instruct and

Qwen2-7B-Instruct were used for levels 3 through 5. For each question, we selected 2 correct and 6 incorrect solutions exhibiting the lowest cosine similarity scores, and split each intermediate step of the solutions for automatic annotation.

### 4.2.2 Intermediate Steps Annotation

When annotating intermediate steps, the  $MC_{PPL}$  values sampled by different LLMs can exhibit significant variability due to differences in their capabilities and biases. To mitigate this variability and more precisely reflect the true accuracy of state  $s_t$ , we employed LLaMA3.1-8B-Instruct and Qwen2-7B-Instruct from multiple diverse sources as completers and calculated the average of their sampled  $MC_{PPL}$  values as the final estimate. As our model was not fine-tuned on mathematical datasets, we standardized the solution output format using a two-shot prompting approach (Appendix A.1 for details). In total, we constructed a PRM training dataset containing 50,000 intermediate steps, with an approximate 1:1 ratio of correct to incorrect steps. To accelerate sampling, we utilized the vLLM inference framework (Kwon et al., 2023). The specific comparative analysis of step distribution can be found in the Appendix A.2

## 4.3 PRM Supervision Performance

### 4.3.1 PRM Training

Because our approach requires no additional supervised fine-tuning data, we chose Qwen2-math-1.5B-base (Yang et al., 2024a), a math-specialized LLM, as our base model. Unlike Math-Shepherd, which uses the complete solution sequence  $s_{1:T}$  as input and calculates softmax values over the vocabulary at special tokens of each step for loss computation, we treat PRM training as binary text classification which is same as PRM800k. We added a classification head to the base model’s output layer and employ a standard binary classification loss:

$$Loss = \frac{1}{R} \sum_{i=1}^N \hat{y}_i \log y_i + (1 - \hat{y}_i) \log (1 - y_i) \quad (6)$$

Where  $\hat{y}_i$  represents the correctness label and  $y_i$  is the prediction score of the PRMs.  $R$  represents the number of training samples.

When evaluating an intermediate step  $a_t$ , we input only the preceding steps  $s_t$ . This is preferable because inference-time iterative search algorithms (e.g., beam search) do not yet have subsequent content. We changed the format of the MATH-

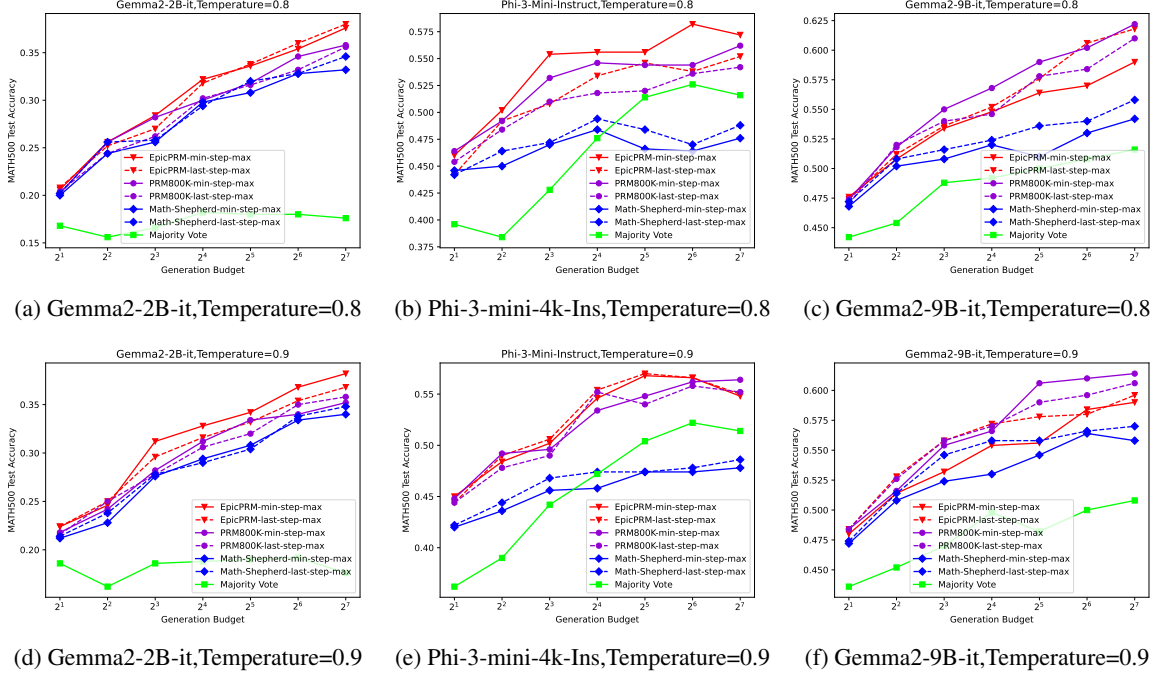


Figure 5: Results of PRM weighted Best-of-N search for models of different scales at commonly used sampling temperatures.

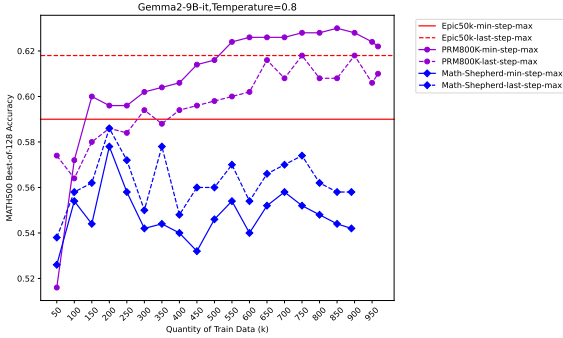


Figure 6: Comparison of the supervision performance of PRMs trained with varying datasets: subsets of PRM800k and MATH-Shepherd with at least 50k scales versus ours Epic50k.

Shepherd dataset to a format suitable for binary text classification tasks.

### 4.3.2 Experiment Setup

We evaluated the effectiveness of the PRM on Gemma2-2B-it, Gemma2-9B-it (Team et al., 2024), Phi-3-mini-4k-Instruct (3.8B) (Abdin et al., 2024), which are from different origins than the data generation models, using the **PRM weighted Best-of-N search** on MATH500. Specifically, a basic model was used as a generator, which sampled  $N$  complete solutions on the MATH500 test problem. We then used PRM to score these  $N$  solutions and selected one solution as the final an-

swer. There are two selection strategies: one takes the minimum PRM score as the solution’s score ( $Value = \min_{t \rightarrow T} PRM(q_i|s_t)$ ); the other takes the PRM score of the last step as the solution’s score ( $Value = PRM(q_i|s_T)$ ). Baseline process supervision datasets included PRM800k (Lightman et al., 2023), Math-Shepherd (Wang et al., 2024a), and Self-Consistency (Wang et al., 2022). To align with our training methodology, we modified the Math-Shepherd data format accordingly.

### 4.3.3 PRM Weighted BON Search Result

Figure 5 presents a comparison of PRM performance across various process-supervision datasets. Despite containing only 50k annotated intermediate steps, Epic50k demonstrates superior performance compared to Math-Shepherd across three distinct model sizes and sampling temperatures. Notably, Epic50k even surpasses PRM800k, which is trained on 800k manually annotated steps, in two of the evaluated models. Given the inherent precision of manually annotated data, it is unexpected that Epic50k, with a dataset size representing only 6.25% of PRM800k, would consistently outperform it across all models. Table 1 shows a comparative analysis of the Best-of-128 search results for different generators under varying PRMs supervisions. Epic50k’s consistent outperformance of existing automatic annotation methods sufficiently demonstrates the efficacy of our approach.

Generator Models	Size	Verifiers	Temperature	MATH500
Gemma2-2B-it	2B	2-shot CoT (Greedy)	0	0.202
		+Self-Consistency mj@128		0.176
		+Math-Shepherd Best-of-128	0.8	0.346
		+PRM800k Best-of-128		0.358
		+Epic50k (Ours) Best-of-128		<b>0.380</b>
Phi-3-mini-4k-Instruct	3.8B	2-shot CoT (Greedy)	0	0.378
		+Self-Consistency mj@128		0.516
		+Math-Shepherd Best-of-128	0.8	0.488
		+PRM800k Best-of-128		0.562
		+Epic50k (Ours) Best-of-128		<b>0.572</b>
Gemma2-9B-it	9B	2-shot CoT (Greedy)	0	0.456
		+Self-Consistency mj@128		0.516
		+Math-Shepherd Best-of-128	0.8	0.558
		+PRM800k Best-of-128		<b>0.622</b>
		+Epic50k (Ours) Best-of-128		0.618

Table 1: Comparison of the supervision effectiveness of PRMs trained on varying training datasets, evaluated using a Best-of-128 search with a temperature parameter set to 0.8.

Algorithm	Verified steps	Sampled number	Generated tokens
Sequential Search	4204	209.81K	76.67M
Binary Search	2620(-37.68%)	133.78K(-36.24%)	42.52M(-44.54%)
<b>Adaptive Binary Search(Ours)</b>	<b>2539(-39.56%)</b>	<b>70.40K(-66.45%)</b>	<b>27.30M(-64.39%)</b>

Table 2: Cost of three search algorithms when annotating the same 800 erroneous steps.

#### 4.4 Comparison of Different Data Scales

We compared Epic50k with randomly selected subsets of PRM800k and MATH-Shepherd using Gemma2-9b-it as the generator with a sampling temperature of 0.8, consistent with Figure 5c. The results, shown in Figure 6, The horizontal axis represents the size of the randomly selected training data, while the vertical axis depicts the MATH500 Best-of-128 accuracy. Epic50k consistently outperforms MATH-Shepherd across all data sizes and is only surpassed by PRM800k when the data exceeds 500k. This highlights Epic50k’s superior data efficiency. We posit that constructing a larger-scale training data using the EpicPRM framework would likely result in PRM supervision exceeding that of PRM800k. However, due to computational constraints, we limited our dataset to 50k.

#### 4.5 Data Annotation Cost

##### 4.5.1 Compared Search Algorithms

We also compared our method with the binary search algorithm and the sequential search algorithm with fixed sampling times to verify the advantage of our method in annotation efficiency. We annotated 800 randomly selected solutions. Both the

binary search algorithm and the sequential search algorithm fixed the sampling times to 48, while our method sampled between 16 and 72 times until at least 10 correct solutions were sampled. Finally, we compared the number of sampling times and the number of tokens generated by the three search methods to annotate 800 incorrect solutions.

##### 4.5.2 Annotated Cost Comparison Results

Table 2 presents the overhead compared with different search algorithms during annotating intermediate steps. Our adaptive binary search algorithm dynamically adjusts the number of samples and the initial search position based on the problem’s difficulty. Compared to other search algorithms, it achieves significant reductions in 3 key metrics: the number of verified steps, total sampled number, and total generated tokens. In particular, the reduction in total generated tokens (a metric directly related to cost) demonstrates the strong effectiveness of our approach. We construct Epic50k using 4 NVIDIA A100 80G GPUs for approximately 260 hours, which is quite cost-effective and cheap.

In terms of validation steps, our method shows only a modest improvement over conventional binary search algorithms. This is because the average

Model	GSM8K	MATH	Olympiad-Bench	Omni-MATH	Average
Math-Shepherd-PRM-7B	47.9	29.5	24.8	23.8	31.5
RLHFlow-PRM-Mistral-8B (Dong et al., 2024)	50.4	33.4	13.8	15.8	28.4
RLHFlow-PRM-Deepseek-8B	38.8	33.8	16.9	16.9	26.6
Skywork-PRM-1.5B (Liu et al., 2024)	59.0	48.0	19.3	19.2	36.4
Skywork-PRM-7B	70.8	53.6	22.9	21.0	42.1
Qwen2-1.5B-PRM800k	34.0	55.3	34.2	41.0	41.1
Qwen2-1.5B-MATH-Shepherd	48.9	34.1	9.8	13.7	26.6
Qwen2-1.5B-Epic50k (Ours)	55.6	36.1	20.2	30.0	35.5
Qwen2.5-1.5B-Instruct-PRM800k	49.1	59.3	32.5	35.9	44.2
Qwen2.5-1.5B-Instruct-Shepherd	44.1	29.8	12.1	17.6	25.9
Qwen2.5-1.5B-Instruct-Epic50k	58.1	53.2	32.2	40.6	46.0

Table 3: Evaluation results of open-source process reward models (PRMs) on PROCESSBENCH. We report the F1 score of the respective accuracies on erroneous and correct samples.

solution length in our dataset is just 10 steps, and most solutions can be completed with a standard binary search in fewer than 4 steps. Adjusting the initial search position provides noticeable advantages only for solutions with longer step lengths.

#### 4.6 Out of Domain Generalization

To evaluate our PRM on stronger generators and out of domain benchmarks, we conducted experiments on PROCESSBENCH (Zheng et al., 2024), a PRM assessment benchmark covering GSM8K, MATH, OlympiadBench, and Omni-MATH. Solutions were generated by LLaMA and Qwen models, including Qwen2.5-Math-7B/72B-Instruct, and labeled step-wise by human experts. Results are shown in Table 3. Epic50k demonstrates strong generalization across benchmarks of varying difficulty. While it lags behind human-labeled PRM800k, it performs competitively with open-source PRMs, excelling at Olympiad-level problems where it matches Skywork-PRM-7B. Though Qwen2-1.5B-Epic50k scores slightly below Skywork-PRM-1.5B (by <1%). Skywork leverages Qwen2.5-MATH-1.5B-Instruct, whereas our Qwen2-1.5B-Epic50k is based on the previous-generation Qwen2-MATH-1.5B-Base. Given the disparity in base models, this advantage is expected. Notably, despite being trained only on MATH, Qwen2-1.5B-Epic50k surpasses Skywork-PRM-1.5B on OlympiadBench and Omni-MATH, highlighting strong generalization from easier to harder problems. In order to be consistent with the base model of Skywork-PRM-1.5B, we conducted comparative experiments based on the Qwen2.5-1.5B-Instruct model. Our Epic50k performs significantly better than other methods.

## 5 Related Work

Improving the mathematical reasoning ability of Large Language Models (LLMs) can significantly accelerate advancements in artificial intelligence across various scientific domains, such as Translation (He et al., 2025), Scientific agent reasoning (Ren et al., 2025), Long-text reasoning (Wan et al., 2025; Chen et al., 2025b,a,c), Multilingual (Ghosh et al., 2025; Yang et al., 2023), Multimodal (Huang et al., 2025; Jian et al., 2024). **Prompt-based approaches.** Particularly Chain-of-Thought (CoT) (Wei et al., 2022), have significantly advanced mathematical reasoning (Perez et al., 2021). Subsequent work has explored CoT fine-tuning (Ouyang et al., 2022; Li et al., 2024; Mitra et al., 2024; Yu et al., 2023b) and extended prompt engineering to problem decomposition (Zhou et al., 2022; Hao et al., 2023) and programming (Chen et al., 2022; Li et al., 2023; Zhou et al., 2023; Chen et al., 2024a; Yang et al., 2024b). Selecting optimal solutions from multiple sampled paths has proven more effective than prompt engineering alone. Self-Consistency (Wang et al., 2022) uses majority voting across generated CoT trajectories, while RAP (Hao et al., 2023) uses self-evaluation and feedback for sequential subproblem solving. However, existing research (Huang et al., 2023) indicates that models encounter significant difficulty in self-assessing without external information.

**Application of PRMs.** Prior work (Huang et al., 2022; Wang et al., 2024a; Cobbe et al., 2021) has shown that using verifiers for problem-solving path assessment and selection is significantly more effective than self-consistency alone. PRMs have also been demonstrated to provide stronger supervision



than ORMs (Lightman et al., 2023). While Math-Shepherd (Wang et al., 2024a) uses a specialized PRM, OVM (Yu et al., 2023a) employs iterative step sampling with a step-level verifier to identify optimal steps. Thus, high-quality process supervision data is crucial for training PRMs, which can further enhance generator models’ mathematical abilities via methods like SFT (Zhang et al., 2024), GROP (Shao et al., 2024), and SVPO (Chen et al., 2024b). Existing data construction methods include costly manual annotation (e.g., PRM800k (Lightman et al., 2023)) and automatic annotation via Monte Carlo estimation (e.g., Math-Shepherd (Wang et al., 2024a), MiPS (Wang et al., 2024b)), with OmegaPRM (Luo et al., 2024) optimizing the latter using binary search within MCTS.

## 6 Conclusion

This study introduces EpicPRM, an open-source framework designed for the automatic annotation of intermediate reasoning steps in mathematical problem-solving. EpicPRM offers significant improvements in both precision and efficiency compared to existing automatic annotation methods. Notably, it achieves superior training results with less than 10% of the data volume required by current state-of-the-art training datasets.

DeepSeek-R1 (DeepSeek-AI et al., 2025) attempted to apply process supervision but failed due to the high cost of PRM training data and its limited supervisory effectiveness. However, this does not disprove the feasibility of process supervision. In contrast, Microsoft’s rStar-MATH (Guan et al., 2025) demonstrated that models with 7B parameters or fewer can achieve reasoning capabilities comparable to OpenAI-o1 using process supervision. Our method further reduces PRM training costs while improving its supervisory effectiveness. This advancement provides more practical and efficient solutions for improving the reasoning capabilities of language models.

## Limitations

EpicPRM has made significant progress compared to existing automatic annotation methods, but it still has some limitations. Although we have significantly improved annotation precision, it has not yet reached the level of human expert annotation. Consequently, our automatically annotated data still contains some inherent noise. Furthermore, our current method relies on human-provided gold

standard answers, meaning that complete independence from human supervision has not yet been achieved.

Theoretically, EpicPRM can be applied to any task that utilizes Chain-of-Thought (CoT) prompting and PRM-based process supervision. However, our experiment evaluation has been limited to the domain of mathematical reasoning. Therefore, the effectiveness and generalizability of EpicPRM in other reasoning domains, such as commonsense reasoning or logical inference, remain open questions. Future research will investigate the performance of EpicPRM across a wider range of tasks and domains to establish its broader applicability. This includes exploring potential adaptations or modifications that may be necessary to optimize performance in different contexts.

Given the high costs involved, we have not been able to perform a comprehensive analysis to assess the impact of adjusting the threshold in Section 3.1.2 on data quality, nor have we explored the possibility of task-specific thresholds for optimal annotation. These inquiries will be the focus of our future work.

## Acknowledgements

We thank our colleagues Jianghao Chen, Jian Pu, Wen Yang, Cong Li, Tengxiao Xi, Tianyu Peng, Ziliang Pang, Junhong Wu for their insightful and constructive feedback. We thank Qian Li and Zhenggang Piao for their special assistance. In addition, we thank all reviewers for their valuable comments and recognition of our work. This work is supported by National Key R&D Program of China N0.2022ZD0160602 and the Strategic Priority Research Program of Chinese Academy of Sciences under Grant N0.XDA04080400.

## References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- AI@Meta. 2024. [Llama 3 model card](#).
- Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. 2024a. Alphamath almost zero: process supervision without process. *arXiv preprint arXiv:2405.03553*.
- Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. 2024b. Step-level value preference optimiza-

- tion for mathematical reasoning. *arXiv preprint arXiv:2406.10858*.
- Jianghao Chen, Zhenlin Wei, Zhenjiang Ren, Ziyong Li, and Jiajun Zhang. 2025a. LR<sup>2</sup> bench: Evaluating long-chain reflective reasoning capabilities of large language models via constraint satisfaction problems. *arXiv preprint arXiv:2502.17848*.
- Jianghao Chen, Junhong Wu, Yangyifan Xu, and Jiajun Zhang. 2025b. Ladm: Long-context training data selection with attention-based dependency measurement for llms. *arXiv preprint arXiv:2503.02502*.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. 2025c. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2022. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. *Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning*. Preprint, arXiv:2501.12948.
- Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, and Tong Zhang. 2024. RLhf workflow: From reward modeling to online rlhf. *arXiv preprint arXiv:2405.07863*.
- Akash Ghosh, Debayan Datta, Sriparna Saha, and Chirag Agarwal. 2025. The multilingual mind: A survey of multilingual reasoning in language models. *arXiv preprint arXiv:2502.09457*.
- Xinyu Guan, Li Lyna Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, and Mao Yang. 2025. *rstar-math: Small llms can master math reasoning with self-evolved deep thinking*. Preprint, arXiv:2501.04519.
- 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. *arXiv preprint arXiv:2305.14992*.
- Mingui He, Yilun Liu, Shimin Tao, Yuanchang Luo, Hongyong Zeng, Chang Su, Li Zhang, Hongxia Ma, Daimeng Wei, Weibin Meng, et al. 2025. R1-t1: Fully incentivizing translation capability in llms via reasoning learning. *arXiv preprint arXiv:2502.19735*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2022. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*.
- 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*.
- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. 2025. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*.
- Pu Jian, Donglei Yu, and Jiajun Zhang. 2024. Large language models know what is key visual entity: An llm-assisted multimodal retrieval for vqa. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10939–10956.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Chen Li, Weiqi Wang, Jingcheng Hu, Yixuan Wei, Nanling Zheng, Han Hu, Zheng Zhang, and Houwen Peng. 2024. Common 7b language models already possess strong math capabilities. *arXiv preprint arXiv:2403.04706*.
- Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, and Brian Ichter. 2023. Chain of code: Reasoning with a language model-augmented code emulator. *arXiv preprint arXiv:2312.04474*.
- Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. 2022. Making large language models better reasoners with step-aware verifier. *arXiv preprint arXiv:2206.02336*.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yahui Zhou. 2024. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint arXiv:2410.18451*.
- Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, Lei Meng, Jiao Sun, et al. 2024. Improve mathematical reasoning in language models by automated process supervision. *arXiv preprint arXiv:2406.06592*.
- Arindam Mitra, Hamed Khanpour, Corby Rosset, and Ahmed Awadallah. 2024. Orca-math: Unlocking the potential of slms in grade school math. *arXiv preprint arXiv:2402.14830*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. True few-shot learning with language models. *Advances in neural information processing systems*, 34:11054–11070.
- Shuo Ren, Pu Jian, Zhenjiang Ren, Chunlin Leng, Can Xie, and Jiajun Zhang. 2025. Towards scientific intelligence: A survey of llm-based scientific agents. *arXiv preprint arXiv:2503.24047*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, YK Li, Yu Wu, and Daya Guo. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv e-prints*, pages arXiv–2408.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*.
- Fanqi Wan, Weizhou Shen, Shengyi Liao, Yingcheng Shi, Chenliang Li, Ziyi Yang, Ji Zhang, Fei Huang, Jingren Zhou, and Ming Yan. 2025. Qwenlong-11: Towards long-context large reasoning models with reinforcement learning. *arXiv preprint arXiv:2505.17667*.
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. 2024a. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9426–9439.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Zihan Wang, Yunxuan Li, Yuexin Wu, Liangchen Luo, Le Hou, Hongkun Yu, and Jingbo Shang. 2024b. Multi-step problem solving through a verifier: An empirical analysis on model-induced process supervision. *arXiv preprint arXiv:2402.02658*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024a. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Wen Yang, Kai Fan, and Minpeng Liao. 2024b. Markov chain of thought for efficient mathematical reasoning. *arXiv preprint arXiv:2410.17635*.

Wen Yang, Chong Li, Jiajun Zhang, and Chengqing Zong. 2023. Bigtranslate: Augmenting large language models with multilingual translation capability over 100 languages. *arXiv preprint arXiv:2305.18098*.

Fei Yu, Anningzhe Gao, and Benyou Wang. 2023a. Outcome-supervised verifiers for planning in mathematical reasoning. *arXiv preprint arXiv:2311.09724*.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhen-guo Li, Adrian Weller, and Weiyang Liu. 2023b. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*.

Dan Zhang, Sining Zhoubian, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. 2024. Rest-mcts\*: Llm self-training via process reward guided tree search. *arXiv preprint arXiv:2406.03816*.

Chujie Zheng, Zhenru Zhang, Beichen Zhang, Runji Lin, Keming Lu, Bowen Yu, Dayiheng Liu, Jingen Zhou, and Junyang Lin. 2024. [Processbench: Identifying process errors in mathematical reasoning](#). *Preprint*, arXiv:2412.06559.

Aojun Zhou, Ke Wang, Zimu Lu, Weikang Shi, Sichun Luo, Zipeng Qin, Shaoqing Lu, Anya Jia, Linqi Song, Mingjie Zhan, et al. 2023. Solving challenging math word problems using gpt-4 code interpreter with code-based self-verification. *arXiv preprint arXiv:2308.07921*.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

## A Appendix

### A.1 Prompt for 2-shot

#### System prompt

**<System>** You are a powerful agent with broad math knowledge and good at accurate calculation on math equations. Below is an instruction that describes a task. Continue to finish the response that appropriately completes the request within a maximum of 40 steps. When outputting each step, mark the sequence number of each step at the beginning, and explicitly state the final answer after the final step following the format 'The final answer is:'. After outputting the final answer only once, be sure to stop outputting.

#### Example 1

**<User>** Instruction: If the lengths of two sides of a right triangle are 5 and 12 units, what is the least possible length, in units, of the third side? Express your answer in simplest radical form.

**<Assistant>** Response: Let's think step by step.

Step 1: I know that the Pythagorean theorem relates the lengths of the sides of a right triangle by the equation  $a^2 + b^2 = c^2$ , where  $c$  is the hypotenuse and  $a$  and  $b$  are the legs.  
Step 2: Since I don't know which side is the hypotenuse, I'll try both possibilities and see which one gives me a smaller value for the third side.

Step 3: If I assume that the hypotenuse is 12, then the other leg must satisfy  $5^2 + b^2 = 12^2$ , or  $b^2 = 144 - 25 = 119$ .

Step 4: Taking the square root of both sides, I get  $b = \sqrt{119}$ , which is already in simplest radical form.

Step 5: If I assume that the hypotenuse is the unknown side, then it must satisfy  $5^2 + 12^2 = c^2$ , or  $c^2 = 25 + 144 = 169$ .

Step 6: Taking the square root of both sides, I get  $c = \sqrt{169} = 13$ .

Step 7: Comparing the two values, I see that  $\sqrt{119}$  is smaller than 13, since 119 is smaller than 169. The final answer is 119.



### Example 2

**<User>** Instruction: A square has sides of length 10, and a circle centered at one of its vertices has radius 10. What is the area of the union of the regions enclosed by the square and the circle? Express your answer in terms of  $\pi$ .

Response: Let's think step by step.

Step 1:I want to find the area of the shaded region in this picture, where the blue is the square and the red is the circle.

Step 2:I notice that the circle and the square share a quarter of the circle's area, which is  $\frac{1}{4}\pi r^2$ , where  $r = 10$ .

**<Assistant>** Step 3:So I can subtract that from the sum of the areas of the circle and the square to get the area of the union.

Step 4:The area of the circle is  $\pi r^2 = 100\pi$ , and the area of the square is  $s^2 = 100$ , where  $s = 10$ .

Step 5:So the area of the union is  $100\pi + 100 - \frac{1}{4}100\pi = 100 + \frac{3}{4}100\pi$ .

Step 6:The final answer is:  $100 + \frac{3}{4}100\pi$ .

## A.2 Step Distribution

We analyzed the step distribution of three datasets-PRM800k, MATH-Shepherd, and our Epic50k-by examining both the number of steps per solution and the length of each step. Regarding the number of steps per solution (as illustrated in Figure 7), our Epic50k dataset exhibits a concentration around 10 steps, whereas PRM800k and MATH-Shepherd are primarily clustered around 5 steps. This discrepancy is primarily attributed to the higher proportion of level 4-5 difficulty problems in Epic50k, which inherently necessitates more steps per solution. Subsequently, concerning the length of each step (shown in Figure 8), we observe that Epic50k and MATH-Shepherd exhibit step lengths predominantly ranging from 25 to 50 tokens. In contrast, PRM800k demonstrates a comparatively shorter step length, concentrated around approximately 25 tokens.

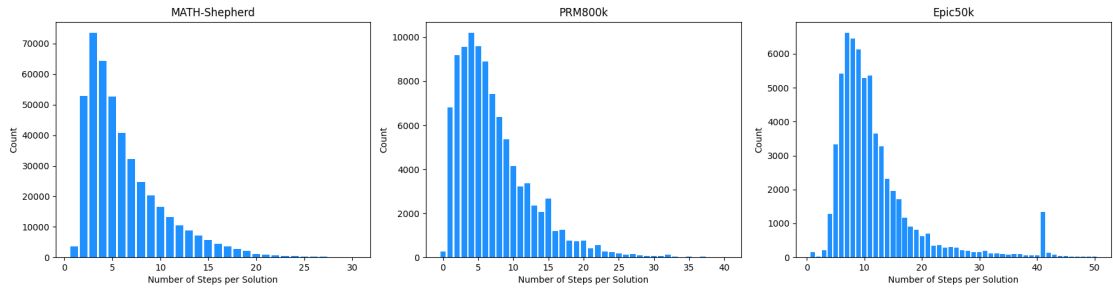


Figure 7: Distribution of steps per solution.

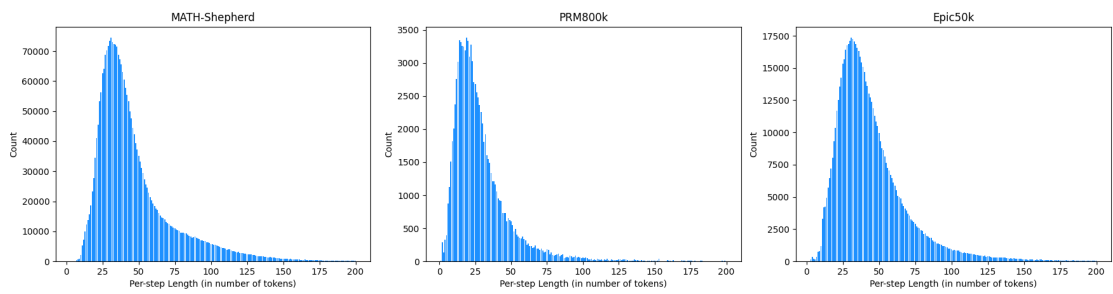


Figure 8: Length distribution of each step.