

Promoting Efficient Reasoning with Verifiable Stepwise Reward

Chuhuai Yue^{1,2}, Chengqi Dong^{1,3}, Yinan Gao⁴,
Hang He^{1,5}, Jiajun Chai¹, Wei Lin¹, Guojun Yin^{1*}

¹Meituan,

²Beijing Institute of Technology

³University of Science and Technology of China,

⁴Fudan University,

⁵East China Normal University

Abstract

Large reasoning models (LRMs) have recently achieved significant progress in complex reasoning tasks, aided by reinforcement learning with verifiable rewards. However, LRMs often suffer from overthinking, expending excessive computation on simple problems and reducing efficiency. Existing efficient reasoning methods typically require accurate task assessment to preset token budgets or select reasoning modes, which limits their flexibility and reliability. In this work, we revisit the essence of overthinking and identify that encouraging effective steps while penalizing ineffective ones is key to its solution. To this end, we propose a novel rule-based verifiable stepwise reward mechanism (VSRM), which assigns rewards based on the performance of intermediate states in the reasoning trajectory. This approach is intuitive and naturally fits the step-by-step nature of reasoning tasks. We conduct extensive experiments on standard mathematical reasoning benchmarks, including AIME24 and AIME25, by integrating VSRM with PPO and Reinforce++. Results show that our method achieves substantial output length reduction while maintaining original reasoning performance, striking an optimal balance between efficiency and accuracy. Further analysis of overthinking frequency and pass@k score before and after training demonstrates that our approach indeed effectively suppresses ineffective steps and encourages effective reasoning, fundamentally alleviating the overthinking problem.

Code — <https://github.com/lbenwu1/VSRM-Efficient-LRMs/tree/main>

Introduction

Large language models (LLMs) have rapidly emerged as powerful tools for tackling a wide range of general-purpose problems, achieving unprecedented performance in natural language processing and beyond. However, traditional LLMs struggle in tasks that require complex reasoning. Large reasoning models (LRMs), such as O1 (OpenAI 2024) and DeepSeek-R1 (DeepSeek-AI 2025), utilize test-time scaling techniques to allocate more computational resources during inference to generate longer CoTs, significantly enhancing their ability to decompose complex prob-

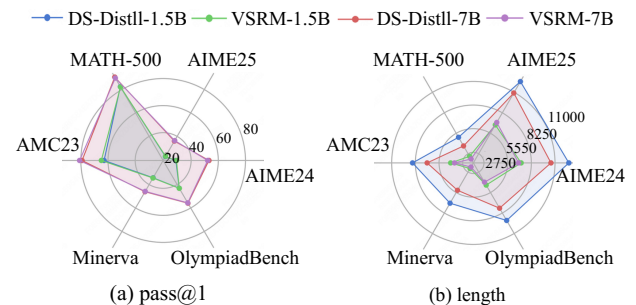


Figure 1: Proposed method substantially reduces token consumption across multiple benchmarks while maintaining model performance.

lems, solve them step-by-step, and explore multiple solutions. These advancements have led to substantial breakthroughs in coding and mathematical tasks (Chai et al. 2025; Xu et al. 2025b).

Nevertheless, LRMs are still far from being practically deployable in real-world applications due to their inability to adapt computational depth based on problem complexity. This often leads to the phenomenon of overthinking: allocating excessive computational resources to simple problems, leading to longer inference times, and sometimes even failing to ensure accuracy. This flaw stems from the training mechanism of reinforcement learning with verifiable rewards (RLVR)—for instance, the binary reward function used in R1 encourages the model to over-generate content in the pursuit of correctness, fostering a preference for *being overly redundant rather than potentially wrong*.

Existing efforts to optimize LRMs for efficient reasoning generally fall into two main categories. The first category involves setting output length constraints for tasks, essentially imposing a preset token budget. The latter, known as difficulty-adaptive reasoning, allocates reasoning modes based on the complexity of the problem. While this latter strategy appears more logical, it still inherently involves implicit token budgeting. Both methods necessitate pre-evaluating tasks to determine output targets, which not only depends heavily on the accuracy of problem anticipation but also lacks flexibility when addressing dynamic scenarios.

*Corresponding author

Rethinking the essence of the overthinking phenomenon, Chen et al. (2025) conduct extensive analysis and identify that a key manifestation of overthinking is the model spending substantial computation on intermediate steps that contribute little to final accuracy. Given that complex reasoning tasks inherently unfold in a step-by-step manner, we argue that the fundamental solution to overthinking lies in: encouraging intermediate steps that improve accuracy while penalizing the meaningless ones.

Building on the above insights, we propose a verifiable stepwise reward mechanism (VSRM) to address the overthinking problem in LRMs and enable more efficient reasoning. To provide timely reward signals throughout the reasoning trajectory, we discard the widely used result-based reward model in RLVR settings and instead adopt a stepwise reward framework. Traditional process-level reinforcement learning typically relies on training a process reward model (PRM). However, despite prior successes, PRMs often suffer from instability and limited accuracy.

To circumvent these issues, our VSRM design eliminates the need for PRMs by leveraging rule-based reward calculation, offering greater interpretability. Specifically, we segment the full reasoning trajectory generated during rollout (referred to as main rollout) into multiple sub-trajectories (referred to as sub-rollouts) of varying lengths, using special tokens that indicate reasoning steps as segmentation points. Each sub-rollout is treated as a prompt to generate multiple responses. Our proposed VSRM then assigns stepwise rewards at each truncation point based on the variation in correctness of sub-rollout responses, thereby encouraging/penalizing specific steps. We evaluate our approach on mathematical problems using extensive benchmarks and multiple base models. Experimental results in Figure 1 demonstrate that our method, besides its high interpretability, effectively addresses the essence of the overthinking problem, significantly compressing output length while maintaining or even surpassing the performance of base models. In summary, our contributions can be outlined as follows:

- We re-think the essence of the overthinking phenomenon and identified the root solution: encouraging beneficial intermediate steps while penalizing ineffective ones.
- We propose VSRM, which assigns verifiable stepwise rewards to the model in a rule-based manner, implementing encouragement and penalization of specific steps, fundamentally optimizing the model.
- Our extensive experiments on multiple diverse base models show that our method not only possesses high interpretability but also achieves significant output length compression while maintaining performance.

Related Work

Stimulating Reasoning Ability

The introduction of O1 marks a major advance in reasoning performance and the beginning of the LRM era, inspiring efforts to replicate such strong reasoning abilities. DeepSeek-R1, for example, achieves comparable results using a simple rule-based reward with the group relative policy optimization (GRPO)(Shao et al. 2024) algorithm, and

its open-source release has established RLVR as an effective paradigm for improving LLM reasoning. Subsequent models, including the Kimi K series(Kimi Team 2025a,b), QwQ(Qwen Team 2025), and O3(OpenAI 2025), further advance these capabilities. RLVR assigns scores to trajectories based on pre-designed rules, rewarding desirable behaviors and penalizing undesirable ones. This encourages models to generate long CoTs to maximize correctness, fostering advanced reasoning behaviors such as search and backtracking(Fu et al. 2025; Lin et al. 2025). However, this also leads to a preference for redundancy over risk of error, resulting in overthinking—wasting computational resources and potentially impairing performance, thus limiting the practical applicability of LRMs.

Attempts Towards Efficient Reasoning

Overthinking issue is first identified and analyzed by Chen et al. (2025), who observe that LRMs generate lengthy outputs that neither improve accuracy nor introduce new solution strategies especially for easy prompts. To address this, various works explore efficient reasoning from different angles. Huang et al. (2025b), Xu et al. (2025a), and Han et al. (2025) focus on prompt engineering to encourage concise responses, while Yang et al. (2025b) and Fan et al. (2025) target decoding strategies to suppress unnecessary CoT extensions. Although these training-free approaches can alleviate overthinking, they cannot fundamentally eliminate the problem. Ma et al. (2025), Ye et al. (2025), and Yu et al. (2025) employ supervised fine-tuning (SFT) on well-curated datasets to help models memorize efficient reasoning patterns, while Zhang et al. (2025b) and Luo et al. (2025a) explicitly incorporate efficiency as an optimization objective.

A more mainstream approach is to use reinforcement learning for better generalization. Kimi Team (2025a), Shen et al. (2025), Yeo et al. (2025), and Cheng et al. (2025) introduce length penalties in the reward function to suppress overly long reasoning traces. Arora et al. (2025) and Zhang et al. (2025) use intra-group length regularization to encourage shorter correct answers. Yi et al. (2025), Hou et al. (2025), and Qi et al. (2025) optimize performance under a fixed token budget to balance efficiency and effectiveness. Approaches such as Zhang et al. (2025a), Huang et al. (2025a), and Wu et al. (2025) assign predefined thinking patterns based on task difficulty, which essentially reflects a length budget. All these methods heavily rely on presetting a rough length estimation for different problems and optimize the model by comparing the gap between this estimation and the actual rollout. Such estimation lacks reliability and flexibility. Moreover, whether explicitly or implicitly, incorporating length compression as part of the optimization objective may potentially impair the model’s reasoning capability.

Our work revisits the essence of overthinking and, leveraging the sequential nature of reasoning tasks, encourages effective steps while penalizing ineffective ones through stepwise rewards. The proposed method is highly interpretable and promotes efficient reasoning without compromising performance.

Methodology

Rethink the Essence of Overthinking Issue

Based on the analysis in the previous section, in order to enable LRMs to perform efficient reasoning elegantly, we first need to rethink the essence of the overthinking issue.

Chen et al. (2025) use QwQ and R1 to generate responses for multiple evaluation benchmarks and observe a general phenomenon of “one problem, multiple solutions,” especially for simpler questions. To gain a finer-grained understanding of overthinking and more precisely locate its root cause, we build upon this observation and take a step further.

Reasoning tasks—particularly in mathematical reasoning and code generation—typically require step-by-step answers. Leveraging this, we use DeepSeek-distilled-Qwen-1.5B (DeepSeek-AI 2025) and DeepScaleR-1.5B (Luo et al. 2025b) (referred to as DS-Distill-1.5B and DeepScaleR, respectively) to generate answers for MATH-500 (Lightman et al. 2023) problems. From the 500 generated responses, we randomly sample 10% and annotate the reasoning steps for a case study. We find that the reasoning process of LRMs can generally be divided into three stages: problem restatement, step-by-step answering, and answer summary. The overthinking issue mainly occurs during the step-by-step answering stage. As shown in Figure 2, different steps are highlighted in different colors. In the example, the response is excessively verbose, expending substantial computational resources on a simple question—how many integers less than 0 are there in $[-500, 0]$ —and ultimately arriving at an incorrect conclusion. Many clearly ineffective steps neither improve accuracy nor help the model reach the correct answer. Such cases are not isolated in our study.

Model	Detected Overthinking
DeepScaleR	279/500
DS-Distill-1.5B	312/500

Table 1: Detecting ineffective steps via DeepSeek-R1

To obtain a more comprehensive view, we further prompt DeepSeek-R1 to analyze all 500 responses generated by both models and statistically examine the occurrence of ineffective steps. The prompt used is provided in the Appendix, and results are summarized in Table 1. Comparing DeepSeek-R1’s predictions with our 50 manual annotations, we find that although it does not identify all ineffective steps precisely, it is able to detect overthinking phenomena, making the statistics generally trustworthy.

Based on the analysis of R1, both models exhibit significant ineffective steps in over half of the samples. This indicates that, in general scenarios, the generation of numerous ineffective steps by the models constitutes the primary manifestation of the overthinking phenomenon. Therefore, we propose that the key to alleviating overthinking is to accurately distinguish between effective and ineffective steps, and to set the optimization objective as encouraging effective steps and penalizing ineffective ones, thereby achieving efficient reasoning. Compared to directly compressing

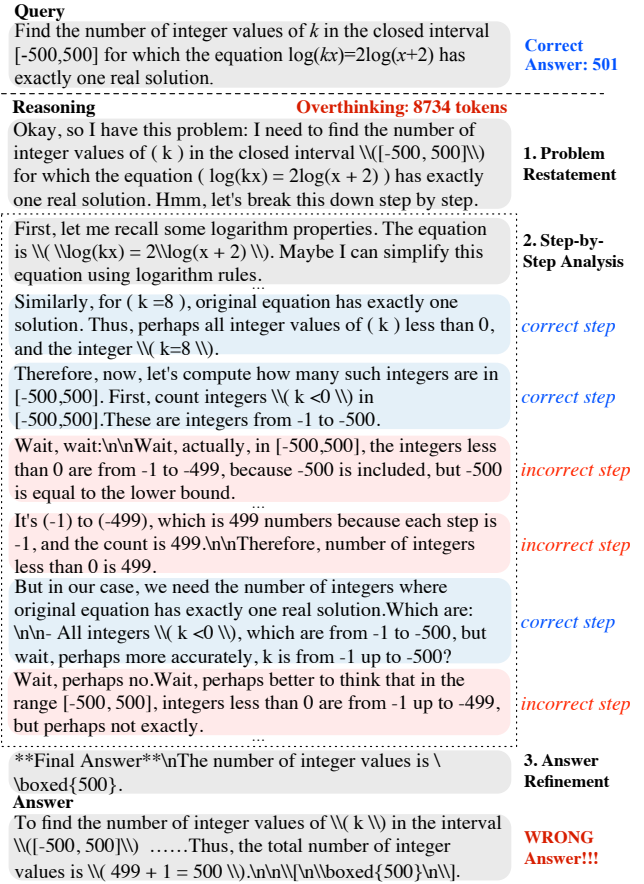


Figure 2: Illustration of the step-by-step answering process by DeepScaleR, with *correct* and *incorrect* steps highlighted in different colors.

the length of the CoT, this is more intuitive and is closely aligned with the step-by-step nature of reasoning tasks.

Make Stepwise Reward Verifiable

The most commonly used outcome-based reward mechanism in RLVR is incompatible with our optimization objective mentioned above as it cannot provide timely intermediate rewards at the reasoning process’s intermediate states to reward or penalize different steps. To address this, we propose to apply a stepwise reward mechanism, which naturally aligns with our optimization objective.

A common approach to implementing stepwise reward is to train a PRM that decomposes the reasoning CoT and assigns a score to each stage, guiding the model to generate more interpretable and rigorous reasoning processes (Lightman et al. 2023). This approach has achieved promising results in tasks like mathematical reasoning. However, PRMs suffer from significant limitations in large-scale reinforcement learning scenarios: they are difficult to train and their reliability cannot be guaranteed (DeepSeek-AI 2025). This dilemma motivates us to design a verifiable stepwise reward mechanism, which combines the advantages of

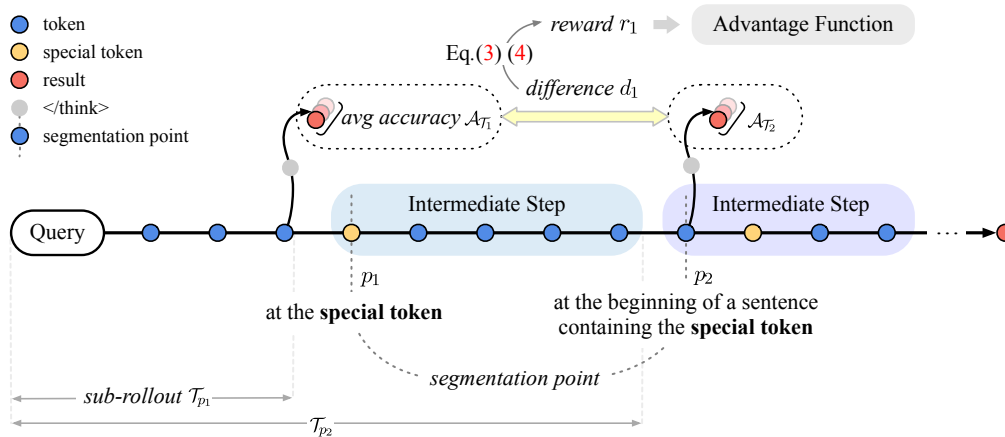


Figure 3: For a given main rollout \mathcal{T} , the process of obtaining its complete reward information involves the following steps: segmenting \mathcal{T} into sub-rollouts; generating multiple candidate answers for each sub-rollout; assigning stepwise rewards based on the average accuracy of all sub-rollouts via proposed method.

stepwise rewards with the reliability and usability of verifiable outcome-based rewards. In this way, we circumvent the shortcomings of PRMs while providing interpretable stepwise rewards for each reasoning step.

Step Separation. The first phase of implementing VSRM is the separation of reasoning steps. Chen et al. (2025) and Cheng et al. (2025) attempt to prompt or train an additional model specifically to identify different solutions or correct answers in the CoT. However, introducing another model inevitably brings additional uncertainty. Therefore, we design a comprehensive rule-based step segmentation algorithm.

Through the analysis of the generated responses mentioned above, we observe that the boundaries between reasoning steps are often accompanied by certain special tokens, such as “however”, “thus”, “so”, “but” “wait”, etc. These tokens indicate that the model has completed one reasoning step and is about to proceed to the next, whether it is a progression or a transition. (Wang et al. 2025) also reported similar phenomena in their work and conducted detailed analysis, which verifies the correctness of our observation. This demonstrates that extracting special tokens from the CoT and using them for step segmentation is feasible.

Specifically, for a complete rollout generated by reference model during reinforcement learning, which will be denoted as main rollout \mathcal{T} in the following, we first use regular expressions to extract the reasoning part (the content within `<think>...</think>`). After obtaining the full thinking content, we skip a certain number of initial tokens to avoid splitting the problem restatement part. During segmentation, we use a predefined list of special tokens to divide \mathcal{T} . To ensure readability, we introduce two additional rules: (1) there must be at least a minimum distance between adjacent segmentation points, and (2) each segmentation point should align with the beginning of a complete sentence containing the special token. These rules help avoid overly dense segmentation when multiple special tokens appear in a single sentence, and ensure that each resulting step contains complete semantic information. The entire step segmenta-

tion procedure, as illustrated in Figure 3, results in a set of segmentation points $\mathcal{P} = \{p_1, p_2, \dots, p_K\}$. Further details can be found in the Appendix.

Assigning Reward to Intermediate Stage. Based on \mathcal{P} , we can extract the reasoning content from the start up to each p_i and combine it with the input question. In this way, a set of sub-rollouts $\mathcal{T}_{\mathcal{P}} = \{\mathcal{T}_{p_1}, \mathcal{T}_{p_2}, \dots, \mathcal{T}_{p_k}\}$ are derived from \mathcal{T} , each representing a distinct segment of the overall reasoning process. Thanks to the rules established during step segmentation, each \mathcal{T}_{p_i} retains complete semantic information and does not contain incomplete sentences. We append a `</think>` token to every \mathcal{T}_{p_i} to mark the end of the reasoning and then use them as new queries for the model to generate corresponding answers. In this way, the reward at position p_i in \mathcal{T} can be directly evaluated by the outcome of \mathcal{T}_{p_i} . This effectively transforms the challenge of assigning stepwise rewards to intermediate states into a verifiable and explainable procedure, where the reward for each sub-rollout serves as a proxy for the reward of the intermediate state of the main rollout.

The simplest approach to assigning reward signals to each \mathcal{T}_{p_i} is to adopt the commonly used binary reward mechanism, where a correct answer receives a reward of 1 and an incorrect answer receives 0. To ensure the stability of the reward signal and reduce the high variance caused by randomness, we generate multiple candidate answers for each \mathcal{T}_i , and use the average accuracy as its reward, denoted as $\mathcal{A}_{\mathcal{T}_i} \in \mathcal{A}_{\mathcal{T}}$. For a set of N candidate answers generated from \mathcal{T}_i , $\mathcal{A}_{\mathcal{T}_i}$ is empirically estimated as:

$$\mathcal{A}_{\mathcal{T}_i} = \frac{1}{N} \sum_{j=i}^N \mathbb{I}(\text{IsCorrect}(\text{LRM}(\mathcal{T}_i)_j)) \quad (1)$$

where $\mathbb{I}(\cdot)$ is the indicator function.

However, it is suboptimal to directly assign \mathcal{A}_i as the stepwise reward for each step, since the reward at every position would always be non-negative, making it ineffective in penalizing ineffective steps. Therefore, we instead use the dif-

ference in average accuracy between adjacent steps as the reward for each step, as follows:

$$d_{i-1} = \mathcal{A}_i - \mathcal{A}_{i-1}, \quad \text{for } i \in [1, \dots, k] \quad (2)$$

For the boundary case r_k , we compute the difference between the average accuracy of \mathcal{T} (denoted as $\mathcal{A}_{\mathcal{T}}$) and \mathcal{A}_k .

Assigning the stepwise reward as the difference in average accuracy between consecutive steps encourages the main rollout to evolve towards better states. In this scheme, steps that contribute to an improvement in accuracy are rewarded, while the rest are penalized or suppressed. However, in real-world reasoning processes, arriving at the correct conclusion often requires the accumulation of several intermediate steps, rather than relying on a single decisive action. Consequently, under this difference-based reward assignment, \mathcal{A}_i tends to remain unchanged across many steps, resulting in $d_i = 0$ at most positions. This leads to sparse reward signals, thus hindering the training process.

To address this issue, when the difference in average accuracy between consecutive steps is negligible, i.e., $|d_{i-1}| < \epsilon$ for a small threshold ϵ , we further examine subsequent steps to determine if a significant change in accuracy occurs within a limited lookahead window of L_{\max} steps. Specifically, let q be the smallest integer such that $q \leq L_{\max}$ and $|d_{i+q-1}| \geq \epsilon$. If such a q exists, we assign a discounted reward to step $i - 1$ as follows:

$$r_{i-1} = \text{sgn}(d_{i+q-1}) \cdot |d_{i+q-1}| \cdot \gamma^q, \quad (3)$$

where $\gamma \in (0, 1)$ is a discount factor and $\text{sgn}(\cdot)$ is an indicator function. If no significant change is found within L_{\max} steps, we set $r_{i-1} = 0$. For the case where the end of the sequence is reached without observing a significant change, we use the difference between the final accuracy and the current accuracy, also discounted appropriately:

$$r_{i-1} = \text{sgn}(d_k) \cdot |d_k| \cdot \gamma^{k-i+1}. \quad (4)$$

This decay-based propagation mechanism alleviates the excessive sparsity of reward signals, enabling the model to more rapidly cultivate efficient reasoning behaviors. More details can be found in the Appendix.

Reinforcement Learning with VSRM

Through our proposed VSRM, we are now able to assign stepwise rewards to each main trajectory generated during RL training. We combine the traditional binary outcome reward, the format reward, and stepwise reward to ensure that the LRM’s outputs conform to verifiable standards, avoid uncontrolled content generation, and fundamentally guarantee result correctness. Ultimately, we obtain a reward list whose length matches the total response length, which is used for subsequent advantage estimation. The overall reward can be expressed as:

$$R_{\mathcal{T}} = [\dots, r_1, \dots, r_t, \dots, r_k, \dots, r_{\mathcal{T}}^{\text{result}} + r_{\mathcal{T}}^{\text{format}}] \quad (5)$$

where r_t denotes the stepwise reward, r_t^{result} is the outcome reward (assigned only at the final position) and r_t^{format} is the format reward (also assigned only at the final position). For other positions, reward remains zero.

Our proposed VSRM is fully compatible with RL algorithms that support stepwise rewards, such as PPO(Schulman et al. 2017) and Reinforce++(Hu 2025), and can be seamlessly integrated. In the next section, we will discuss the effectiveness of VSRM through experiments.

Experiments

Experimental Setup

To thoroughly evaluate the effectiveness of our approach, we conduct experiments using several widely adopted LRMs as base models, including DS-Distill-1.5B, DS-Distill-7B(DeepSeek-AI 2025), and DeepScaleR(Luo et al. 2025b). We employ the VeRL(Sheng et al. 2024) framework and apply both PPO and Reinforce++ algorithms, which are naturally compatible with stepwise reward, for post-training of the base models. For training, we use the dataset introduced by DeepScaleR, which is compiled from historical AIME problems, AMC problems, the Omni-MATH dataset, and the Still dataset, covering a wide range of difficulty levels.

In terms of evaluation, we follow previous work and select a comprehensive set of benchmarks, including MATH-500, AIME24, AIME25(Art of Problem Solving 2024), AMC23(Mathematical Association of America 2023), Minerva(Lewkowycz et al. 2022), and OlympiadBench(Chaoqun et al. 2024), providing broader coverage than previous studies. During evaluation, we set the maximum generation length at 20,000 tokens—substantially higher than the maximum response length during training—to ensure that model performance is not constrained by this hyperparameter. The temperature and top-p are set to 0.6 and 0.95, respectively, consistent with Hou et al. (2025)’s work and DeepSeek-R1. For all benchmarks, we report the average pass@1 over N runs. Specifically, for MATH-500 and OlympiadBench, where the benchmark sizes are relatively large, we set N to 16; for the other benchmarks, we set N to 32 to reduce randomness.

Additional details about implementation can be found in the Appendix.

Main Results

Table 2 presents a comparison of VSRM with several related works across six benchmarks, organized into three sections according to the base model used: DS-Distill-1.5B, DS-Distill-7B, and DeepScaleR. The compared methods include LC-R1(Cheng et al. 2025), ThinkPrune(Hou et al. 2025), AdaptThink(Zhang et al. 2025a), Efficient-Reasoning(Arora et al. 2025), and GPRO-LEAD(Zhang et al. 2025).

For experiments with DS-Distill-1.5B as the base model, our method consistently achieve the best or second-best performance across all benchmarks, while largely preserving or even slightly improving the original model’s capabilities on tasks such as AIME24, AIME25, AMC, and Minerva. Efficient Reasoning produce results comparable to ours on several benchmarks and achieve the best score on OlympiadBench; however, its output length was much longer, indicating a lack of efficiency. AdaptThink achieve slightly better output length compression than our method on MATH-500,

Method	AIME24		AIME25		MATH-500		AMC23		Minerva		OlympiadBench	
	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓
DS-Distill-1.5B	28.8	12605	22.8	12444	82.4	4960	62.9	8577	34.2	6929	43.1	9249
+ LC-R1	26.3 ↓2.5	7098 ↓5507	20.9 ↓1.9	6942 ↓5502	80.4 ↓2.0	2473 ↓2487	62.8 ↓0.1	4389 ↓4188	32.8 ↓1.4	2512 ↓4417	42.0 ↓1.1	4632 ↓4617
+ ThinkPrune-2k	26.7 ↓2.1	7085 ↓5520	19.7 ↓2.1	6918 ↓5526	81.7 ↓0.7	2426 ↓2534	63.8 ↑0.9	4224 ↓4353	33.9 ↓0.3	2667 ↓4262	42.9 ↓0.2	4752 ↓4497
+ AdaptThink	28.7 ↓0.1	8055 ↓4550	21.8 ↓1.0	8155 ↓4289	80.4 ↓2.0	1963 ↓2997	63.3 ↑0.4	4069 ↓4508	32.3 ↓1.9	1912 ↓5017	42.6 ↓0.5	4563 ↓4686
+ Efficient Reasoning	<u>29.2</u> ↑0.4	9189 ↓3416	<u>22.9</u> ↑0.1	8590 ↓3854	<u>82.0</u> ↓0.4	2621 ↓2339	<u>64.7</u> ↑1.8	5202 ↓3375	<u>34.4</u> ↑0.2	3230 ↓3699	43.8 ↑0.7	5755 ↓3494
+ VSRM-PPO (ours)	<u>29.2</u> ↑0.4	<u>7065</u> ↓5540	23.0 ↑0.2	6621 ↓5823	82.2 ↓0.2	<u>2400</u> ↓2560	64.9 ↑2.0	4153 ↓4424	<u>34.4</u> ↑0.2	2273 ↓4656	<u>43.2</u> ↑0.1	4499 ↓4750
+ VSRM-R++ (ours)	29.5 ↑0.7	6958 ↓5647	<u>22.9</u> ↑0.1	<u>6892</u> ↓5552	81.7 ↓0.7	2597 ↓2363	<u>64.7</u> ↑1.8	<u>4119</u> ↓4458	34.6 ↑0.4	<u>2172</u> ↓4751	43.1 ↑0.0	<u>4544</u> ↓4705
DS-Distill-7B	53.1	10529	36.0	10909	90.6	3796	79.1	6884	46.1	5195	55.9	7602
+ LC-R1	51.7 ↓1.4	6820 ↓3809	35.7 ↓0.3	7458 ↓3451	88.3 ↓2.3	1776 ↓2020	79.1 ↑0.0	3686 ↓3198	44.4 ↓1.7	1834 ↓3361	55.5 ↓0.4	4193 ↓3409
+ AdaptThink	<u>52.1</u> ↓1.0	8679 ↓1830	35.0 ↓1.0	9807 ↓1102	88.9 ↓1.7	2199 ↓1597	80.7 ↑1.6	5130 ↓1754	45.2 ↓0.9	2869 ↓2326	55.4 ↓0.5	5915 ↓1687
+ Efficient Reasoning	51.9 ↓1.2	8667 ↓1862	<u>36.2</u> ↑0.2	9100 ↓1809	89.8 ↓0.8	2408 ↓1388	80.7 ↑1.6	4933 ↓1951	45.7 ↓0.4	2903 ↓2292	<u>55.6</u> ↓0.3	5599 ↓2003
+ GRPO-LEAD	<u>52.1</u> ↓1.0	9023 ↓1506	36.1 ↑0.1	9842 ↓1067	89.5 ↓1.1	3152 ↓644	80.7 ↑1.6	5584 ↓1300	46.3 ↑0.2	3930 ↓1265	56.1 ↑0.2	6372 ↓1230
+ VSRM-PPO (ours)	51.8 ↓1.3	6675 ↓3854	<u>36.2</u> ↑0.2	6920 ↓3989	<u>89.6</u> ↓1.0	<u>2024</u> ↓1772	81.0 ↑1.9	3759 ↓3125	45.5 ↓0.6	<u>2002</u> ↓3193	55.3 ↓0.6	4168 ↓3434
+ VSRM-R++ (ours)	52.2 ↓0.9	<u>6773</u> ↓3756	36.4 ↑0.4	<u>6953</u> ↓3956	89.8 ↓0.8	2044 ↓1752	<u>80.9</u> ↑1.8	<u>3704</u> ↓3180	<u>45.8</u> ↓0.3	2068 ↓3127	<u>55.6</u> ↓0.3	4103 ↓3499
DeepScaleR	40.7	8515	31.3	8029	86.5	3105	73.6	5487	38.9	5102	50.0	5647
+ ThinkPrune-2k	36.5 ↓4.2	6108 ↓2407	26.0 ↓5.3	5486 ↓2543	84.5 ↓2.0	1946 ↓1159	71.8 ↓1.8	3346 ↓2141	37.9 ↓1.0	2188 ↓2914	48.1 ↓1.9	3323 ↓2324
+ VSRM-PPO (ours)	<u>39.2</u> ↓1.4	6744 ↓1771	31.3 ↑0.0	6133 ↓1896	<u>85.5</u> ↓1.0	2391 ↓714	<u>73.5</u> ↓0.1	4074 ↓1413	39.5 ↑0.6	<u>2231</u> ↓2871	50.1 ↑0.1	4414 ↓1233
+ VSRM-R++ (ours)	39.3 ↓1.4	<u>6668</u> ↓1847	<u>30.9</u> ↓0.4	<u>6103</u> ↓1926	85.8 ↓0.7	<u>2275</u> ↓830	74.7 ↑1.1	4082 ↓1405	<u>38.4</u> ↓0.5	2319 ↓2783	<u>49.4</u> ↓0.6	<u>4389</u> ↓1258

Table 2: Extensive experiments are conducted on several widely used benchmarks across multiple mathematical domains, comparing our method with several related approaches. Here, VSRM-PPO refers to training the base model by integrating our proposed VSRM with the PPO algorithm, while VSRM-R++ denotes the combination of VSRM with the Reinforce++ algorithm.

AMC23, and Minerva, but this high level of conciseness restrict the model’s ability to explore different solution strategies. As a result, AdaptThink struggle in terms of overall performance and consistently lag behind our method. As for the other methods, both their performance and efficiency are generally less competitive compared to VSRM.

A similar trend is observed when using DS-Distill-7B as the base model. Both GRPO-LEAD and Efficient Reasoning adopt a strategy of sacrificing less compression to achieve better performance, resulting in performance levels comparable to ours. However, their efficiency is much lower than that of VSRM, with only minimal reductions in token count—the smallest among all evaluated methods. These two sets of experiments together indicate that, on distilled models, our method not only alleviates the overthinking problem but also has the potential to slightly improve overall model performance. Other methods also fall short of our approach when considering both performance and efficiency in a comprehensive manner.

DeepScaleR has already undergone systematic and thorough reinforcement learning, and its performance is nearly saturated. Consequently, further fine-tuning on this model often results in performance degradation, which explains why few related works use DeepScaleR as a base model. However, in practical model training pipelines, such as for Qwen3 (Yang et al. 2025a), training is typically conducted in multiple stages. Thus, efficiency optimization must ultimately target already strong LLMs that have been reinforced. In this work, we use DeepScaleR as the base model for our experiments. Results show that VSRM compresses output length across all benchmarks while maintaining performance close to the original DeepScaleR. By contrast, although ThinkPrune-2k achieves greater output length reduc-

tion, it compromises DeepScaleR’s strong reasoning ability, which runs counter to the original motivation of addressing the overthinking problem.

Across all experimental results, our method consistently reduces redundant output and mitigates the overthinking problem for models of various sizes and types, while preserving the original reasoning capability. Among all compared methods, VSRM achieves an excellent balance between performance and efficiency, demonstrating its superiority in fundamentally addressing the overthinking issue.

Ablation Study

In this section, We design ablation experiments to verify the contribution of the individual components. All experiments are trained for the same number of steps to ensure fairness, and the results are shown in Table 3.

First, we perform standard PPO-based reinforcement learning to the base model. We observe some improvement in performance on certain benchmarks, along with a slight reduction in output length, likely because the format constraints of RL reward discourage excessively long outputs.

To explore the impact of different outcome reward designs, we adopt the length-penalized reward function proposed by Yeo et al. (2025), instead of the simple binary outcome reward. The results show that explicit length penalties do not further compress output length as expected; in fact, the compression is less effective than with VSRM-PPO. This is likely because the more complex outcome reward interferes with the effectiveness of the process-level rewards, leading to suboptimal results.

Finally, we evaluate the impact of the propagation decay mechanism in VSRM, which is designed to deliver denser reward signals and guide the model toward correct reasoning

Method	AIME24		AIME25		MATH-500		AMC23		Minerva		OlympiadBench	
	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓	pass@1↑	Length↓
DS-Distill-1.5B	28.8	12605	22.8	12444	82.4	4960	62.9	8577	34.2	6929	43.1	9249
Original PPO	29.0 ↑0.2	11111 ↓1494	22.8 ↑0.0	10545 ↓1899	82.3 ↓0.1	3968 ↓992	63.9 ↑1.0	7294 ↓1283	34.4 ↑0.2	5205 ↓1724	43.4 ↑0.3	7927 ↓1322
VSRM w/ penalty	28.6 ↓0.2	7533 ↓5072	20.9 ↓1.9	7369 ↓5075	80.9 ↓1.5	2615 ↓2345	64.1 ↑1.2	4695 ↓3882	34.2 ↑0.0	2712 ↓4217	42.9 ↓0.2	5102 ↓4147
VSRM w/o decay	28.8 ↑0.0	8638 ↓3967	22.6 ↓0.2	8451 ↓3993	81.1 ↓1.3	2990 ↓1970	63.4 ↑0.5	5618 ↓2959	34.1 ↓0.1	3288 ↓3641	43.1 ↑0.0	5935 ↓3314
VSRM	29.2 ↑0.4	7065 ↓5540	23.0 ↑0.2	6621 ↓5823	82.2 ↓0.2	2400 ↓2560	64.9 ↑2.0	4153 ↓4424	34.4 ↑0.2	2273 ↓4656	43.2 ↑0.1	4499 ↓4750

Table 3: Ablation experiments for our proposed VSRM(-PPO). We investigate the impact of each component by incrementally adding them to the original PPO.

paths more efficiently. Results show that while simple differential rewards can serve as stepwise feedback, their ability to suppress output length and maintain performance is significantly limited without propagation decay. This highlights the role of propagation decay in providing finer-grained supervision, enabling the model to better distinguish between effective and ineffective steps.

These ablation experiments illustrate the relationships and roles of each component of our proposed method, and demonstrate the rationality and effectiveness of VSRM.

Analysis on Reasoning Behavior

Model	Detected Overthinking
VSRM-PPO(DeepScaleR)	132/500
VSRM-PPO(DS-Distill-1.5B)	126/500

Table 4: Detecting ineffective steps via DeepSeek-R1

Through the above experimental analysis, the effectiveness of our proposed method has been thoroughly validated. To further investigate the impact of VSRM on reasoning patterns, we conduct an empirical analysis in this section.

First, we revisit the example in Figure 2, where DeepScaleR repeatedly self-reflects on the number of negative integers in $[-500, 0)$ and ultimately reaches an incorrect answer, consuming many tokens in the process. We examined all 16 outputs generated by the model after VSRM-PPO fine-tuning for this problem. As long as there are no errors in the first half of reasoning, the model no longer repeatedly self-reflects and quickly summarizes the correct answer. Representative responses are shown in the Appendix. To obtain broader and more robust conclusions, we again use R1 to detect overthinking in models trained with VSRM. As shown in Table 4, overthinking is notably reduced compared to the base models (Table 1). Since MATH-500 contains some challenging (level 5) problems, some self-reflection is expected. Overall, VSRM effectively reduces ineffective steps, shortens output length, and fundamentally alleviates the overthinking issue. The above experiments demonstrate the effectiveness of VSRM in suppressing ineffective steps. Equally important, however, is its encouragement of effective steps, which we assess by evaluating the model’s exploration capability on challenging benchmarks.

To this end, we select four relatively difficult benchmarks—AIME24, AIME25, Minerva, and OlympiadBench—and compare DS-Distill-1.5B with its VSRM-PPO




Figure 4: The increase in performance as the number of k grows reflects the model’s diversity in exploring different problem-solving approaches.

fine-tuned counterpart. We examine the change in the pass@k metric as k increases. The pass@k metric reflects the model’s ability to explore multiple approaches and strategies to eventually arrive at the correct answer. If the model’s exploration ability is restricted, pass@k will not show significant improvement as k increases; conversely, a steady increase indicates that the model’s exploration capacity is preserved. As shown in Figure 4, both the base model and VSRM-enhanced model exhibit a stable increase in pass@k as k grows. Moreover, the trend for our model is consistent with, and in some cases even surpasses, that of the base model. This demonstrates that VSRM does not hinder the model’s ability to explore valuable reasoning paths; on the contrary, it encourages the exploration of effective steps.

Together, these experiments provide strong evidence that VSRM not only effectively suppresses overthinking, but also preserves the model’s exploration ability, achieving output conciseness without sacrificing reasoning performance.

Conclusion

In this work, we address overthinking problem in LRMs by setting the optimization objective to encourage effective intermediate steps and penalize ineffective ones. To this end, we introduce a verifiable stepwise reward mechanism that evaluates and rewards reasoning steps based on their effectiveness. Experiments on multiple benchmarks show that our method significantly reduces overthinking while maintaining or slightly improving reasoning performance. We hope to inspire further advances in efficient reasoning and practical LRM deployment.

References

- Arora et al. 2025. Training Language Models to Reason Efficiently. arXiv:2502.04463.
- Art of Problem Solving. 2024. American Invitational Mathematics Examination. https://artofproblemsolving.com/wiki/index.php/American_Invitational_Mathematics_Examination. Accessed: 2025-03-28.
- Chai, J.; Yin, G.; Xu, Z.; Yue, C.; Jia, Y.; Xia, S.; Wang, X.; Jiang, J.; Li, X.; Dong, C.; He, H.; and Lin, W. 2025. RLFactory: A Plug-and-Play Reinforcement Learning Post-Training Framework for LLM Multi-Turn Tool-Use. arXiv:2509.06980.
- Chaoqun, H.; Renjie, L.; Yuzhuo, B.; Shengding, H.; Zhen, T.; Junhao, S.; Jinyi, H.; Xu, H.; Yujie, H.; Yuxiang, Z.; Jie, L.; Lei, Q.; Zhiyuan, L.; and Maosong, S. 2024. Olympiad-Bench: A Challenging Benchmark for Promoting AGI with Olympiad-Level Bilingual Multimodal Scientific Problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3828–3850. Bangkok, Thailand: Association for Computational Linguistics.
- Chen, X.; Xu, J.; Liang, T.; He, Z.; Pang, J.; Yu, D.; Song, L.; Liu, Q.; Zhou, M.; Zhang, Z.; Wang, R.; Tu, Z.; Mi, H.; and Yu, D. 2025. Do NOT Think That Much for $2+3=?$: On the Overthinking of Long Reasoning Models. In *Forty-second International Conference on Machine Learning*.
- Cheng, Z.; Chen, D.; Fu, M.; and Zhou, T. 2025. Optimizing Length Compression in Large Reasoning Models. arXiv:2506.14755.
- DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. arXiv:2501.12948.
- Fan, C.; Zhang, Y.; Jia, J.; Hero, A.; and Liu, S. 2025. CycliCReflEx: Improving Large Reasoning Models via Cyclical Reflection Token Scheduling. arXiv:2506.11077.
- Fu, Y.; Chen, T.; Chai, J.; Wang, X.; Tu, S.; Yin, G.; Lin, W.; Zhang, Q.; Zhu, Y.; and Zhao, D. 2025. SRFT: A Single-Stage Method with Supervised and Reinforcement Fine-Tuning for Reasoning. arXiv:2506.19767.
- Han, T.; Wang, Z.; Fang, C.; Zhao, S.; Ma, S.; and Chen, Z. 2025. Token-Budget-Aware LLM Reasoning. arXiv:2412.18547.
- Hou, B.; Zhang, Y.; Ji, J.; Liu, Y.; Qian, K.; Andreas, J.; and Chang, S. 2025. ThinkPrune: Pruning Long Chain-of-Thought of LLMs via Reinforcement Learning. arXiv:2504.01296.
- Hu, J. 2025. Reinforce++: A simple and efficient approach for aligning large language models. *arXiv preprint arXiv:2501.03262*.
- Huang, S.; Wang, H.; Zhong, W.; Su, Z.; Feng, J.; Cao, B.; and Fung, Y. R. 2025a. AdaCtrl: Towards Adaptive and Controllable Reasoning via Difficulty-Aware Budgeting. arXiv:2505.18822.
- Huang, X.; Vangani, T. K.; Liu, Z.; Zou, B.; and Aw, A. T. 2025b. AdaCoT: Rethinking Cross-Lingual Factual Reasoning through Adaptive Chain-of-Thought. arXiv:2501.16154.
- Kimi Team. 2025a. Kimi k1.5: Scaling Reinforcement Learning with LLMs. arXiv:2501.12599.
- Kimi Team. 2025b. Kimi K2: Open Agentic Intelligence. <https://moonshotai.github.io/Kimi-K2/>.
- Lewkowycz, A.; Andreassen, A. J.; Dohan, D.; Dyer, E.; Michalewski, H.; Ramasesh, V. V.; Slone, A.; Anil, C.; Schlag, I.; Gutman-Solo, T.; Wu, Y.; Neyshabur, B.; Gur-Ari, G.; and Misra, V. 2022. Solving Quantitative Reasoning Problems with Language Models. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Lightman, H.; Kosaraju, V.; Burda, Y.; Edwards, H.; Baker, B.; Lee, T.; Leike, J.; Schulman, J.; Sutskever, I.; and Cobbe, K. 2023. Let’s Verify Step by Step. arXiv:2305.20050.
- Lin, Z.; Wang, X.; Cao, J.; Chai, J.; Yin, G.; Lin, W.; and He, R. 2025. ResT: Reshaping Token-Level Policy Gradients for Tool-Use Large Language Models. arXiv:2509.21826.
- Luo, H.; Shen, L.; He, H.; Wang, Y.; Liu, S.; Li, W.; Tan, N.; Cao, X.; and Tao, D. 2025a. O1-Pruner: Length-Harmonizing Fine-Tuning for O1-Like Reasoning Pruning. arXiv:2501.12570.
- Luo, M.; Tan, S.; Wong, J.; Shi, X.; Tang, W. Y.; Roongta, M.; Cai, C.; Luo, J.; Li, L. E.; Popa, R. A.; and Stoica, I. 2025b. DeepScaleR: Surpassing O1-Preview with a 1.5B Model by Scaling RL. <https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2>.
- Ma, X.; Wan, G.; Yu, R.; Fang, G.; and Wang, X. 2025. CoT-Valve: Length-Compressible Chain-of-Thought Tuning. arXiv:2502.09601.
- Mathematical Association of America. 2023. AMC Contests. <https://maa.org/student-programs/amc/>. Accessed: 2025-03-28.
- OpenAI. 2024. OpenAI o1: Reasoning-Enhanced Model. <https://openai.com/o1/>.
- OpenAI. 2025. OpenAI o3: Most Advanced Reasoning Model. <https://openai.com/index/introducing-o3-and-o4-mini/>.
- Qi, P.; Liu, Z.; Pang, T.; Du, C.; Lee, W. S.; and Lin, M. 2025. Optimizing Anytime Reasoning via Budget Relative Policy Optimization. arXiv:2505.13438.
- Qwen Team. 2025. QwQ-32B: Embracing the Power of Reinforcement Learning. <https://qwenlm.github.io/blog/qwq-32b/>.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y. K.; Wu, Y.; and Guo, D. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. arXiv:2402.03300.
- Shen, Y.; Zhang, J.; Huang, J.; Shi, S.; Zhang, W.; Yan, J.; Wang, N.; Wang, K.; Liu, Z.; and Lian, S. 2025. DAST: Difficulty-Adaptive Slow-Thinking for Large Reasoning Models. arXiv:2503.04472.

Sheng, G.; Zhang, C.; Ye, Z.; Wu, X.; Zhang, W.; Zhang, R.; Peng, Y.; Lin, H.; and Wu, C. 2024. HybridFlow: A Flexible and Efficient RLHF Framework. *arXiv preprint arXiv:2409.19256*.

Wang, S.; Yu, L.; Gao, C.; Zheng, C.; Liu, S.; Lu, R.; Dang, K.; Chen, X.; Yang, J.; Zhang, Z.; Liu, Y.; Yang, A.; Zhao, A.; Yue, Y.; Song, S.; Yu, B.; Huang, G.; and Lin, J. 2025. Beyond the 80/20 Rule: High-Entropy Minority Tokens Drive Effective Reinforcement Learning for LLM Reasoning. *arXiv:2506.01939*.

Wu, S.; Xie, J.; Zhang, Y.; Chen, A.; Zhang, K.; Su, Y.; and Xiao, Y. 2025. ARM: Adaptive Reasoning Model. *arXiv:2505.20258*.

Xu, S.; Xie, W.; Zhao, L.; and He, P. 2025a. Chain of Draft: Thinking Faster by Writing Less. *arXiv:2502.18600*.

Xu, Z.; Xia, S.; Yue, C.; Chai, J.; Tian, M.; Wang, X.; Lin, W.; Li, H.; and Yin, G. 2025b. MTIR-SQL: Multi-turn Tool-Integrated Reasoning Reinforcement Learning for Text-to-SQL. *arXiv:2510.25510*.

Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; Zheng, C.; Liu, D.; Zhou, F.; Huang, F.; Hu, F.; Ge, H.; Wei, H.; Lin, H.; Tang, J.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Zhou, J.; Lin, J.; Dang, K.; Bao, K.; Yang, K.; Yu, L.; Deng, L.; Li, M.; Xue, M.; Li, M.; Zhang, P.; Wang, P.; Zhu, Q.; Men, R.; Gao, R.; Liu, S.; Luo, S.; Li, T.; Tang, T.; Yin, W.; Ren, X.; Wang, X.; Zhang, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Wang, Z.; Cui, Z.; Zhang, Z.; Zhou, Z.; and Qiu, Z. 2025a. Qwen3 Technical Report. *arXiv:2505.09388*.

Yang, C.; Si, Q.; Duan, Y.; Zhu, Z.; Zhu, C.; Li, Q.; Lin, Z.; Cao, L.; and Wang, W. 2025b. Dynamic Early Exit in Reasoning Models. *arXiv:2504.15895*.

Ye, Y.; Huang, Z.; Xiao, Y.; Chern, E.; Xia, S.; and Liu, P. 2025. LIMO: Less is More for Reasoning. *arXiv:2502.03387*.

Yeo, E.; Tong, Y.; Niu, X.; Neubig, G.; and Yue, X. 2025. Demystifying Long Chain-of-Thought Reasoning in LLMs. In *ICLR 2025 Workshop on Navigating and Addressing Data Problems for Foundation Models*.

Yi et al. 2025. ShorterBetter: Guiding Reasoning Models to Find Optimal Inference Length for Efficient Reasoning. *arXiv:2504.21370*.

Yu, B.; Yuan, H.; Li, H.; Xu, X.; Wei, Y.; Wang, B.; Qi, W.; and Chen, K. 2025. Long-Short Chain-of-Thought Mixture Supervised Fine-Tuning Eliciting Efficient Reasoning in Large Language Models. *arXiv:2505.03469*.

Zhang, J.; Lin, N.; Hou, L.; Feng, L.; and Li, J. 2025a. AdaptThink: Reasoning Models Can Learn When to Think. *arXiv:2505.13417*.

Zhang, J.; Zhu, Y.; Sun, M.; Luo, Y.; Qiao, S.; Du, L.; Zheng, D.; Chen, H.; and Zhang, N. 2025b. LightThinker: Thinking Step-by-Step Compression. *arXiv:2502.15589*.

Zhang et al. 2025. GRPO-LEAD: A Difficulty-Aware Reinforcement Learning Approach for Concise Mathematical Reasoning in Language Models. *arXiv:2504.09696*.