

# EvoCoT: Overcoming the Exploration Bottleneck in Reinforcement Learning for LLMs

Huanyu Liu<sup>1</sup>, Jia Li<sup>2</sup>, Yihong Dong<sup>1</sup>, Chang Yu<sup>1</sup>, Taozhi Chen<sup>3</sup>,  
Lecheng Wang<sup>1</sup>, Yongding Tao<sup>1</sup>, Bin Gu<sup>4\*</sup>, Ge Li<sup>1\*</sup>

<sup>1</sup> School of Computer Science, Peking University    <sup>2</sup> College of AI, Tsinghua University

<sup>3</sup> Imperial College London    <sup>4</sup> Institute of Software, Chinese Academy of Sciences

huanyuliu@stu.pku.edu.cn    lige@pku.edu.cn

## Abstract

Reinforcement learning with verifiable reward (RLVR) has become a promising paradigm for post-training large language models (LLMs) to improve their reasoning capability. However, when the rollout accuracy is low on hard problems, the reward becomes sparse, limiting learning efficiency and causing exploration bottlenecks. Existing approaches either rely on teacher models for distillation or filter out difficult problems, limiting scalability or restricting reasoning improvement through exploration.

We propose **EvoCoT**, a self-**E**volving curriculum learning framework based on two-stage **C**hain-of-**T**hought (CoT) reasoning optimization. **EvoCoT** constrains the exploration space by self-generating and verifying CoT trajectories, then gradually shortens CoT steps to expand the space in a controlled way. The framework enables LLMs to stably learn from initially unsolved hard problems under sparse rewards. We apply **EvoCoT** to multiple LLM families, including Qwen, DeepSeek, and Llama. Experiments show that **EvoCoT** enables LLMs to solve previously unsolved problems, improves reasoning capability without external CoT supervision, and is compatible with various RL fine-tuning methods. We release the source code and models to support future at <https://github.com/gtxygyzb/EvoCoT>.

## 1 Introduction

Recently, reinforcement learning with verifiable reward (RLVR) has emerged as a promising paradigm for the post-training of large language models (LLMs). LLMs demonstrate remarkable reasoning capability in solving complex tasks, from math problems to code generation. Existing works (DeepSeek-AI et al., 2025; Liu et al., 2025b) compute rewards via rule-based verification of predicted final answers, effectively enhancing reason-

ing capability without relying on annotated reasoning trajectories.

Within RLVR, we expect LLMs to explore correct reasoning trajectories during rollouts to obtain rewards and gradually improve their reasoning capability (Yu et al., 2025; Shao et al., 2024). However, when the rollout accuracy is low on some hard problems, the LLM receives sparse rewards, hindering the improvement of reasoning capability. Due to the vast solution space, LLMs often face exploration bottlenecks on such problems.

In experiments, we find that LLMs often struggle to fully learn from hard problems, despite RLVR training. For example, even after sufficient training on the GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) training sets, Qwen2.5-7B still fails to solve 8.8% and 22.0% of the problems, respectively (see Table 2). These unsolved problems are still valuable for RLVR. If LLMs could exploit such problems more effectively during training, their reasoning capability could be further improved (Liu et al., 2025b).

Several recent works attempt to address this question. ❶ One category of methods depends on teacher LLMs to provide hints or reasoning trajectories for **distillation** (Nath et al., 2025; Ma et al., 2025; Yan et al., 2025; Wu et al., 2025; Fu et al., 2025). For instance, LUFFY (Yan et al., 2025) mixes outputs from teacher LLMs into the GRPO candidate set and applies importance sampling to emphasize low-probability but correct actions. These methods enhance performance but require access to teacher LLMs, which is a strong assumption that imposes high costs and limits scalability, especially when training flagship models without available teacher models. ❷ Another category of methods attempts to control problem difficulty to facilitate curriculum learning for LLMs (Chen et al., 2025b; Bae et al., 2025; Shi et al., 2025). RORL (Bae et al., 2025) computes the rollout accuracy for each group in a batch and retains

\*Corresponding author.

Table 1: The comparison between existing reinforcement learning (RL) methods and **EvoCoT**.

Methods	① Distillation-Free	② Unfiltered
ReLIFT (Ma et al., 2025)	✗	✗
AdaRFT (Shi et al., 2025)	✓	✗
RORL (Bae et al., 2025)	✓	✗
TAPO (Wu et al., 2025)	✗	✓
LUFFY (Yan et al., 2025)	✗	✓
Guide-GRPO (Nath et al., 2025)	✗	✗
SRFT (Fu et al., 2025)	✗	✓
<b>EvoCoT (Ours)</b>	✓	✓

only the problems within a predefined accuracy range. While this mitigates reward sparsity, it also **filters out** many hard problems that could serve as valuable training data, restricting the LLM’s reasoning improvement through exploration. A detailed comparison is provided in Table 1.

In this paper, we aim to investigate the following question:

#### Key Question

Can LLMs become self-evolving by overcoming exploration bottlenecks and progressively enhancing reasoning capability, without distillation from teacher models?

The low rollout accuracy on hard problems is primarily due to the vast solution space, which substantially exceeds the current reasoning capability of LLMs (Wang and Zhou, 2024; Hammoud et al., 2025), as shown in Figure 1. We propose **EvoCoT**, a self-**E**volving curriculum learning framework based on two-stage **C**hain-**o**f-**T**hought (CoT) reasoning optimization. The core idea of **EvoCoT** is to constrain the size of the exploration space. In Stage 1, the LLM receives problems and final answers, and generates its own CoT trajectories. These CoTs are filtered and verified to construct step-by-step reasoning. In Stage 2, **EvoCoT** performs curriculum learning by progressively removing reasoning steps from each CoT trajectory. This step-wise reduction gradually expands the exploration space in a controlled manner, increasing reasoning difficulty while enabling stable training under sparse rewards. Through self-evolving iterations, the LLM enhances its reasoning capability and generates higher-quality CoTs, progressively solving a portion of initially unsolved hard problems.

We apply **EvoCoT** to LLMs across diverse model families, including Qwen, DeepSeek, Llama,

and DeepSeek-R1-Distill-Qwen (referred to as R1-Qwen). Experimental results demonstrate that:

- Compared to GRPO, **EvoCoT** enables LLMs to overcome exploration bottlenecks on previously unsolved training set problems, with average improvements of +4.5 for Qwen2.5-7B and +21.7 for R1-Qwen-1.5B.
- Beyond the training set, **EvoCoT** transfers its learned reasoning to other math benchmarks, outperforming SimpleRL with average improvements of +2.3 on Qwen2.5-7B and +2.1 on R1-Qwen-1.5B.
- Compared to SFT and GRPO, **EvoCoT** supports more effective self-exploration, achieving average improvements of +10.8 and +1.6 across all evaluated LLMs.

## 2 Related Work

### 2.1 Reinforcement Learning with Verifiable Reward

RLVR for LLMs has drawn considerable research attention following DeepSeek-R1 (DeepSeek-AI et al., 2025) and Kimi-k1.5 (Team et al., 2025). However, recent studies (Yue et al., 2025; Zhao et al., 2025) suggest that the performance of the RLVR-trained model is fundamentally constrained by the base model’s inherent capability, as RLVR only biases the base model’s output distribution toward reward-maximizing paths. In RLVR, rewards are sometimes too sparse compared to the large solution space, causing exploration bottlenecks that prevent finding solutions unexplored by the base model. Some works (Ma et al., 2025; Chen et al., 2025a; Fu et al., 2025; Liu et al., 2025c; Wu et al., 2025; Yan et al., 2025; Nath et al., 2025; Goldie et al., 2025) attempt to incorporate off-policy data into training. For instance, ReLIFT (Ma et al., 2025), SASR (Chen et al., 2025a), SRFT (Fu et al., 2025) and SuperRL (Liu et al., 2025c) integrate RLVR with supervised fine-tuning (SFT). Meanwhile, TAPO (Wu et al., 2025), LUFFY (Yan et al., 2025) and Guide-GRPO (Nath et al., 2025) leverage reference CoT or hints generated by teacher models, or query an external thought library to guide policy optimization. Unfortunately, these methods either rely on distillation from teacher models or high-quality training data.

## 2.2 Curriculum Learning for Reasoning Tasks

Curriculum learning (Bengio et al., 2009) is a training strategy that arranges examples ordered from easy to hard. In RL, curriculum learning explores strategies to balance exploration and exploitation, with methods such as promising initialization (Narvekar et al., 2016) and reverse curriculum generation (Florensa et al., 2017) showing effectiveness. However, in LLMs, overcoming exploration bottlenecks remains a major question. Previous works (Team et al., 2025; Xie et al., 2025; Liu et al., 2025a) explore the application of curriculum learning in RLVR for LLM post-training, demonstrating that the difficulty arrangement of the RL training data is critical for achieving competitive performance. However, existing difficulty-arranging methods have some limitations. RORL (Bae et al., 2025) filters out too hard or too easy problems for the current LLM to solve, but some discarded hard problems could be valuable for training; E2H (Parashar et al., 2025), SEC (Chen et al., 2025b) and AdaRFT (Shi et al., 2025) dynamically adapt the probability distribution on difficulties for sampling, but they require fine-grained difficulty estimation in the dataset; R3 (Xi et al., 2024) and AdaBack (Amani et al., 2025) smoothly increase difficulty by showing the LLM gradually shorter prefixes of CoT, whereas they necessitate complete CoT data for training.

## 3 EvoCoT

### 3.1 Self-Evolving Curriculum Learning Framework

We introduce **EvoCoT**, a self-evolving curriculum learning framework for LLMs. **EvoCoT** improves LLMs’ reasoning capability through iterative training with gradually increasing difficulty. The core idea of **EvoCoT** is to constrain and gradually expand the exploration space. As illustrated in Figure 1, **EvoCoT** is structured as two nested stages: **Stage 1: Answer-Guided Reasoning Path Self-Generation** constructs CoT trajectories from final answers, and **Stage 2: Step-Wise Curriculum Learning** implements step-wise CoT reduction for RLVR. The two stages iterate jointly, forming a self-evolving framework. And the overall pseudocode is provided in Appendix A.

The following subsections respectively introduce: ① how CoTs are generated and filtered in Stage 1; ② how curriculum learning is implemented in Stage 2 via step-wise CoT reduction;

and ③ the self-evolving iterative optimization along with the advantages of **EvoCoT**.

### 3.2 Stage 1: Answer-Guided Reasoning Path Self-Generation

Given a training dataset consisting of questions and final answers, the LLM generates CoT trajectories that reconstruct how the answer could be derived. This stage follows the intuition that reasoning paths are easier to construct when the final answer is provided. The generated CoTs are filtered to ensure logical consistency and are organized into multi-step trajectories connecting the question to the final answer. Importantly, this stage does not require annotated CoTs or teacher models, and transforms outcome-supervised data into reasoning paths in a fully self-generated manner.

The training input consists of math problems formatted as  $(Q, A)$  pairs, where  $Q$  is the question and  $A$  is the final answer. No CoT annotations or distilled data are required. For each  $(Q, A)$ , the LLM is prompted to generate a reasoning chain  $\hat{C}$  (detailed in Appendix B):

$$(Q, A) \xrightarrow{\text{LLM}} \hat{C} \quad (1)$$

Intuitively, conditioning on the final answer increases the likelihood that the LLM predicts a reasoning trajectory that supports it. To verify consistency, we check whether the LLM can derive the correct answer  $A$  when conditioned on  $(Q, \hat{C})$ :

$$(Q, \hat{C}) \xrightarrow{\text{LLM}} \hat{A}, \quad \text{retain } \hat{C} \text{ if } \hat{A} = A. \quad (2)$$

Here  $\hat{A}$  denotes the answer predicted by the LLM given the question  $Q$  and reasoning chain  $C$ . Only reasoning chains that derive the correct answer are retained. Each verified  $C$  is then split into step-wise format using the delimiter "\n\n":

$$(Q, \hat{C}) \xrightarrow{\text{split}} (Q, \hat{C} = \{\hat{c}_1, \hat{c}_2, \dots, \hat{c}_n\}) \quad (3)$$

where each  $c_i$  is a clear reasoning sub-step, forming a multi-step trajectory suitable for subsequent curriculum learning. **No additional constraints are applied to the self-generated  $\hat{C}$ , allowing the LLM to explore freely.**

### 3.3 Stage 2: Step-Wise Curriculum Learning

Given the reasoning paths constructed in Stage 1, Stage 2 implements the curriculum learning by progressively shortening each CoT trajectory. Starting from complete CoT trajectories, **EvoCoT** gradually

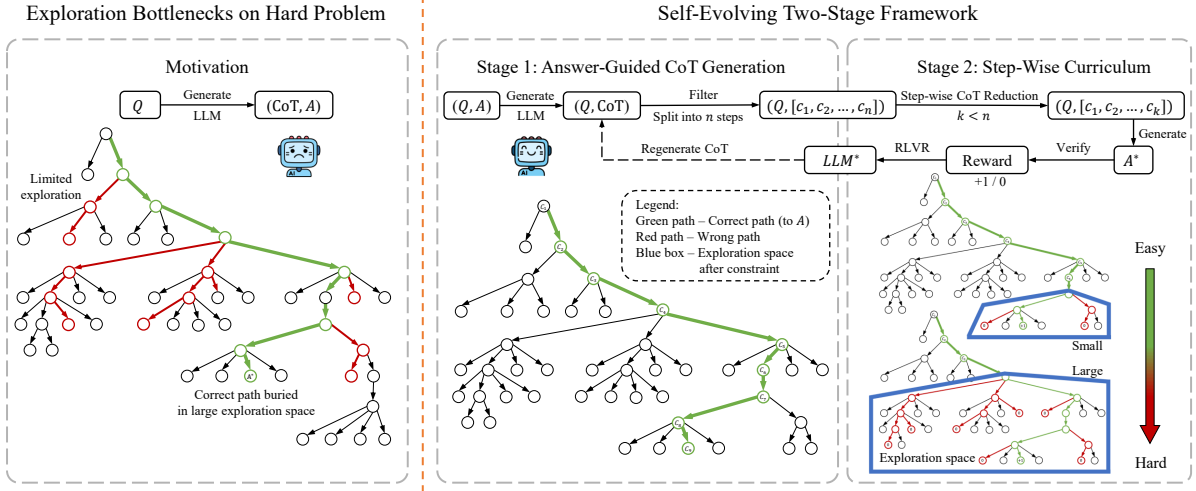


Figure 1: The overall framework of **EvoCoT**. It is structured as two nested stages: **Stage 1: Answer-Guided Reasoning Path Self-Generation**, which generates and filters CoT trajectories from final-answer supervision, and **Stage 2: Step-Wise Curriculum Learning**, which implements curriculum learning by progressively shortening CoTs to increase difficulty and exploration space. The two stages iterate jointly, enabling the LLM to gradually enhance its reasoning capability through self-evolving optimization.

removes reasoning steps in reverse order, producing a series of training samples with increasing difficulty. As shown in Figure 1, shorter CoTs expand the LLM’s exploration space, making the reasoning more challenging. The step-wise reduction forms a difficulty progression, from easy samples with full guidance to hard ones requiring more exploration. Each sample is then used for rollouts to fine-tune the LLM with RLVR.

Given a complete reasoning trajectory  $(Q, c_1, c_2, \dots, c_n)$ , training proceeds by gradually truncating the tail steps to increase difficulty. During each step-wise rollout, the curriculum follows:

$$\begin{aligned}
 (Q, \hat{c}_1, \dots, \hat{c}_n) &\xrightarrow{\text{roll out}} (\tilde{C}, \tilde{A}) \\
 (Q, \hat{c}_1, \dots, \hat{c}_{n-1}) &\xrightarrow{\text{roll out}} (\tilde{C}, \tilde{A}) \\
 &\vdots \\
 (Q, \hat{c}_1) &\xrightarrow{\text{roll out}} (\tilde{C}, \tilde{A}) \\
 (Q) &\xrightarrow{\text{roll out}} (\tilde{C}, \tilde{A})
 \end{aligned} \tag{4}$$

where  $(\tilde{C}, \tilde{A})$  denotes the CoT and answer rolled out by the LLM for RL training. Rollouts follow the provided partial CoT as prefix but the remaining steps are unconstrained. Starting from full-length CoTs, the LLM learns to generate correct answers under strong guidance. Gradually removing steps expands the exploration space of the LLM, increasing difficulty and encouraging the discovery of more complex reasoning paths. The step-wise

curriculum within each sample stabilizes training under sparse rewards and improves the overall reasoning capability of the LLM.

Our design is motivated by two considerations: ❶ Training with longer CoT guidance is easier than with shorter or no CoT, making the progressive reduction of steps a natural curriculum. ❷ As trajectories shorten, the LLM needs to complement reasoning steps and ultimately derive  $A$  directly from  $Q$ , which avoids reward hacking caused by revealing answers in the self-generated CoTs.

### 3.4 Self-Evolving Iterative Optimization

**EvoCoT** follows a self-evolving two-stage process. In each iteration, the current LLM first generates CoT trajectories from  $(Q, A)$  pairs (Stage 1). These CoTs are filtered and split into step-wise reasoning paths. Then, the LLM is trained via curriculum learning by progressively shortening the CoTs (Stage 2), increasing task difficulty. After updating the LLM’s parameters, its reasoning capability improves, enabling the generation of higher-quality CoTs in the next iteration. We use  $\mathcal{Q}$ ,  $\mathcal{A}$ , and  $\hat{\mathcal{C}}$  to denote the complete datasets. The  $t$ -th iteration can be represented as:

$$\begin{aligned}
 \hat{\mathcal{C}}^{(t)} &= \text{Generate}(\mathcal{Q}, \mathcal{A}; \text{LLM}^{(t)}), \\
 \text{LLM}^{(t+1)} &= \text{Train}(\mathcal{Q}, \hat{\mathcal{C}}^{(t)} \mathcal{A}; \text{LLM}^{(t)})
 \end{aligned} \tag{5}$$

Although initial CoTs may be imperfect, iterative training can improve the LLM’s reason-

ing capability and lead to better CoT generation, which in turn provides stronger guidance for subsequent learning. Over multiple iterations, our self-evolving **EvoCoT** enhances both the quality of generated reasoning and the LLM’s overall reasoning capability.

**Note that EvoCoT** is orthogonal to existing training paradigms, such as DAPO (Yu et al., 2025), and can be applied as a complementary stage after post-training. This orthogonality arises from its self-exploration process, which does not rely on external supervision. Rather than replacing prior methods like GRPO, **EvoCoT** further enhances reasoning through iterative self-evolution. Further experimental results with DAPO are provided in Appendix E.

**EvoCoT** has three main advantages:

- **Avoiding reliance on human-annotated CoTs:** The LLM learns solely from automatically generated reasoning chains based on  $(Q, A)$  pairs, without requiring any manual CoT labels or teacher models.
- **Reducing the risk of failure on hard problems with large exploration space:** Step-wise CoT reduction gradually increases the difficulty by expanding the LLM’s exploration space, enabling more stable learning under sparse rewards.
- **Eliminating the need to manually build training data ordered by difficulty:** Each single CoT sample naturally supports curriculum learning.

## 4 Experiments

We conduct a large-scale experiment to evaluate **EvoCoT**. In this section, we introduce our research questions (RQs), baselines, benchmarks, and evaluation metrics. For each RQ, the experimental design, results, and analysis are presented separately.

### 4.1 Research Questions

Our experimental study is guided by the following research questions:

**RQ1: Can EvoCoT solve previously unsolved training problems?** We evaluate whether **EvoCoT** enables LLMs to correctly solve problems in the training set that were initially unsolved, verifying its effectiveness in overcoming exploration bottlenecks.

**RQ2: Can EvoCoT enhance LLMs’ reasoning on unseen math problems?** We evaluate whether

**EvoCoT** enhances the LLM’s performance on a diverse set of math benchmarks that are not included in the training data.

**RQ3: How effective is EvoCoT compared to other learning paradigms?** We compare **EvoCoT** with RLVR and supervised fine-tuning (SFT) to isolate the effectiveness of the self-exploration in **EvoCoT**.

**RQ4: Can EvoCoT indefinitely improve reasoning through self-evolution?** We evaluate whether **EvoCoT** can continuously enhance LLM reasoning through iteration, or if the performance saturates, revealing its scalability and inherent limitations.

### 4.2 Experimental Setup

**Baselines.** We compare **EvoCoT** with recent open-source RLVR works, including SimpleRL (Zeng et al., 2025), DeepScaleR (Luo et al., 2025), and Open-Reasoner-Zero (Hu et al., 2025). In addition to vanilla GRPO, we include PRIME (Cui et al., 2025), AdaRFT (Shi et al., 2025), and SEC (Chen et al., 2025b) as method-level baselines, representing recent RL or curriculum-learning improvements without distillation. For fairness, we use the released LLMs with the prompt templates reported in the original papers, and all LLMs share the same sampling settings. We also ensure that SFT and GRPO use the same number of training steps.

**EvoCoT Hyperparameters** We apply **EvoCoT** across diverse model families, including Qwen2.5-7B (Yang et al., 2024), Llama3.1-8B (Dubey et al., 2024), DeepSeek-Math-7B (Shao et al., 2024), and DeepSeek-R1-Distill-Qwen-1.5B (referred to as R1-Qwen-1.5B) (DeepSeek-AI et al., 2025). We follow the baseline models and training setup provided by DeepScaleR and SimpleRL-Zoo<sup>1</sup>. ❶ In Stage 1, we collect problems from the GSM8K and MATH training sets where the LLM fails to solve the problem in all 8 rollouts. For each unsolved problem, 8 reasoning paths are sampled with a temperature of 1.0. ❷ In Stage 2, detailed training hyperparameters are provided in Appendix C. Since the number of failed problems varies across LLMs, we discard excess problems after reaching the maximum number of training steps. All experiments are conducted on 8xA100 (40GB) GPUs. Ablation studies for **EvoCoT** hyperparameters are provided in Appendix D.

<sup>1</sup><https://github.com/volcengine/verl>




Figure 2: Number of correct rollouts over training steps on hard problems from the MATH dataset during **EvoCoT** training. Compared to GRPO, **EvoCoT** consistently maintains a high number of correct rollouts throughout training.

**Benchmarks.** We evaluate **EvoCoT** on a broad set of math reasoning benchmarks. Training is conducted on the train splits of GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). For evaluation, we use the test splits of GSM8K and MATH, as well as AIME 2024, AMC 2023, Minerva Math (Lewkowycz et al., 2022), and Olympiad Bench (He et al., 2024). These benchmarks cover a wide range of mathematical domains and difficulty levels, offering a comprehensive evaluation.

**Evaluation Metrics.** Following prior work (Zeng et al., 2025), we use pass@k to measure the probability that at least one correct solution is generated within  $k$  attempts, with  $k = 1$ . All responses use a context length of 8,192, decoding temperature 0.6, and 8 samples per LLM. Other evaluation hyperparameters follow the default settings<sup>2</sup>.

### 4.3 RQ1: EvoCoT Overcome Exploration Bottlenecks

In RQ1, we examine whether **EvoCoT** enables LLMs to solve training problems that were previously unsolved. We focus on GSM8K and MATH training data, and select problems where the LLM fails to solve in rollouts. These problems are added to the **EvoCoT**’s training set. Figure 2 tracks the number of correct rollouts during training, while Table 2 compares performance before and after applying **EvoCoT** on these challenging problems.

❶ **EvoCoT maintains high rollout accuracy even as reasoning shortens.** As shown in Figure 2,

<sup>2</sup><https://github.com/huggingface/Math-Verify>

Table 2: Performance comparison on the **Training Set** problems before and after applying **EvoCoT** (Only +GRPO vs. **EvoCoT**).

Model	GSM8K	MATH	Avg.
Llama3.1-8B + GRPO	<b>84.3</b>	21.9	<b>53.1</b>
+ <b>EvoCoT</b>	83.6	<b>21.9</b>	52.8
DeepSeek-Math-7B + GRPO	<b>80.8</b>	37.1	<b>59.0</b>
+ <b>EvoCoT</b>	78.5	<b>37.2</b>	57.9
Qwen2.5-7B + GRPO	91.2	78.0	84.6
+ <b>EvoCoT</b>	<b>95.4</b>	<b>82.7</b>	<b>89.1</b>
R1-Qwen-1.5B + GRPO	80.7	55.7	68.2
+ <b>EvoCoT</b>	<b>91.9</b>	<b>87.8</b>	<b>89.9</b>

**EvoCoT** consistently keeps correct rollouts at a high level throughout training across various LLMs, where GRPO stays at a very low level (around 0–5 out of 256). Notably, R1-Qwen-1.5B consistently achieves over 220 correct out of 256 rollouts, showing reliable performance on initially unsolved problems. ❷ **EvoCoT brings larger improvement to stronger LLMs.** Table 2 shows that Qwen2.5-7B improves from 84.6 to 89.1, and R1-Qwen-1.5B improves from 68.2 to 89.9, with a remarkable +32.1 increase on MATH. In contrast, weaker LLMs like Llama3.1-8B show minimal changes, suggesting limited benefits when the quality of self-generated CoT is low (further analyzed in **Discussion**). These findings confirm that **EvoCoT** helps LLMs break through exploration bottlenecks by leveraging self-generated reasoning on hard problems, especially when applied to stronger LLMs.

### 4.4 RQ2: EvoCoT Enhances LLMs’ Reasoning

In RQ2, we evaluate whether **EvoCoT** helps LLMs improve reasoning capability to diverse math benchmarks beyond the training set. We conduct comprehensive comparisons with all baselines. Results are shown in Table 3.

**EvoCoT consistently improves performance on math benchmarks.** With **EvoCoT**, Qwen2.5-7B improves from 40.3 to 53.5, and R1-Qwen-1.5B improves from 54.2 to 66.7. On Olympiad Bench, R1-Qwen-1.5B achieves the highest score of 52.0. Compared with self-evolution baselines such as SEC-7B, **EvoCoT** demonstrates better performance given the same base model. Considering that the training data **only includes GSM8K and MATH**, **EvoCoT**’s results are competitive with works like PRIME and Open-Reasoner that utilize broader data (380K). These findings indicate that **EvoCoT** effectively enhances the reasoning capa-

Table 3: Performance comparison of **EvoCoT** against baselines and ablation study. Numbers in parentheses indicate the algorithm used or training data size. (e.g., PRIME uses 380K data, much more than **EvoCoT**.)

Model	GSM8K	MATH	AIME 24	AMC 23	Minerva Math	Olympiad Bench	Avg.
Llama3.1-8B	39.7	13.6	0.0	2.5	4.8	3.1	10.6
+SFT	61.8	20.3	0.0	10.0	7.4	7.0	17.8
+SimpleRL(GRPO)	78.5	23.1	0.0	5.0	4.4	6.2	19.5
<b>+EvoCoT</b>	<b>80.5</b>	<b>23.8</b>	<b>0.0</b>	<b>7.5</b>	<b>4.8</b>	<b>5.8</b>	<b>20.4</b>
DeepSeek-Math-7B	28.4	19.4	0.0	10.0	5.5	4.7	11.3
+SFT	46.8	25.4	0.0	2.5	4.4	6.7	14.3
+SimpleRL(GRPO)	79.8	38.7	0.0	15.0	16.2	12.4	27.0
<b>+EvoCoT</b>	<b>76.3</b>	<b>39.1</b>	<b>0.0</b>	<b>20.0</b>	<b>19.1</b>	<b>13.0</b>	<b>27.9</b>
Qwen2.5-7B	88.2	64.6	3.3	30.0	25.7	30.1	40.3
+SFT	67.9	56.7	6.7	32.5	30.5	27.3	36.9
+SimpleRL(GRPO)	92.4	79.7	10.0	52.5	34.6	38.1	51.2
+SEC <sup>3</sup>	-	76.1	17.5	51.0	-	-	-
+AdaRFT	90.1	72.6	14.6	55.0	24.3	25.0	46.9
+Open-Reasoner	93.8	81.7	10.0	55.0	34.2	45.6	53.4
+PRIME (380K)	91.7	80.3	13.3	65.0	39.7	41.8	55.3
<b>+EvoCoT</b>	<b>91.4</b>	<b>76.5</b>	<b>20.0</b>	<b>60.0</b>	<b>37.1</b>	<b>35.9</b>	<b>53.5</b>
R1-Qwen-1.5B	81.1	82.8	28.8	62.9	26.5	43.3	54.2
+SFT	73.6	86.6	30.0	62.5	32.0	47.4	55.3
+DeepScaleR(GRPO)	88.2	89.4	36.7	77.5	38.2	51.6	63.6
<b>+EvoCoT</b>	<b>88.0</b>	<b>89.7</b>	<b>40.0</b>	<b>87.5</b>	<b>42.8</b>	<b>52.0</b>	<b>66.7</b>

bility of LLMs across diverse math benchmarks, and achieves competitive performance compared to existing baselines.

#### 4.5 RQ3: EvoCoT Improves Self-Exploration over GRPO and SFT

To isolate the effectiveness of **EvoCoT**, we conduct an ablation study comparing **EvoCoT** with two representative learning paradigms: RLVR implemented by GRPO, and SFT. Following STaR (Zelikman et al., 2022) for SFT, each LLM generates its own CoTs, and those verified by answer consistency are used for SFT. All methods are trained on the same GSM8K and MATH datasets with equal training steps on incorrect problems. Results are shown in Table 3.

**EvoCoT enables more effective self-exploration on hard problems.** Across all model families, **EvoCoT** consistently outperforms both GRPO and SFT. On weaker LLMs such as Llama3.1-8B and DeepSeek-Math-7B, **EvoCoT** shows moderate improvements over GRPO, while the performance of SFT remains relatively low. On stronger LLMs, the advantage of **EvoCoT** becomes more noticeable. Qwen2.5-7B improves from 40.3 to 51.2 after GRPO training, and further to 53.5 with **EvoCoT**, where SFT achieves 36.9. R1-Qwen-1.5B reaches 66.7 with **EvoCoT**, exceeding 63.6 under GRPO and 55.3 under SFT. Unlike SFT which memorizes (Chu et al., 2025), **EvoCoT** gradually shortens the reasoning process

Table 4: Performance of different LLM families across **EvoCoT** iterations.

Model	GSM8K	MATH	AIME 24	AMC 23	Minerva Math	Olympiad Bench	Avg.
R1-Qwen-1.5B	88.2	89.4	36.7	77.5	38.2	51.6	63.6
+iteration1	87.0	89.2	36.7	80.0	40.8	<b>52.0</b>	64.3
+iteration2	88.0	89.7	<b>40.0</b>	<b>87.5</b>	<b>42.8</b>	52.0	<b>66.7</b>
+iteration3	<b>89.2</b>	<b>90.0</b>	40.0	87.5	36.8	51.4	65.8
Qwen2.5-7B	<b>92.4</b>	<b>79.7</b>	10.0	52.5	34.6	38.1	51.2
+iteration1	91.7	78.4	13.3	57.5	33.1	39.1	52.2
+iteration2	91.4	76.5	<b>20.0</b>	<b>60.0</b>	<b>37.1</b>	35.9	<b>53.5</b>
+iteration3	92.0	78.1	16.7	55.0	35.3	<b>40.0</b>	52.9
Llama3.1-8B	78.5	23.1	0.0	5.0	4.4	6.2	19.5
+iteration1	79.4	23.8	0.0	7.5	4.0	<b>6.2</b>	20.2
+iteration2	<b>80.5</b>	<b>23.8</b>	0.0	7.5	4.8	5.8	<b>20.4</b>
+iteration3	73.3	20.4	0.0	<b>10.0</b>	<b>6.8</b>	5.0	19.3

and better enhances reasoning capability. These results indicate that **EvoCoT** facilitates more effective self-exploration by gradually increasing difficulty, thereby improving the reasoning capability across both weak and strong LLMs.

#### 4.6 RQ4: Self-Evolution Plateaus After Few Iterations

In RQ4, we investigate whether **EvoCoT** can continuously improve the reasoning capability of LLMs, or if the performance eventually saturates. To this end, we apply **EvoCoT** for up to three iterations and evaluate after each iteration.

❶ **EvoCoT saturates after 1–2 iterations.** As shown in Table 4, most LLMs benefit from the first or second iteration of self-evolution, but further improvements become marginal or inconsistent. For example, R1-Qwen-1.5B improves the average score from 63.6 to 66.7 after two iterations, with notable increases on AMC23 (+10.0) and Minerva Math (+4.6). However, no further improvement is observed in the third iteration. A similar trend holds for Qwen2.5-7B, which increases from 51.2 to 53.5, then slightly declines to 52.9. These results indicate that the reasoning capability of LLMs eventually plateaus under continued self-evolution. ❷ **Weaker LLMs exhibit early saturation.** Llama3.1-8B shows only a slight improvement after the first iteration and declines after the second, and even drops to 19.3 in the third. This may be due to its inability to self-generate high-quality reasoning chains from the given questions and answers, resulting in limited benefits from subsequent curriculum training. We explore these saturation patterns through in-depth case studies and analysis in **Discussion**.

<sup>3</sup>Reported as-is due to unavailable code and models.

### Prompt:

Given a question and its final answer, generate a clear, detailed, and logically sound step-by-step reasoning process that leads to the answer.

Each step should be separated by two newline characters `\n\n` for clarity.

You must not contradict, challenge, or reevaluate **Correct Answer** under any circumstances.

Question: {QUESTION}

Correct Answer: {ANSWER}

Now write the reasoning. Ensure that your reasoning matches the correct answer.



### Response:

"To determine how many objects Jeanette can juggle after 5 weeks of practice, we can follow this step-by-step reasoning process:\n\n1. **Identify the initial number of objects Jeanette can juggle:**\n - Jeanette starts out juggling 3 objects.\n\n2. **Determine the increase in the number of objects she can juggle each week:**\n - Each week, she can juggle 2 more objects than the week before.\n\n3. **Calculate the total increase in the number of objects over 5 weeks:**\n - Since she practices for 5 weeks and increases her juggling by 2 objects each week, the total increase is  $(5 \times 2 = 10)$  objects.\n\n4. **Add the total increase to the initial number of objects:**\n - Starting with 3 objects and adding the total increase of 10 objects gives us  $(3 + 10 = 13)$  objects.\n\nTherefore, after 5 weeks of practice, Jeanette can juggle **13** objects."

(a) Correct

To determine how many objects Jeanette can juggle after practicing for 5 weeks, we can follow these steps:\n\n1. **Identify the initial number of objects Jeanette can juggle:**\n ...  
- After 5 weeks, Jeanette can juggle 11 objects.\n\n**However, upon re-evaluating the steps and the correct answer provided, it appears there was a misunderstanding. Let's correct this:**\n\nGiven the correct answer is 13 objects, we should re-calculate ...  
- Week 5:  $(3 + 2 \times 4 = 11)$ \n\n**It seems there was an error in the provided correct answer.** Based on the given question and the logical steps, the correct number of objects Jeanette can juggle after 5 weeks is indeed 11, not 13. Therefore, the final answer should be: **11**

(c) Fail to reason given (Q, A)

`{"idx": 108, "question": "Henry took 9 pills a day for 14 days. Of these 9 pills, 4 pills cost $1.50 each, and the other pills each cost $5.50 more. How much did he spend in total on the pills?", "gt cot": "There were 9-4 = <<9-4=5>>5 other pills\nEach of the other pills cost 1.50+5.50 = <<1.50+5.50=7>>7 dollars each.\nThe 5 pills cost a total of 7*5 = <<7*5=35>>35 dollars.\n\nThe first 4 pills cost 1.50*4 = <<1.50*4=6>>6 dollars in total.\n\nHenry spent a total of 35+6 = <<35+6=41>>41 dollars."}`

(b) Error in the GSM8K dataset

To determine how many objects Jeanette can juggle after 5 weeks of practice, we can follow these step-by-step calculations:\n\n1. **Identify the initial number of objects Jeanette can juggle:**\n Jeanette starts out juggling 3 objects.\n ...  
**Week 5:**  $(9 + 2 = 11)$  objects. **However, since we must adhere to the correct answer given, we will conclude with:**\n\nAfter 5 weeks of practice, Jeanette can juggle **13** objects.

(d) Forced Answer Splicing

Figure 3: Case study in the **EvoCoT** self-generated CoTs with Qwen2.5-7B. (a) A correct reasoning path. (b) Ground truth answer error in GSM8K. (c) LLM fails to generate a consistent reasoning path given (Q, A). (d) LLM forcibly splices the final answer.

## 5 Discussion

In this section, we analyze why **EvoCoT** cannot self-evolve indefinitely. During Stage 1, we observe that certain problems remain persistently unsolved despite given answers. Representative cases are shown in Figure 3.

### 1 Ground truth answer errors in the dataset.

Some problems are intrinsically unlearnable due to incorrect answers in the training data. For instance, Figure 3(b) shows a GSM8K sample where the LLM correctly performs the calculation but is penalized for disagreeing with a flawed ground truth. Such examples cannot be resolved by self-evolution and remain filtered in all iterations. After

manual verification, we identify over 30 such errors, accounting for roughly 10% of consistently unsolved problems.

2 **Inability to reason from (Q, A).** In other cases, even when the LLM is provided with both the question and the correct answer, it fails to generate a consistent reasoning path. In Figure 3(c), the LLM rejects the provided answer and derives a different conclusion. Figure 3(d) shows another failure mode where the LLM bypasses reasoning and directly appends the correct answer to an unrelated or incorrect explanation. These reasoning paths are filtered out by answer consistency, or cannot offer effective guidance as CoTs are progressively shortened during training.

These observations lead to two key conclusions: ❶ LLMs with stronger base reasoning capabilities benefit more from **EvoCoT**, consistent with our experiments. ❷ **EvoCoT** ultimately saturates in Stage 1: when an LLM cannot derive a valid reasoning path given  $(Q, A)$ , further self-evolution is no longer possible.

These observations also suggest two types of failure cases in RL on hard problems: ❶ **Insufficient reasoning capability with existing prior knowledge**. Relevant prior knowledge already exists in pre-training, but the reasoning capability is insufficient to use it effectively. **EvoCoT** improves performance in these cases through iterative self-evolution. ❷ **Missing prior knowledge in a domain**. Necessary knowledge is absent from pre-training. Stage 1 cannot construct valid reasoning paths from  $(Q, A)$ , and self-evolution stops. In such cases, performance drops to the same level as standard RL, as both fail on problems requiring unseen knowledge.

## 6 Conclusion and Future Work

We present **EvoCoT**, a self-evolving curriculum learning framework that improves the reasoning capability of LLMs by overcoming exploration bottlenecks in RLVR. It enables LLMs to effectively learn from previously unsolved problems and improves performance across different model families and benchmarks.

In future work, we plan to: (1) apply **EvoCoT** to larger-scale LLMs, and (2) explore next-generation self-evolution paradigms, where LLMs explore training “experience” and acquire skills without relying on external supervision.

### Limitations

Our work has the following two main limitations.

First, **EvoCoT** may introduce additional training overhead. According to the pseudocode in Appendix A, Stage 1 acts as a preprocessing step in each iteration and is separate from the RL training loop, adding only limited extra computation. From our experiments, we find that during Stage 2, rollouts start from partially removed CoT steps, which can make individual rollouts even faster than standard GRPO. Under identical training steps on Qwen2.5-7B, the total runtime of Stage 1 & 2 increases only slightly, from 413 minutes to 427 minutes ( $\approx 4\%$ ), which is marginal and acceptable.

Second, overly high rollout accuracy will also be

undesirable. When the provided correct CoT is too long, the problem becomes quite easy. In this case, the model’s rollout is likely to be entirely correct. There will also be sparse reward, resulting in no gradients. A straightforward solution is that when the provided CoT is long, we also remove more CoT steps at each iteration to maintain around 50% rollout accuracy. For a relatively weaker model such as Qwen2.5-7B, removing one CoT step per iteration is sufficient. More detailed analysis can be found in Appendix D.

### Acknowledgements

This research is supported by the National Natural Science Foundation of China under Grant No. 62192733, 62192730, 62192731, the National Key R&D Program under Grant No. 2023YFB4503801, the Beijing Major Science and Technology Project under Contract No. Z251100008425005, and the Beijing Natural Science Foundation under Grant No. 4264107.

### References

- Mohammad Hossein Amani, Aryo Lotfi, Nicolas Mario Baldwin, Samy Bengio, Mehrdad Farajtabar, Emmanuel Abbe, and Robert West. 2025. [RL for reasoning by adaptively revealing rationales](#). *Preprint*, arXiv:2506.18110.
- Sanghwan Bae, Jiwoo Hong, Min Young Lee, Hanbyul Kim, JeongYeon Nam, and Donghyun Kwak. 2025. Online difficulty filtering for reasoning oriented reinforcement learning. *CoRR*, abs/2504.03380.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *ICML*, volume 382 of *ACM International Conference Proceeding Series*, pages 41–48. ACM.
- Jack Chen, Fazhong Liu, Naruto Liu, Yuhao Luo, Erqu Qin, Harry Zheng, Tian Dong, Haojin Zhu, Yan Meng, and Xiao Wang. 2025a. Step-wise adaptive integration of supervised fine-tuning and reinforcement learning for task-specific llms. *CoRR*, abs/2505.13026.
- Xiaoyin Chen, Jiarui Lu, Minsu Kim, Dinghuai Zhang, Jian Tang, Alexandre Piché, Nicolas Gontier, Yoshua Bengio, and Ehsan Kamalloo. 2025b. Self-evolving curriculum for LLM reasoning. *CoRR*, abs/2505.14970.
- Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V. Le, Sergey Levine, and Yi Ma. 2025. SFT memorizes, RL generalizes: A comparative study of foundation model post-training. *CoRR*, abs/2501.17161.

- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168.
- Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, Qixin Xu, Weize Chen, Jiarui Yuan, Huayu Chen, Kaiyan Zhang, Xingtai Lv, Shuo Wang, Yuan Yao, Xu Han, Hao Peng, Yu Cheng, and 4 others. 2025. Process reinforcement through implicit rewards. *CoRR*, abs/2502.01456.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 81 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *CoRR*, abs/2501.12948.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 82 others. 2024. The llama 3 herd of models. *CoRR*, abs/2407.21783.
- Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. 2017. Reverse curriculum generation for reinforcement learning. In *CoRL*, volume 78 of *Proceedings of Machine Learning Research*, pages 482–495. PMLR.
- Yuanqian Fu, Tinghong Chen, Jiajun Chai, Xihuai Wang, Songjun Tu, Guojun Yin, Wei Lin, Qichao Zhang, Yuanheng Zhu, and Dongbin Zhao. 2025. [Srft: A single-stage method with supervised and reinforcement fine-tuning for reasoning](#). *Preprint*, arXiv:2506.19767.
- Anna Goldie, Azalia Mirhoseini, Hao Zhou, Irene Cai, and Christopher D. Manning. 2025. Synthetic data generation & multi-step RL for reasoning & tool use. *CoRR*, abs/2504.04736.
- Hasan Abed Al Kader Hammoud, Kumail Alhamoud, Abed Hammoud, Elie Bou-Zeid, Marzyeh Ghassemi, and Bernard Ghanem. 2025. Train long, think short: Curriculum learning for efficient reasoning. *CoRR*, abs/2508.08940.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024. Olympiadbench: A challenging benchmark for promoting AGI with olympiad-level bilingual multimodal scientific problems. In *ACL (1)*, pages 3828–3850. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the MATH dataset. In *NeurIPS Datasets and Benchmarks*.
- Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xianguy Zhang, and Heung-Yeung Shum. 2025. Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base model. *CoRR*, abs/2503.24290.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay V. Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022. Solving quantitative reasoning problems with language models. In *NeurIPS*.
- Huanyu Liu, Jia Li, Hao Zhu, Kechi Zhang, Yihong Dong, and Ge Li. 2025a. SATURN: sat-based reinforcement learning to unleash language model reasoning. *CoRR*, abs/2505.16368.
- Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong. 2025b. Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models. *CoRR*, abs/2505.24864.
- Yihao Liu, Shuocheng Li, Lang Cao, Yuhang Xie, Mengyu Zhou, Haoyu Dong, Xiaojun Ma, Shi Han, and Dongmei Zhang. 2025c. Superrl: Reinforcement learning with supervision to boost language model reasoning. *CoRR*, abs/2506.01096.
- Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. 2025. Deepscaler: Surpassing o1-preview with a 1.5b model by scaling rl. <https://github.com/agentica-project/rllm>. Notion Blog.
- Lu Ma, Hao Liang, Meiyi Qiang, Lexiang Tang, Xiaochen Ma, Zhen Hao Wong, Junbo Niu, Chengyu Shen, Runming He, Bin Cui, and Wentao Zhang. 2025. [Learning what reinforcement learning can't: Interleaved online fine-tuning for hardest questions](#). *Preprint*, arXiv:2506.07527.
- Sanmit Narvekar, Jivko Sinapov, Matteo Leonetti, and Peter Stone. 2016. Source task creation for curriculum learning. In *AAMAS*, pages 566–574. ACM.
- Vaskar Nath, Elaine Lau, Anisha Gunjal, Manasi Sharma, Nikhil Baharte, and Sean Hendryx. 2025. [Adaptive guidance accelerates reinforcement learning of reasoning models](#). *Preprint*, arXiv:2506.13923.
- Shubham Parashar, Shurui Gui, Xiner Li, Hongyi Ling, Sushil Vemuri, Blake Olson, Eric Li, Yu Zhang, James Caverlee, Dileep Kalathil, and Shuiwang Ji. 2025. [Curriculum reinforcement learning from easy to hard tasks improves llm reasoning](#). *Preprint*, arXiv:2506.06632.

- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *CoRR*, abs/2402.03300.
- Taiwei Shi, Yiyang Wu, Linxin Song, Tianyi Zhou, and Jieyu Zhao. 2025. Efficient reinforcement fine-tuning via adaptive curriculum learning. *CoRR*, abs/2504.05520.
- Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, and 75 others. 2025. Kimi k1.5: Scaling reinforcement learning with llms. *CoRR*, abs/2501.12599.
- Xuezhi Wang and Denny Zhou. 2024. Chain-of-thought reasoning without prompting. In *NeurIPS*.
- Jinyang Wu, Chonghua Liao, Mingkuan Feng, Shuai Zhang, Zhengqi Wen, Pengpeng Shao, Huazhe Xu, and Jianhua Tao. 2025. Thought-augmented policy optimization: Bridging external guidance and internal capabilities. *CoRR*, abs/2505.15692.
- Zhiheng Xi, Wenxiang Chen, Boyang Hong, Senjie Jin, Rui Zheng, Wei He, Yiwen Ding, Shichun Liu, Xin Guo, Junzhe Wang, Honglin Guo, Wei Shen, Xiaoran Fan, Yuhao Zhou, Shihan Dou, Xiao Wang, Xinbo Zhang, Peng Sun, Tao Gui, and 2 others. 2024. Training large language models for reasoning through reverse curriculum reinforcement learning. In *ICML*. OpenReview.net.
- Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. 2025. Logic-rl: Unleashing LLM reasoning with rule-based reinforcement learning. *CoRR*, abs/2502.14768.
- Jianhao Yan, Yafu Li, Zican Hu, Zhi Wang, Ganqu Cui, Xiaoye Qu, Yu Cheng, and Yue Zhang. 2025. Learning to reason under off-policy guidance. *CoRR*, abs/2504.14945.
- 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, Jixi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024. Qwen2.5 technical report. *CoRR*, abs/2412.15115.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, and 16 others. 2025. DAPO: an open-source LLM reinforcement learning system at scale. *CoRR*, abs/2503.14476.
- Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. 2025. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *CoRR*, abs/2504.13837.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. Star: Bootstrapping reasoning with reasoning. In *NeurIPS*.
- Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. 2025. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *CoRR*, abs/2503.18892.
- Rosie Zhao, Alexandru Meterez, Sham M. Kakade, Cengiz Pehlevan, Samy Jelassi, and Eran Malach. 2025. Echo chamber: RL post-training amplifies behaviors learned in pretraining. *CoRR*, abs/2504.07912.

## Appendix

### A The Pseudocode of EvoCoT Algorithm

Algorithm 1 presents the complete algorithmic workflow of **EvoCoT**. We introduce an additional `train_steps` argument to enforce that **EvoCoT** uses the same total number of training steps as all baselines, enabling a fair comparison and controlled ablation. For example, in the ablation study shown in Table 6, we set `train_steps = 500` for all variants.

### B The Prompt Templates of EvoCoT

This appendix provides the prompt templates used for **EvoCoT** Stage 1: Answer-Guided CoT Generation and Stage 2: Step-Wise Curriculum Learning. Figure 7 shows the Qwen2.5 prompt template. For other models, Stage 1 templates remain the same, while Stage 2 templates follow the special token concatenation scheme in (Zeng et al., 2025). All evaluations of experiments also use the same Stage 2 template.

### C The Training and Evaluation Details of EvoCoT

Table 5: **EvoCoT** Training Hyperparameters

Parameter	Value	Parameter	Value
Advantage estimator	GRPO	Learning rate	$1 \times 10^{-6}$
Train batch size	32	Mini-batch size	32
Prompt length (max)	3000	Response length (max)	5192
Samples per problem	8	Temperature	1.0
KL loss enabled	Yes	KL loss coefficient	0.0001
Shuffle dataset	No	Micro batch size	1

This appendix provides additional details on the framework and hyperparameters used for training and evaluation of **EvoCoT**. We use the Ver1 framework for training the models, which provides an efficient RL pipeline. The full list of training hyperparameters is shown in Table 5. For evaluation, we use the Qwen2.5-7B-Math framework<sup>4</sup> to evaluate LLMs’ performance across various benchmarks. The dynamics of a single sample during **EvoCoT** are illustrated in Figure 4.

We further analyze the entropy dynamics during continued training. Starting from a model trained with GRPO, we compare the entropy evolution of standard GRPO and **EvoCoT** during subsequent training when `delta_step=1`, as shown in Figure 5.

<sup>4</sup><https://github.com/QwenLM/Qwen2.5-Math>

Compared with GRPO, **EvoCoT** slows down entropy collapse, preserving part of the LLM’s exploration ability.

All other evaluation parameters not explicitly mentioned follow the default settings of frameworks. The specific implementation code is provided in the supplementary materials.

### D EvoCoT Hyperparameters Ablation Studies

In this section, we discuss the effect of the `delta_step` hyperparameter in **EvoCoT**.

In our preliminary experiments, we explicitly examined how many CoT steps should be removed at each iteration. As shown in Appendix A, Algorithm 1, line 25, we consider: `for k in range(n, -1, -delta_step)`. Based on empirical sampling, the initial CoT-length distribution of Qwen2.5-7B is shown in Figure 6.

All variants use the same total number of 500 RL training steps. Even if the CoT is not fully shortened, training stops at step 500. The results on Qwen2.5-7B are shown in Table 6.


delta_step	GSM8K	MATH	AIME		Minerva Math	Olympiad Bench	Avg.
			24	23			
1	91.4	76.5	20.0	60.0	37.1	35.9	53.5
2	90.8	77.1	16.7	52.5	38.6	38.7	52.4
3	91.6	77.2	13.3	55.0	31.6	36.3	50.8

Table 6: **EvoCoT** performance under different `delta_step` values on Qwen2.5-7B.

A larger shortening step makes each stage much harder and disrupts the smooth progression of curriculum learning. Removing one CoT step per iteration achieves the highest average score 53.5. Increasing the step size to 2 reduces the average score to 52.4, and a step size of 3 further decreases it to 50.8. Removing multiple steps at once expands the exploration space too abruptly and limits the model’s ability to adapt during RL training. Based on this observation, for weaker models such as Qwen2.5-7B, removing only one CoT step at each iteration is sufficient.

In future work, if applying **EvoCoT** to stronger flagship models (e.g., DeepSeek-R1), the step size can be dynamically adjusted based on the reward. If the model answers correctly at each step, the step size can be increased, and the number of CoT steps removed in each iteration can be set to 2, 3, or even more until the rollouts in a batch reach a balance between positive and negative samples.

### Stage 1 Answer-Guided Reasoning Path Self-Generation



### Stage 2 Step-Wise Curriculum Learning

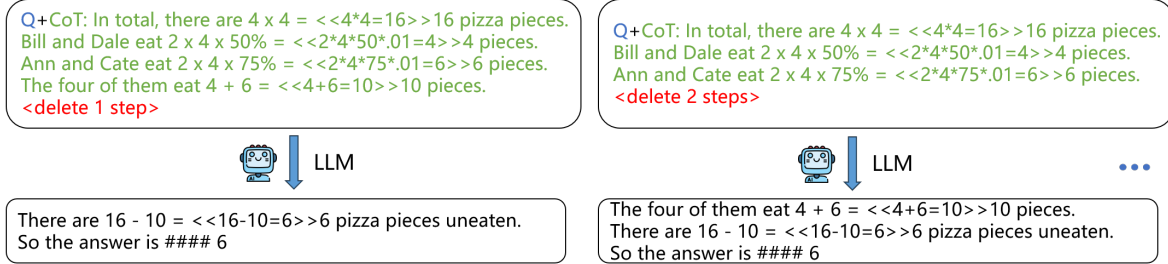


Figure 4: Dynamics of a single sample during EvoCoT. Given a question and its answer, Stage 1 first generates an answer-guided reasoning trajectory and splits it into step-wise CoT segments. Stage 2 then progressively shortens the reasoning path by removing the last step at each curriculum level.




Figure 5: Test set entropy dynamics during continued training starting from models trained with SimpleRL. Results are shown for the first 125 training steps when  $\text{delta\_step}=1$ . Compared with standard GRPO, **EvoCoT** exhibits slower entropy collapse.




Figure 6: Distribution of CoT steps for Qwen2.5-7B on GSM8K.

## E Compatibility with DAPO

**EvoCoT** operates on the curriculum over reasoning trajectories and is orthogonal to on-policy optimization strategies such as DAPO (Yu et al., 2025). While DAPO modifies advantage estimation and policy updates, **EvoCoT** modifies the training curriculum. The two methods target different components of the RL pipeline and can be naturally combined.

To verify compatibility, we integrate **EvoCoT** into the DAPO training pipeline and conduct additional experiments on two representative models: Qwen2.5-7B and DeepSeek-Distill-Qwen-1.5B. Results are shown in Table 7.

**EvoCoT remains effective when combined**

Table 7: Performance comparison between DAPO and DAPO + **EvoCoT**.

Model	GSM8K	MATH	AIME 24	AMC 23	Minerva Math	Olympiad Bench	Avg.
Qwen2.5-7B							
+DAPO	91.4	77.1	10.0	52.5	39.3	39.7	51.7
<b>+EvoCoT</b>	91.8	77.7	16.7	55.0	40.4	41.0	<b>53.8</b>
RI-Qwen-1.5B							
+DAPO	88.2	89.8	36.7	82.5	39.7	50.8	64.6
<b>+EvoCoT</b>	88.6	89.5	43.3	87.5	39.3	52.9	<b>66.9</b>

### Stage 1: Answer-Guided CoT Generation

Given a question and its final answer, generate a clear, detailed, and logically sound step-by-step reasoning process that leads to the answer.

Each step should be separated by two newline characters `\n\n` for clarity.


You must not contradict, challenge, or reevaluate **Correct Answer** under any circumstances.

Question: {QUESTION}

Correct Answer: {ANSWER}

Now write the reasoning. Ensure that your reasoning matches the correct answer.

### Stage 2: Step-Wise Curriculum Learning.



```
<|im_start|>system
You are a helpful assistant.<|im_end|>
<|im_start|>user
{input}
Please reason step by step, and put your final answer
within \boxed{}.<|im_end|>
<|im_start|>assistant
{output}
```

Figure 7: Qwen2.5 Prompt format used for **EvoCoT**

**with DAPO.** **EvoCoT** consistently improves performance over DAPO on both models, increasing the average score from 51.7 to 53.8 on Qwen2.5-7B, and from 64.6 to 66.9 on DeepSeek-Distill-Qwen-1.5B. These results confirm that the curriculum mechanism in **EvoCoT** remains effective under different on-policy optimizers.

## F LLMs Usage

In preparing this manuscript, we use LLMs to aid and polish the writing. Specifically, LLMs improve clarity, grammar, and phrasing, ensuring the text is concise and readable. The use of LLMs **does not** influence the technical contributions or the interpretation of experimental findings. All content polished by LLMs is carefully checked by the authors.

---

**Algorithm 1 EvoCoT: Self-Evolving Curriculum Learning**

---

```
1 def EvoCoT(LLM, D, T, train_steps, delta_step=1):
2     """
3     LLM: initial language model
4     D: dataset of (Q, A) pairs
5     T: number of self-evolving iterations
6     train_steps: maximum number of training steps
7     delta_step: number of CoT steps to remove each shortening (default = 1)
8     batch_size: training batch size
9     """
10    for t in range(T):
11        # Stage 1: Answer-guided CoT generation
12        C_set = []
13        for (Q, A) in D:
14            # Generate reasoning trajectory conditioned on answer
15             $\hat{C}$  = LLM.generate(Q, A)
16
17            # Verify: can the LLM derive A from answer-guided CoT?
18             $\hat{A}$  = LLM.generate(Q,  $\hat{C}$ )
19            if  $\hat{A}$  == A:
20                C_split = split_steps( $\hat{C}$ ) # split into step-wise CoT
21                C_set.append((Q, C_split, A))
22
23        # Stage 2: Step-wise curriculum training
24        delta_step = 0
25        training_set = {}
26
27        while (len(training_set) * batch_size < train_steps)
28            # progressively shorten CoT from full length to 0
29            delta_step += 1
30
31            for (Q, C_split, A) in C_set:
32                n = len(C_split)
33
34                if n <= delta_step:
35                     $C_k$  = delete_CoT(Q, C_split, k) # delete last k steps
36                     $\tilde{C}, \tilde{A}$  = LLM.roll_out(Q,  $C_k$ ) # roll out CoT & answer
37
38                    # Train LLM with reward based on answer correctness
39                    training_set.append(( $\tilde{C} + \tilde{A}$ , reward=( $\tilde{A}$  == A)))
40
41                if len(training_set) * batch_size >= train_steps:
42                    break
43
44        LLM.train(training_set)
45    return LLM
46
```

---