

# DRQA: Dynamic Reasoning Quota Allocation for Controlling Overthinking in Reasoning Large Language Models

Anonymous ACL submission

## Abstract

Reasoning large language models (RLLMs), such as OpenAI-O3 and DeepSeek-R1, have recently demonstrated remarkable capabilities by performing structured and multi-step reasoning. However, recent studies reveal that RLLMs often suffer from overthinking, i.e., producing unnecessarily lengthy reasoning chains even for simple questions, leading to excessive token consumption and computational inefficiency. Interestingly, we observe that when processing multiple questions in batch mode, RLLMs exhibit more resource-efficient behavior by dynamically compressing reasoning steps for easier problems, due to implicit resource competition. Inspired by this, we propose *Dynamic Reasoning Quota Allocation (DRQA)*, a novel method that transfers the benefits of resource competition from batch processing to single-question inference. Specifically, DRQA leverages batch-generated preference data and reinforcement learning to train the model to allocate reasoning resources adaptively. By encouraging the model to internalize a preference for responses that are both accurate and concise, DRQA enables it to generate concise answers for simple questions while retaining sufficient reasoning depth for more challenging ones. Extensive experiments on a wide range of mathematical and scientific reasoning benchmarks demonstrate that DRQA significantly reduces token usage while maintaining, and in many cases improving, answer accuracy. By effectively mitigating the overthinking problem, DRQA offers a promising direction for more efficient and scalable deployment of RLLMs, and we hope it inspires further exploration into fine-grained control of reasoning behaviors.

## 1 Introduction

Reasoning large language models (RLLMs), such as OpenAI-O3 (OpenAI, 2025) and DeepSeek-R1 (DeepSeek-AI et al., 2025), have recently showcased remarkable capabilities in com-

plex problem solving and decision-making, achieving state-of-the-art performance across a wide range of tasks. However, recent studies have revealed that LLMs often generate unnecessarily lengthy reasoning chains, even for simple questions like “2+3=?” (Sui et al., 2025; Chen et al., 2025). While extended reasoning can improve accuracy on complex tasks, this tendency to *overthink* leads to excessive token usage and growing computational and economic costs, posing significant challenges for the scalable and practical deployment of RLLMs in real-world scenarios.

Interestingly, we uncover an emergent efficiency behavior during *batch inference*: when multiple questions are processed together, the total output length is significantly shorter than the combined length of the responses generated for each question individually. For example, as shown in Figure 1, answering three questions together yields only 648 tokens in total, compared to 1205 tokens when answered separately. This suggests that under a shared context window, questions implicitly compete for a global reasoning quota, prompting the model to prioritize essential logic and suppress redundancy, an effect we refer to as “*resource competition pressure*”.

This observation raises a core research question: can the benefits of resource competition in batch inference be transferred to single-question settings? If so, RLLMs could dynamically adjust their reasoning behaviors, offering concise responses for simple questions while allocating more resources to more complex ones. To this end, we introduce **Dynamic Reasoning Quota Allocation (DRQA)**, a novel approach that brings the advantages of resource competition into single-question inference, enabling more efficient and adaptive reasoning. Specifically, we first collect diverse reasoning chains under batch inference settings and analyze how the model automatically allocates the length of reasoning chains to problems of varying difficulty

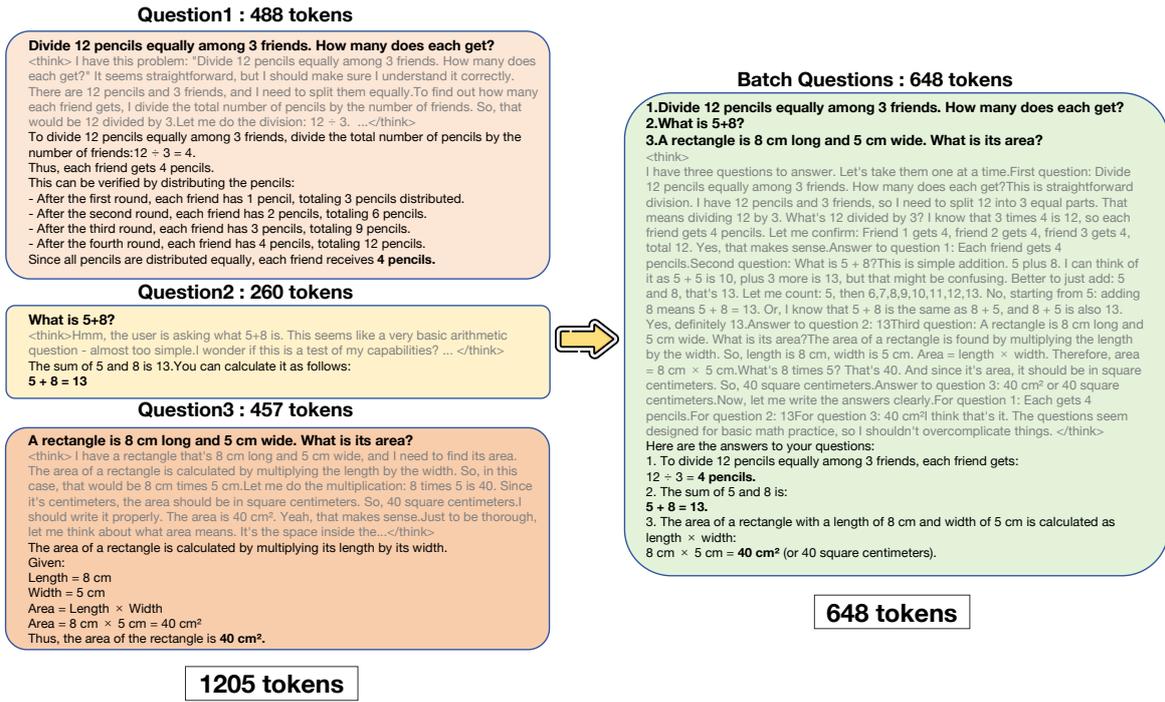


Figure 1: Comparison between batch inference and single-question inference using Deepseek-R1. Answering three questions together results in significantly fewer tokens than answering each question individually.

in the presence of resource competition. We then construct a preference dataset and introduce a reinforcement learning objective that enables the model to distinguish and learn the advantages of “concise and accurate” reasoning chains over those that are “verbose or incorrect”. By indirectly encouraging the model to favor the “concise and accurate” patterns that emerge from batch inference, we enhance its overall reasoning capabilities.

We evaluate the effectiveness of DRQA across a diverse set of reasoning benchmarks, including GSM8K (Cobbe et al., 2021), MATH-500 (Hendrycks et al., 2021), AIME 2024 and 2025 (MAA Committees), AMC (AI-MO, 2024), and GPQA-Diamond (Rein et al., 2023). Experimental results show that DRQA reduces token usage by over 30% while consistently maintaining or improving answer accuracy, offering an effective and scalable solution to the overthinking problem. In summary, our main contributions are:

- To the best of our knowledge, we for the first time systematically investigate how “resource competition pressure” can enhance the reasoning efficiency of RLLMs during batch inference.
- We propose DRQA, a novel method that transfers this efficiency mechanism to single-question inference by leveraging batch-generated preference

data and reinforcement learning. This enables the model to generate *concise answers for simple questions* while maintaining *deep reasoning for complex ones*.

- With extensive experiments, we demonstrate the effectiveness of DRQA compared to existing ones and analyze the results thoroughly.

## 2 Resource Competition During Batch Inference

**Batch Inference Encourages Efficient Reasoning.** As discussed in the introduction, a major challenge for RLLMs is their tendency to overthink, producing unnecessarily long reasoning chains even for simple questions. To investigate whether batch inference can encourage more efficient reasoning, we conduct a series of controlled experiments. Specifically, we randomly select 500 samples from the DeepScaleR dataset (Luo et al., 2025c) and evaluate several mainstream LLMs under two settings: (i) querying one question at a time (Vanilla), and (ii) querying two questions per prompt (Batch-2). As shown in Table 1, models including DeepSeek-R1 (DeepSeek-AI et al., 2025), Qwen3-32B (think) (Yang et al., 2025a), and Doubao-Seed-1.6 (Seed, 2025) consistently generate shorter outputs in the ‘Batch-2’ setting,

Model	Vanilla	Batch-2
Deepseek-R1	5640.4	4035.2
Qwen3-32B (think)	7761.6	5274.7
Doubao-Seed-1.6	5288.1	3898.2

Table 1: Comparison of average output token lengths across different models under the ‘Vanilla’ and ‘Batch-2’ settings.

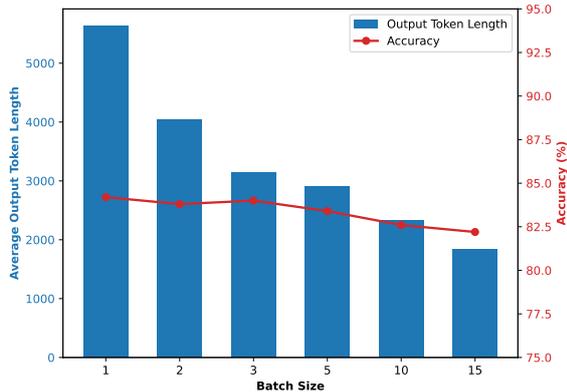


Figure 2: Impact of batch size on output length and accuracy using DeepSeek-R1.

suggesting that batch inference naturally promotes more concise reasoning and that this effect generalizes well across different model architectures.

**Scaling Up Batch Size Further Enhances Efficiency.** To further analyze the effect, we vary the batch size using DeepSeek-R1 as a case study, testing batches of 2, 3, 5, 10 and 15 questions. As shown in Figure 2, increasing the batch size leads to a continuous and substantial reduction in the average output length per question. Notably, this compression is achieved with only minimal degradation in answer accuracy, indicating that when more questions are packed into a single context window, the model tends to prioritize conciseness, allocating fewer tokens to simpler problems while preserving reasoning quality for more challenging ones. We refer to this emergent behavior as *resource competition pressure*.

These findings provide compelling empirical evidence that RLLMs are capable of implicit reasoning compression when facing context constraints. The behavior of allocating reasoning resources based on task complexity, without any explicit instruction, points to a promising direction for mitigating the overthinking problem commonly observed in single-question inference. Building on

this insight, our work is driven by a central research question: *can we transfer the benefits of resource competition from batch inference to single-question settings?* If so, models could learn to reason adaptively, producing concise answers for simple queries while maintaining sufficient reasoning depth for more complex ones. To this end, we introduce Dynamic Reasoning Quota Allocation (DRQA), detailed in the following section.

### 3 Methodology

Our goal is to enable RLLMs to assess question complexity and allocate reasoning resources adaptively, even when processing a single query. Ideally, the model should generate short responses for simple problems while preserving sufficient reasoning depth for more challenging ones, thereby improving inference efficiency without compromising answer accuracy. A key challenge in realizing this capability lies in how to effectively transfer “resource competition pressure” from batch inference to single-question settings. We first explore a straightforward solution via supervised fine-tuning (SFT) using batch-generated data. However, this approach revealed inherent limitations in teaching the model to internalize conciseness as a quality criterion. Inspired by recent advancements in Reinforcement Learning with Verifiable Rewards (RLVR) (Lambert et al., 2025; DeepSeek-AI et al., 2025), we introduce Dynamic Reasoning Quota Allocation (DRQA), a reinforcement learning framework that explicitly encourages reasoning that is both accurate and concise. By optimizing an intrinsic reward aligned with these dual objectives, DRQA guides models to dynamically allocate reasoning resources, enabling more efficient and adaptive inference.

#### 3.1 Supervised Fine-Tuning with Batch Data

Our initial approach to transferring the benefits of resource competition into single-question inference is based on imitation learning, where we apply supervised fine-tuning (SFT) to mimic the efficient reasoning patterns exhibited by models during batch inference.

**Method** We use DeepSeek-R1 (DeepSeek-AI et al., 2025) to perform batch inference over multiple questions sampled from DeepScaleR (Luo et al., 2025c) and collect the generated responses, which are consistently more concise than those from single-question inference. Based on these re-

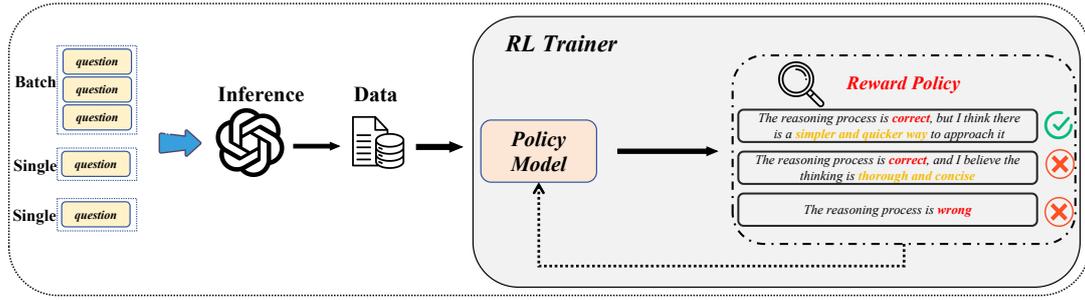


Figure 3: The pipeline of Dynamic Reasoning Quota Allocation (DRQA). Batched questions are input to LLM, producing reasoning chains labeled as A/B/C. Reinforcement learning trains the model to prefer concise and accurate reasoning for efficient resource allocation.

sults, we construct a dataset of “question–concise answer” pairs and apply full-parameter SFT on a Qwen3-8B (Yang et al., 2025a), with the goal of teaching it to generate similarly concise responses in single-question scenarios.

**Experimental Results and Analysis** We evaluate the fine-tuned models on a comprehensive set of reasoning benchmarks, including GSM8K (Cobbe et al., 2021), MATH-500 (Hendrycks et al., 2021), AIME 2024 (MAA Committees), GPQA-Diamond (Rein et al., 2023), AMC (AI-MO, 2024), and AIME 2025 (MAA Committees). The results reported in Table 2 indicate that SFT does lead to substantial reductions in output length. For example, on GSM8K, the average response length drops from 1878.55 to 575.64 tokens, a 69.36% reduction, demonstrating that overthinking is mitigated to some extent.

However, the efficiency gains come at a considerable cost to accuracy, particularly on more challenging tasks. As shown in Table 2, models fine-tuned with two-question batch data experience accuracy drops of 1.67% on GSM8K and 1.00% on MATH-500 compared to vanilla prompting. More notably, the performance degradation becomes increasingly severe with higher batch sizes and task complexity. On AIME 2024, accuracy falls from 74.67% (Vanilla) to 57.33% (Batch-2), 26.00% (Batch-3), and just 9.33% (Batch-5). These results suggest the emergence of catastrophic forgetting (Luo et al., 2025d): in attempting to mimic the surface-level conciseness of batch responses, the model compromises its ability to perform the deeper, more nuanced reasoning necessary for solving complex problems.

In summary, while supervised fine-tuning with batch data effectively mitigates overthinking and

improves inference efficiency, it comes at the cost of reasoning accuracy, especially on complex tasks, highlighting its limitations for real-world deployment. These shortcomings underscore the need for a more principled solution that can balance conciseness with reasoning depth, which motivates our proposed method: Dynamic Reasoning Quota Allocation (DRQA).

### 3.2 Dynamic Reasoning Quota Allocation

Rather than imitating outputs from batch inference, we aim to endow the model with an intrinsic ability to evaluate and generate reasoning chains that are both accurate and concise. To this end, we propose Dynamic Reasoning Quota Allocation (DRQA), a reinforcement learning framework that enables RLLMs to dynamically allocate reasoning resources in single-question inference.

**Core Idea** The core idea of DRQA is to enhance the model’s intrinsic reasoning capabilities by equipping it with the ability to evaluate the quality of its own reasoning chains. Specifically, the model is trained to make two key judgments: (i) whether a given reasoning chain is logically correct, and (ii) if correct, whether it is unnecessarily verbose. *By developing this self-evaluation ability, the model learns to strike a balance between accuracy and conciseness during generation, effectively realizing adaptive resource allocation.*

**Preference Data Construction** To train this evaluation ability, we construct a preference dataset consisting of multiple-choice question-answering samples. Each sample contains a question, a model-generated chain of thought (CoT), and three evaluation options that reflect different levels of reasoning quality:

Method	GSM8K		Math500		AIME2024		GPQA-Diamond		AMC		AIME2025		Overall	
	Acc	tokens	Acc	tokens	Acc	tokens	Acc	tokens	Acc	tokens	Acc	tokens	Acc	tokens
Vanilla	95.67	1878.55	96.00	5270.58	74.67	15468.23	66.67	8685.21	97.50	8608.85	63.33	18058.65	82.31	9661.68
Batch-2	96.67	575.64	95.00	2359.21	57.33	11100.55	53.54	6874.42	90.00	4136.58	45.33	13130.95	72.98	6362.89
Batch-3	93.33	437.23	82.67	1593.53	26.00	5685.36	55.56	3555.65	77.50	4098.10	28.00	7400.53	60.51	3795.07
Batch-5	93.33	336.81	69.67	434.50	9.33	2486.77	46.46	1190.23	42.50	922.25	7.33	2365.41	44.77	1289.33

Table 2: Single-question evaluation results of Qwen3-8B after SFT with data generated by batch inference. Batch-X denotes fine-tuning with data from batches of X questions, and Vanilla refers to the original model without SFT.

- **A**: The reasoning process is correct, but I think there is a simpler and quicker way to approach it.
- **B**: The reasoning process is correct, and I believe the thinking is thorough and concise.
- **C**: The reasoning process is wrong.

The dataset construction process involves three key steps. First, for ease of evaluation, we select all questions in the DeepScaleR (Luo et al., 2025c) dataset whose answers are numbers of various types, resulting in approximately 30,000 samples. Second, for each question, we generate two types of reasoning chains using DeepSeek-R1 (DeepSeek-AI et al., 2025): (1) *vanilla CoTs* obtained by prompting the model with individual questions, and (2) *batch CoTs* generated by prompting the model with batched questions, followed by extracting the corresponding reasoning chain for each question. Finally, we assign labels based on reasoning correctness and conciseness: for vanilla CoTs, we label **A** if the reasoning is correct, and **C** if incorrect; for batch CoTs, we label **B** if the reasoning is correct, and **C** if incorrect. This labeling scheme enables the model to learn nuanced distinctions between correct-but-verbose reasoning (option A), correct-and-concise reasoning (option B), and incorrect reasoning (option C), thereby developing a clearer understanding of what constitutes a high-quality reasoning chain.

**Reinforcement Learning Framework** We use Group Relative Policy Optimization (GRPO) (Shao et al., 2024) to train the model to accurately classify each reasoning chain as A, B, or C, thus encouraging concise and accurate reasoning. Formally, the GRPO objective is defined as maximizing the likelihood of selecting the correct evaluation label:

$$\mathcal{L}_{\text{GRPO}}(\theta) = \mathbb{E}_{\tau \sim \mathcal{D}} \left[ \sum_{a \in \mathcal{G}} \log \pi_{\theta}(a | s) \hat{A}(a, s, a^*) - \beta \text{KL}(\pi_{\theta} \| \pi_{\text{old}}) \right] \quad (1)$$

where  $\tau \sim \mathcal{D}$  denotes a sample from the dataset, with state  $s$  representing the question, reasoning

chain, and multiple-choice options (A, B, C);  $a^*$  is the ground-truth label;  $\mathcal{G} = \{A, B, C\}$  is the set of actions;  $\hat{A}(a, s, a^*)$  is the relative advantage estimate, positive if  $a = a^*$  and negative otherwise;  $\text{KL}(\pi_{\theta} \| \pi_{\text{old}})$  is the KL divergence between the current and old policies, constrains the policy update; and  $\beta$  is a regularization coefficient balancing learning efficiency and policy stability. This training objective encourages the model to assign higher probabilities to correct judgments while mitigating the risk of catastrophic forgetting caused by over-updating, a common issue encountered in SFT. As a result, the model gradually internalizes a preference for reasoning chains that are both correct and concise.

**Summary** DRQA enables the model to move beyond surface-level imitation and develop an intrinsic, reward-driven preference for high-quality reasoning. By balancing accuracy and conciseness, the model learns to allocate reasoning resources more effectively, addressing the limitations of SFT and supporting more efficient and adaptive inference in single-question settings.

## 4 Experiments

In this section, we systematically evaluate the performance of the proposed DRQA algorithm, focusing on its ability to balance reasoning accuracy and efficiency. We compare DRQA against a range of strong baselines and provide an in-depth analysis of the results.

### 4.1 Experimental Setup

**Models** We evaluate all methods using two widely adopted distilled models: DeepSeek-R1-Distill-Qwen-1.5B and DeepSeek-R1-Distill-Qwen-7B. Both models are derived from the more powerful DeepSeek-R1 (DeepSeek-AI et al., 2025) through large-scale distillation, offering a favorable trade-off between computational efficiency and reasoning capability.

**Datasets** For training, we use the dataset described in Section 3.2, constructed by performing batch inference with DeepSeek-R1 on the DeepScaleR (Luo et al., 2025c) training set. This process yields over 50,000 multiple-choice examples annotated with reasoning quality labels.

**Baselines** To assess the effectiveness of DRQA, we compare it against a comprehensive set of strong baselines approaches (refer to Appendix A for detailed descriptions of the baselines). All baselines are either publicly released or carefully reproduced according to their original protocols.

**Evaluation** We evaluate the performance of different methods across a diverse set of benchmarks. For mathematical reasoning, we include GSM8K (Cobbe et al., 2021), MATH-500 (Hendrycks et al., 2021), AIME 2024 and 2025 (MAA Committees), and AMC 2023 (AIMO, 2024). For domain-specific scientific reasoning, we use the high-quality GPQA-diamond subset (Rein et al., 2023). Detailed descriptions of these datasets are provided in Appendix B. We use both accuracy and response length as evaluation metrics and report the average performance across all test sets. For the AIME datasets, which contain only 30 questions each, we repeatedly sample 5 responses for each case and report the average results to ensure more stable and reliable evaluation.

All models are evaluated using a unified inference configuration to ensure fair comparison. Experiments are conducted with the vLLM framework on a computing cluster equipped with eight A800 (40GB) GPUs. The inference parameters are set to a temperature of 0.6 and a maximum generation length of 32K tokens.

**Training Details** We use verl (Sheng et al., 2024) as the training framework. We set the batch size to 256, the number of rollouts to 16, the learning rate to  $1 \times 10^{-6}$ , and the maximum response length to 16K tokens. The model is trained for one epoch, consisting of 204 steps in total.

## 4.2 Main Results

As shown in Table 3, DRQA demonstrates clear superiority in both answer accuracy and response efficiency across all mathematical benchmarks. For example, on GSM8K with the 1.5B model, DRQA achieves an accuracy of 86.67%, outperforming the vanilla baseline by 2 percentage points, while reducing average token usage from 1928.96 to

1427.63, a 25.9% reduction. Similar patterns are observed on more challenging datasets such as AIME 2024 and MATH-500, where DRQA maintains high accuracy while significantly reducing output length. These results highlight DRQA’s effectiveness in dynamically allocating reasoning resources, enabling it to strike a favorable balance between accuracy and efficiency across tasks of varying difficulties. Moreover, DRQA demonstrates strong generalization on out-of-distribution (OOD) benchmarks, as evidenced by its performance on GPQA-Diamond.

We also compare DRQA with aggressive compression methods such as ShorterBetter (Yi et al., 2025) and DAST (Shen et al., 2025), which can reduce output length even further, for example, generating outputs as short as 107.86 tokens on GSM8K. However, these methods often suffer from severe accuracy degradation, with performance drops exceeding 20 percentage points in some cases. This highlights a key limitation of methods that rely solely on length-based reward signals: they tend to compromise the logical integrity of reasoning chains, limiting their practical applicability.

Notably, DRQA remains highly effective on larger models. On GSM8K with the 7B model, DRQA improves accuracy by 1.34% over the baseline while reducing token usage by 23.6%. This demonstrates that DRQA continues to deliver substantial efficiency gains and accuracy improvements at higher model scales. Across all benchmarks, it consistently achieves the most favorable trade-off between accuracy and output efficiency. Compared to strong baselines such as DAST (Shen et al., 2025), O1-Pruner (Luo et al., 2025b), Dynasor-CoT (Fu et al., 2025), and DEER (Xia et al., 2024), DRQA not only matches or surpasses them in length reduction but, more importantly, maintains state-of-the-art reasoning accuracy.

Overall, DRQA achieves an average accuracy improvement of 1.5 percentage points and an average token usage reduction of 31.3% across all evaluated benchmarks and both model sizes. These results provide compelling evidence that DRQA effectively transfers the benefits of “resource competition pressure” from batch inference to single-question settings, establishing a strong foundation for the efficient and scalable deployment of RLLMs.

Method	GSM8K		MATH-500		AIME 2024		GPQA-Diamond		AMC 2023		AIME 2025		Overall	
	Acc	Tokens	Acc	Tokens	Acc	Tokens	Acc	Tokens	Acc	Tokens	Acc	Tokens	Acc <sub>All</sub>	Token <sub>All</sub>
<b>DeepSeek-R1-Distill-Qwen-1.5B</b>														
Vanilla	84.67%	1928.96	83.33%	5536.14	28.67%	14394.61	30.84%	14731.59	72.50%	8830.10	23.67%	15323.3	53.95%	10124.12
O1-Pruner	74.80%	458	82.20%	3212	28.90%	10361	-	-	-	-	-	-	-	-
DAST	77.20%	586	83.00%	2428	26.90%	7745	-	-	-	-	-	-	-	-
ShortBetter	63.67%	<b>107.86</b>	60.33%	<b>1186.27</b>	11.33%	<b>2935.68</b>	21.72%	<b>1433.95</b>	57.50%	<b>1260.43</b>	12.67%	<b>3326.22</b>	37.87% <sup>+16.08</sup>	<b>1708.40</b> <sup>+83.13%</sup>
AdaptThink	86.00%	324.26	83.67%	1244.98	29.33%	7044.06	29.80%	4744.23	72.50%	2441.45	24.67%	7490.79	54.33% <sup>+0.38</sup>	3881.63 <sup>+61.66%</sup>
GRPO	<b>87.33%</b>	1691.19	<b>84.67%</b>	5743.01	<b>32.67%</b>	15017.54	27.78%	13809.53	<b>77.50%</b>	9378.21	24.00%	13082.98	55.66% <sup>+1.71</sup>	9787.08 <sup>-3.3%</sup>
GRPO+Length Penalty	86.00%	722.34	84.67%	2479.14	24.67%	9011.46	26.76%	6148.50	67.50%	3130.51	22.00%	9782.34	51.93% <sup>+2.01</sup>	5212.38 <sup>+48.52%</sup>
SFT	81.67%	2296.54	80.33%	5465.95	25.33%	21337.44	27.27%	18540.94	65.00%	8806.48	19.33%	20258.82	49.82% <sup>+1.13</sup>	12784.36 <sup>+26.28%</sup>
DRQA(our)	86.67%	1427.63	<b>84.67%</b>	3488.08	32.00%	11008.31	<b>31.81%</b>	9148.83	75.00%	5355.03	24.00%	10382.12	<b>55.69%</b> <sup>+1.74</sup>	6801.67 <sup>-32.82%</sup>
<b>DeepSeek-R1-Distill-Qwen-7B</b>														
Vanilla	91.33%	1735.5	90.40%	5099.95	53.33%	13712.6	48.98%	13313.92	90.00%	6349.53	40.00%	14248.11	69.01%	9076.60
DAST	86.70%	459	89.60%	2162	45.60%	7578	-	-	-	-	-	-	-	-
O1-Pruner	87.60%	428	86.60%	2534	49.20%	9719	-	-	-	-	-	-	-	-
Dynasor-CoT	89.60%	1285	89.00%	2971	46.70%	12695	30.50%	7639	85.00%	5980	-	-	-	-
DEER	90.60%	6917	89.80%	2143	49.20%	9839	31.30%	5469	85.00%	4451	-	-	-	-
ShortBetter	70.00%	<b>112.86</b>	68.00%	<b>623.44</b>	41.33%	<b>5005.96</b>	43.43%	<b>1811.43</b>	57.50%	<b>1567.50</b>	30.67%	5393.96	51.82% <sup>+17.19</sup>	<b>2419.19</b> <sup>+73.38%</sup>
AdaptThink	89.67%	296.94	91.67%	1839.59	54.00%	9894.05	51.52%	7128.95	87.50%	3287.95	39.33%	12454.59	68.95% <sup>+0.06</sup>	5817.01 <sup>+35.91%</sup>
AutoL2S	93.33%	444.8	83.33%	3113.93	40.67%	6499.32	45.39%	2553.01	85.00%	2613.05	31.33%	<b>3669.53</b>	63.18% <sup>+5.84</sup>	3148.94 <sup>+65.31%</sup>
GRPO	<b>93.67%</b>	1524.24	<b>92.00%</b>	4532.21	<b>54.67%</b>	12013.92	47.47%	12124.10	87.50%	5130.13	<b>41.33%</b>	12192.12	69.44% <sup>+0.43</sup>	7919.45 <sup>+12.75%</sup>
GRPO+Length Penalty	91.33%	876.25	91.33%	2751.13	52.00%	7213.11	45.96%	7124	<b>92.50%</b>	3256.02	39.67%	6058.40	68.80% <sup>+0.21</sup>	4546.49 <sup>+49.91%</sup>
SFT	92.33%	1317.85	92.00%	3824.43	44.67%	14903.82	46.97%	12385.43	77.50%	5519.55	32.00%	13931.80	64.25% <sup>+4.76</sup>	8647.15 <sup>+4.73%</sup>
DRQA(our)	92.67%	1324.24	91.40%	3902.74	<b>54.67%</b>	10007.18	<b>49.50%</b>	8988.50	<b>92.50%</b>	4463.03	40.67%	9545.44	<b>70.24%</b> <sup>+1.23</sup>	6371.85 <sup>+29.80%</sup>

Table 3: Performance of different methods using two RLLMs. DRQA achieves competitive or superior accuracy while greatly reducing token usage across all datasets and model sizes, striking an excellent balance between performance and efficiency.

### 4.3 Ablation Study

To thoroughly assess the contribution of each core component in DRQA, we conduct a series of ablation studies that isolate the effects of different training paradigms and input conciseness on reasoning performance and efficiency. All experiments are performed using the same benchmark datasets, evaluation metrics, and base model (DeepSeek-R1-Distill-Qwen-7B) as in the main study, with consistent inference configurations to ensure fair comparison.

**Effect of Batch Size in DRQA Data Construction.** We investigate the impact of different batch sizes on model performance. Specifically, we construct preference datasets by prompting DeepSeek-R1 with batches of 2, 3, or 5 questions, then splitting the outputs into individual reasoning chains for downstream RL training. This design allows us to analyze how increasing levels of resource competition influence both answer accuracy and response efficiency within the DRQA framework.

**Replacing Batch Reasoning Data with Qwen2.5-7B Concise Chains** To evaluate the importance of batch-induced resource competition, we consider an alternative setting where the preference dataset is constructed using concise reasoning chains generated directly by Qwen2.5-7B (Qwen et al., 2025), without leveraging batch inference. This comparison allows us to disentangle the effects of resource-driven compression from those achieved solely through the model’s inherent ability

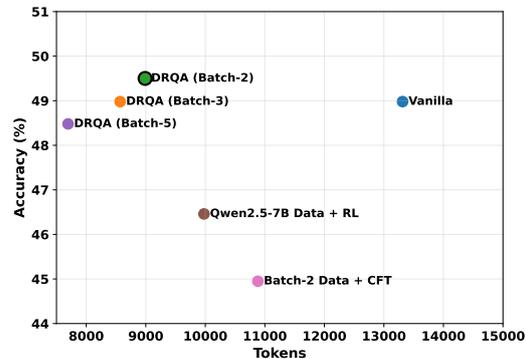


Figure 4: The efficiency-accuracy trade-off on GPQA-diamond for DRQA and ablation variants.

to generate concise outputs.

**Critique Fine-Tuning with Preference data** Beyond reinforcement learning, we also evaluate the Critique Fine-Tuning (CFT) paradigm (Wang et al., 2025) as an alternative training strategy, applying it to the preference data we constructed.

#### 4.3.1 Results and Analysis

Table 4 presents the results of our ablation study. As batch size increases, the model produces increasingly concise outputs, with token usage reduced by up to 45% for larger batches. However, this efficiency gain comes at the cost of declining accuracy, highlighting a trade-off between efficiency and correctness. Notably, a batch size of 2 achieves the best balance, improving accuracy while significantly reducing token consumption compared to the vanilla baseline.

Method	GSM8K		MATH-500		AIME 2024		Overall	
	Acc	tokens	Acc	tokens	Acc	tokens	Acc	tokens
<i>Vanilla</i>	91.33%	1735.5	90.40%	5099.95	53.33%	13712.6	78.35%	6849.35
<i>DRQA (Batch-2)</i>	92.67%	1324.24	91.33%	3902.74	54.67%	10007.18	<b>79.58%</b> <sup>+1.23</sup>	5078.05 <sup>-25.86%</sup>
<i>DRQA (Batch-3)</i>	91.67%	1212.59	90.20%	3311.20	53.33%	8805.24	78.4% <sup>+0.05</sup>	4443.01 <sup>-35.13%</sup>
<i>DRQA (Batch-5)</i>	90.67%	1158.88	89.80%	2675.81	49.33%	7366.80	76.60% <sup>-1.75</sup>	<b>3733.83</b> <sup>-45.49%</sup>
<i>Qwen2.5-7B Data + RL</i>	90.00%	1434.65	89.60%	3313.12	50.67%	12190.59	76.76% <sup>-1.60</sup>	5646.12 <sup>-17.57%</sup>
<i>Batch-2 Data + CFT</i>	89.67%	1361.00	88.20%	3973.54	49.66%	10012.55	75.84% <sup>-2.51</sup>	5115.70 <sup>-25.31%</sup>

Table 4: Ablation experiments across different training paradigms.

When compared to concise reasoning chains generated directly by Qwen2.5-7B (Qwen et al., 2025) without batch inference, we observe that only batch-induced compression achieves both high efficiency and strong accuracy. Similarly, while Critique Fine-Tuning helps reduce output length, it leads to a notable accuracy drop, underscoring the importance of reinforcement learning for preserving reasoning quality. Figure 4 further supports these insights, showing that DRQA achieves the best overall trade-off on the OOD dataset GPQA-Diamond, highlighting its robustness across both in-distribution and out-of-distribution scenarios.

## 5 Related Work

### 5.1 Reasoning Large Language Models

Recent reasoning large language models (RLLMs), such as OpenAI-O3 (OpenAI, 2025), Deepseek-R1 (DeepSeek-AI et al., 2025), and QwQ (Team, 2025) leverage chain-of-thought (Wei et al., 2023) for step-by-step reasoning, achieving state-of-the-art performance across a wide range of tasks, including mathematical reasoning and language understanding. Our work focuses on further improving the efficiency of such reasoning models.

### 5.2 Efficient Reasoning

Reasoning efficiency in RLLMs (Qu et al., 2025; Sui et al., 2025) refers to balancing task quality and computational cost. Models like OpenAI-O3 (OpenAI, 2025) and DeepSeek-R1 (DeepSeek-AI et al., 2025) often generate too long and redundant reasoning chains, over explaining simple problems while sometimes offering shallow reasoning for complex ones. Main approaches for improving efficiency include:

- **Inference time control:** Methods such as TALE (Han et al., 2025), DEER (Yang et al., 2025b) apply token budgets or early exit strategies inspired by dual-system theory.
- **Chain compression and supervised tuning:** To-

kenSkip (Xia et al., 2025), CoT-Valve (Ma et al., 2025), and AutoL2S (Luo et al., 2025a) use supervised fine-tuning or distillation to shorten reasoning chains, often improving conciseness but sometimes at the expense of complex reasoning.

- **Reinforcement learning approaches:** DAST (Shen et al., 2025), O1-Pruner (Luo et al., 2025b), and S-GRPO (Dai et al., 2025) introduce reward functions to penalize lengthy outputs and promote token efficiency, supporting adaptive reasoning with little loss of accuracy.

These methods largely depend on fixed budgets or hand crafted rewards. Our DRQA instead transfers the “resource competition pressure” observed in batch inference to single-question settings, enabling models to automatically adjust reasoning length according to problem complexity, providing brief responses for simple questions and detailed explanations for challenging ones without manual constraints.

## 6 Conclusion

This paper introduces Dynamic Reasoning Quota Allocation (DRQA), a novel approach aimed at addressing the overthinking problem in reasoning large language models (RLLMs). Motivated by the observation that resource competition pressure in batch inference naturally encourages efficient reasoning, DRQA leverages batch-generated data and reinforcement learning to transfer the benefits of resource competition from batch inference to single-question scenarios. Specifically, the model is trained to recognize and prefer reasoning chains that are both accurate and concise, resulting in responses that are both concise and accurate for single-question inference. Extensive experimental results and analysis show that DRQA significantly reduces token consumption while maintaining, or even improving accuracy, effectively alleviating overthinking. DRQA offers a new direction for more efficient and scalable deployment of RLLMs.

## 589 Limitations

590 Although DRQA has demonstrated strong perfor-  
591 mance on both mathematical reasoning and OOD  
592 scientific benchmarks, its effectiveness on broader  
593 tasks and real-world applications remains to be fur-  
594 ther evaluated. In future work, we plan to extend  
595 DRQA to a wider range of domains and explore  
596 its potential for finer-grained control over reason-  
597 ing processes, as well as improvements on other  
598 efficiency dimensions beyond reasoning length.

## 599 References

600 AI-MO. 2024. [Amc 2023](#).

601 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He,  
602 Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi  
603 Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang,  
604 Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025. [Do  
605 not think that much for 2+3=? on the overthinking of  
606 o1-like llms](#). *Preprint*, arXiv:2412.21187.

607 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,  
608 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias  
609 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro  
610 Nakano, Christopher Hesse, and John Schulman.  
611 2021. [Training verifiers to solve math word prob-  
612 lems](#). *Preprint*, arXiv:2110.14168.

613 Muzhi Dai, Chenxu Yang, and Qingyi Si. 2025. [S-grpo:  
614 Early exit via reinforcement learning in reasoning  
615 models](#). *Preprint*, arXiv:2505.07686.

616 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang,  
617 Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
618 Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang,  
619 Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhi-  
620 hong Shao, Zhuoshu Li, Ziyi Gao, and 181 others.  
621 2025. [Deepseek-r1: Incentivizing reasoning capa-  
622 bility in llms via reinforcement learning](#). *Preprint*,  
623 arXiv:2501.12948.

624 Yichao Fu, Junda Chen, Yonghao Zhuang, Zheyu Fu,  
625 Ion Stoica, and Hao Zhang. 2025. [Reasoning without  
626 self-doubt: More efficient chain-of-thought through  
627 certainty probing](#). In *ICLR 2025 Workshop on Founda-  
628 tion Models in the Wild*.

629 Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu  
630 Zhao, Shiqing Ma, and Zhenyu Chen. 2025. [Token-budget-aware llm reasoning](#). *Preprint*,  
631 arXiv:2412.18547.  
632

633 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul  
634 Arora, Steven Basart, Eric Tang, Dawn Song, and  
635 Jacob Steinhardt. 2021. [Measuring mathematical  
636 problem solving with the math dataset](#). *Preprint*,  
637 arXiv:2103.03874.

638 Nathan Lambert, Jacob Morrison, Valentina Pyatkin,  
639 Shengyi Huang, Hamish Ivison, Faeze Brahman,  
640 Lester James V. Miranda, Alisa Liu, Nouha Dziri,

Shane Lyu, Yuling Gu, Saumya Malik, Victoria  
Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le  
Bras, Oyvind Tafjord, Chris Wilhelm, Luca Sol-  
dani, and 4 others. 2025. [Tulu 3: Pushing fron-  
tiers in open language model post-training](#). *Preprint*,  
arXiv:2411.15124. 641  
642  
643  
644  
645  
646

Feng Luo, Yu-Neng Chuang, Guanchu Wang, Hoang  
Anh Duy Le, Shaochen Zhong, Hongyi Liu, Jiayi  
Yuan, Yang Sui, Vladimir Braverman, Vipin Chaud-  
hary, and Xia Hu. 2025a. [Autol2s: Auto long-  
short reasoning for efficient large language models](#).  
*Preprint*, arXiv:2505.22662. 647  
648  
649  
650  
651  
652

Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shi-  
wei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao,  
and Dacheng Tao. 2025b. [O1-pruner: Length-  
harmonizing fine-tuning for o1-like reasoning prun-  
ing](#). *Preprint*, arXiv:2501.12570. 653  
654  
655  
656  
657

Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi,  
William Tang, Manan Roongta, Colin Cai, Jeffrey  
Luo, Tianjun Zhang, Erran Li, Raluca Ada Popa,  
and Ion Stoica. 2025c. [Deepscaler: Surpassing o1-  
preview with a 1.5b model by scaling rl](#). Notion  
Blog. 658  
659  
660  
661  
662  
663

Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou,  
and Yue Zhang. 2025d. [An empirical study of cata-  
strophic forgetting in large language models during  
continual fine-tuning](#). *Preprint*, arXiv:2308.08747. 664  
665  
666  
667

Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan  
Fang, and Xinchao Wang. 2025. [Cot-valve: Length-  
compressible chain-of-thought tuning](#). *Preprint*,  
arXiv:2502.09601. 668  
669  
670  
671

MAA Committees. Aime problems and solutions.  
[https://artofproblemsolving.com/wiki/  
index.php/AIME\\_Problems\\_and\\_Solutions](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions). 672  
673  
674

OpenAI. 2025. [Introducing O3 and O4  
Mini](#). [https://openai.com/index/  
introducing-o3-and-o4-mini/](https://openai.com/index/introducing-o3-and-o4-mini/). Accessed:  
2025-04-16. 675  
676  
677  
678

Xiaoye Qu, Yafu Li, Zhaochen Su, Weigao Sun, Jianhao  
Yan, Dongrui Liu, Ganqu Cui, Daizong Liu, Shuxian  
Liang, Junxian He, Peng Li, Wei Wei, Jing Shao,  
Chaochao Lu, Yue Zhang, Xian-Sheng Hua, Bowen  
Zhou, and Yu Cheng. 2025. [A survey of efficient rea-  
soning for large reasoning models: Language, multi-  
modality, and beyond](#). *Preprint*, arXiv:2503.21614. 679  
680  
681  
682  
683  
684  
685

Qwen, :, An Yang, Baosong Yang, Beichen Zhang,  
Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan  
Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan  
Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin  
Yang, Jiayi Yang, Jingren Zhou, and 25 oth-  
ers. 2025. [Qwen2.5 technical report](#). *Preprint*,  
arXiv:2412.15115. 686  
687  
688  
689  
690  
691  
692

David Rein, Betty Li Hou, Asa Cooper Stickland,  
Jackson Petty, Richard Yuanzhe Pang, Julien Di-  
rani, Julian Michael, and Samuel R. Bowman. 2023.  
[Gpqa: A graduate-level google-proof q&a bench-  
mark](#). *Preprint*, arXiv:2311.12022. 693  
694  
695  
696  
697

698	ByteDance Seed. 2025. <a href="#">Doubao-seed-1.6</a> .	Jingyang Yi, Jiazheng Wang, and Sida Li. 2025. <a href="#">Shorterbetter: Guiding reasoning models to find optimal inference length for efficient reasoning</a> . <i>Preprint</i> , arXiv:2504.21370.	751 752 753 754
699 700 701 702 703 704	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. <a href="#">Deepseekmath: Pushing the limits of mathematical reasoning in open language models</a> . <i>Preprint</i> , arXiv:2402.03300.	Jiajie Zhang, Nianyi Lin, Lei Hou, Ling Feng, and Juanzi Li. 2025. <a href="#">Adaptthink: Reasoning models can learn when to think</a> . <i>Preprint</i> , arXiv:2505.13417.	755 756 757
705 706 707 708 709	Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai Wang, Zhaoxiang Liu, and Shiguo Lian. 2025. <a href="#">Dast: Difficulty-adaptive slow-thinking for large reasoning models</a> . <i>Preprint</i> , arXiv:2503.04472.		
710 711 712 713 714	Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. <a href="#">Hybridflow: A flexible and efficient rlhf framework</a> . <i>arXiv preprint arXiv:2409.19256</i> .		
715 716 717 718 719 720	Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Hanjie Chen, and Xia Hu. 2025. <a href="#">Stop overthinking: A survey on efficient reasoning for large language models</a> . <i>Preprint</i> , arXiv:2503.16419.		
721 722	Qwen Team. 2025. <a href="#">Qwq-32b: Embracing the power of reinforcement learning</a> .		
723 724 725 726	Yubo Wang, Xiang Yue, and Wenhua Chen. 2025. <a href="#">Critique fine-tuning: Learning to critique is more effective than learning to imitate</a> . <i>Preprint</i> , arXiv:2501.17703.		
727 728 729 730 731	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. <a href="#">Chain-of-thought prompting elicits reasoning in large language models</a> . <i>Preprint</i> , arXiv:2201.11903.		
732 733 734 735	Bo Xia, Yilun Kong, Yongzhe Chang, Bo Yuan, Zhiheng Li, Xueqian Wang, and Bin Liang. 2024. <a href="#">Deer: A delay-resilient framework for reinforcement learning with variable delays</a> . <i>Preprint</i> , arXiv:2406.03102.		
736 737 738 739	Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. 2025. <a href="#">Tokenskip: Controllable chain-of-thought compression in llms</a> . <i>Preprint</i> , arXiv:2502.12067.		
740 741 742 743 744 745 746	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. <a href="#">Qwen3 technical report</a> . <i>Preprint</i> , arXiv:2505.09388.		
747 748 749 750	Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Qiaowei Li, Zheng Lin, Li Cao, and Weiping Wang. 2025b. <a href="#">Dynamic early exit in reasoning models</a> . <i>Preprint</i> , arXiv:2504.15895.		

## A Baseline Methods

We consider the following baseline methods in our experiments:

- **GRPO**: We train a model on the DeepScaleR (Luo et al., 2025c) dataset using the Group Relative Policy Optimization algorithm, where only answer correctness is used as the reward signal.
- **GRPO+Length Penalty**: This variant further introduces a length penalty to the reward design: for correct answers, shorter responses yield higher rewards, while for incorrect answers, longer responses incur greater penalties. This encourages the model to produce concise and accurate reasoning.
- **SFT (Supervised Fine-Tuning)**: We perform full-parameter supervised fine-tuning on the model using question-answer pairs generated via batch inference of Deepseek-R1 on the DeepScaleR (Luo et al., 2025c) dataset.
- **AdaptThink (Zhang et al., 2025)**: This approach encourages adaptive selection between direct answer and step-by-step reasoning (Chain-of-Thought) based on question difficulty. Training objectives and sample balancing enable the model to flexibly explore both thinking modes, improving reasoning efficiency and performance.
- **AutoL2S (Luo et al., 2025a)**: A dynamic, model-agnostic framework that annotates each question with both long and short Chain-of-Thought (CoT) solutions. By marking simple questions with <EASY>, the model is trained to automatically select concise CoT for simple problems and detailed reasoning for complex ones.
- **DAST (Shen et al., 2025)**: DAST explicitly quantifies problem difficulty via a token length budget and employs a reward that penalizes redundant reasoning on simple problems while encouraging extensive CoT for difficult ones. This preference data is optimized via SimPO, enabling efficient dynamic control over reasoning path length.
- **O1-Pruner (Luo et al., 2025b)**: Based on reinforcement learning, this method rewards shorter CoT traces without compromising accuracy. It employs an offline PPO-like procedure to prune redundant reasoning while preserving or even improving correctness.
- **ShorterBetter (Yi et al., 2025)**: This RL-based approach defines the optimal length for each question as the shortest possible correct response and leverages this dynamic signal as a reward for GRPO-based training, guiding the model toward concise yet accurate answers.
- **Dynasor-CoT (Fu et al., 2025)**: Without extra training, this method dynamically truncates reasoning by probing intermediate answers, monitoring consistency, and detecting hesitancy tokens. This yields substantial token savings while preserving accuracy.
- **DEER (Xia et al., 2024)**: DEER employs a dynamic early-exit mechanism by monitoring reasoning transitions (such as “Wait”) to induce trial answers. Decisions to terminate CoT generation are based on confidence estimation, reducing reasoning length without additional training.

All baseline models are tested under identical inference configurations and on the same benchmark datasets to guarantee fair and reliable comparison. For each baseline, we use either the officially released model or reproduce the method using released data and code.

## B Dataset Details

### Mathematical Reasoning Datasets

- **GSM8K**: This dataset contains 8,500 English elementary school single-step math reasoning questions. It serves as one of the mainstream benchmarks for evaluating the math reasoning abilities of large language models, focusing on basic arithmetic reasoning skills.

- **MATH-500:** Includes 500 medium-difficulty mathematical problems covering algebra, geometry, number theory, and other areas, designed to test the model’s comprehensive mathematical reasoning ability.
- **AIME 2024/2025:** Originating from the American Invitational Mathematics Examination 2024 and 2025, each set contains 30 high-difficulty math questions, mainly assessing complex mathematical reasoning and problem-solving skills.
- **AMC 2023:** 40 questions from the American Mathematics Competitions (AMC), covering middle to high school levels, examining fundamental and advanced mathematics knowledge and problem solving abilities.

**Scientific Reasoning Dataset** To evaluate model reasoning performance in other domains, we use the high-quality GPQA-diamond subset from the GPQA dataset. GPQA-diamond is a refined version of GPQA, focusing on challenging, high-quality scientific domain questions and designed to provide a comprehensive assessment of scientific understanding and reasoning ability.

## C Prompt Template

### Prompt for Batch Inference

Please answer the following math problems in order and summarize all answers at the end: Your response should be in the following format:

[Solution Process]

Provide a detailed solution for each problem...

[Final Answer]

1.  $\boxed{\text{Answer1}}$

2.  $\boxed{\text{Answer2}}$

...

n.  $\boxed{\text{Answern}}$

Below is the list of questions:

{numbered\_questions}

### Prompt for Evaluation

{origin\_question}\n\n

Please reason step by step, and put your final answer within  $\boxed{\}$ .