

Discovery and Reinforcement of Tool-Integrated Reasoning Chains via Rollout Trees

Kun Li^{1,2,*†}, Zenan Xu^{2*}, Junan Li¹, Zengrui Jin³,
Jinghao Deng⁴, Zexuan Qiu^{1,2}, Bo Zhou^{2‡}

¹The Chinese University of Hong Kong ²LLM Department, Tencent

³Tsinghua University ⁴Tencent

li.kun@link.cuhk.edu.hk {zenanxu, chaysezhou}@tencent.com

Abstract

Tool-Integrated Reasoning has emerged as a key paradigm to augment Large Language Models (LLMs) with computational capabilities, yet integrating tool-use into long Chain-of-Thought (long CoT) remains underexplored, largely due to the scarcity of training data and the challenge of integrating tool-use without compromising the model’s intrinsic long-chain reasoning. In this paper, we introduce **DART** (**D**iscovery **A**nd **R**einforcement of **T**ool-Integrated Reasoning Chains via **R**ollout Trees), a reinforcement learning framework that enables spontaneous tool-use during long CoT reasoning without additional human annotation. DART operates by constructing dynamic rollout trees during training to *discover* valid tool-use opportunities, branching out at promising positions to explore tool-integrated trajectories. Subsequently, a tree-based process advantage estimation identifies and credits specific sub-trajectories where tool invocation positively contributes to the solution, effectively *reinforcing* these beneficial behaviors during training. Extensive experiments on challenging benchmarks like AIME and GPQA-Diamond demonstrate that DART significantly outperforms existing methods, successfully harmonizing tool execution with long CoT reasoning.

1 Introduction

Tool-Integrated Reasoning (TIR) (Schick et al., 2023; Gou et al., 2024; Chen et al., 2025a) has emerged as a promising paradigm to augment large language models (LLMs) with computational capabilities, addressing inherent limitations in precise calculation, symbolic manipulation, and complex equation solving that hinder pure natural language reasoning approaches.

However, existing TIR efforts have predominantly focused on short Chain-of Thought (CoT)

(Wei et al., 2022), largely due to the modest level of implementation complexity—they can be implemented via prompting (Chen et al., 2023; Zhang et al., 2024), supervised finetuning (SFT) with annotated data (Yang et al., 2024; Feng et al., 2025), or reinforcement learning (RL) (Li et al., 2025b; Xue et al., 2025). Long CoT by Large Reasoning Models (LRMs) has been demonstrated as a remarkable capability in complex reasoning (OpenAI, 2024; DeepSeek, 2025). On the other hand, it still heavily relies on natural language as the reasoning medium, and the integration of tool-use into long CoT has been underexplored. This limitation stems primarily from the practice of LRM post-training for reasoning: the SFT employs curated data consisting of pure-text long CoT, and the subsequent RL rewards the models based on the correctness of final answers with little attention to intermediate reasoning behaviors (DeepSeek, 2025; Yang et al., 2025a).

It is infeasible to apply the aforementioned TIR methods, originally devised for short CoT, directly to long CoT, for the following reasons: (1) Prompting with tool-use instructions affects only the answer part instead of the thinking part (long CoT) in LRMs’ responses, hindering the applicability of prompting-based and even pure-RL-based methods (due to the low hit rate in the rollout phase). (2) For supervised finetuning, data annotation incurs non-negligible cost, especially for tool-integrated long CoT. Furthermore, SFT on new data may perturb the original distribution of LRMs, potentially impairing their long CoT capability (seen in § 4).

Recent LLMs (or LRMs) already possess non-trivial tool-use proficiency thanks to pre-training on extensive code- and tool-related data (Gao et al., 2023; Chen et al., 2023; Gou et al., 2024; Wang et al., 2023; Yang et al., 2024). Nevertheless, such capability plays no role during long CoT reasoning. This arises primarily because tool-use behaviors and long CoT reasoning are learned from separate

* Equal contributions.

† Work done while at LLM Department, Tencent.

‡ Corresponding author.

data sources during post-training, resulting in two disjoint output distributions.

Based on all the considerations above, in this paper we explore integrating tool-use behaviors into long CoT without tool-use annotation. We enable spontaneous tool-use in long CoT by fusing the output distributions of tool-use behavior and long CoT, which is a self-bootstrap procedure rather than training on newly curated data. This design offers twofold advantages: first, it eliminates the cost associated with data curation. Second, it preserves and builds on the model’s intrinsic long CoT capabilities by aligning tool-use with, rather than supplanting, its native reasoning patterns.

Specifically, at the core of our approach is a dynamic rollout mechanism in our reinforcement learning process, which elicits and reinforces tool-integrated reasoning trajectories. During the rollout phase, we construct a rollout tree by injecting tool-related hints at the high-entropy yet contextually appropriate positions, effectively producing those long CoT that seamlessly integrate tool invocations. Furthermore, we adopt a tree-based process advantage estimation, which provides a fine-grained signal enabling the model to pinpoint and reinforce specific sub-trajectories where tool invocation genuinely contributes to problem-solving. We validate our approach on three challenging benchmarks. Experimental results demonstrate that our approach significantly outperforms existing methods, and successfully preserves the long CoT ability of LRMs while organically integrating tool-use.

Our contributions are summarized as follows: (i) We introduce DART, a novel RL framework that integrates tool-use into long Chain-of-Thought reasoning without tool-use annotation, eliminating annotation costs and preserving the model’s original reasoning capabilities. (ii) We propose a dynamic tree rollout mechanism to discover and reinforce tool-integrated behavior. (iii) Extensive experiments show that DART achieves superior performance on complex reasoning tasks.

2 Methodology

Our approach comprises two principal sequential modules. First, we introduce rollout tree construction, a dedicated rollout mechanism devised to discover tool-integrated reasoning trajectories (§2.1). Second, we employ tree-based advantage estimation, which leverages the constructed tree to compute the advantage associated with each

sub-trajectories (§2.2) for reinforcement learning. Through these, tool-integrated reasoning behaviors can be reinforced within the policy model.

2.1 Discovery: Rollout Tree with Tool-Integrated Forking

During RL training, for a given question q , we formalize the rollout as multiple reasoning with CoT to generate a group of candidate answers. Since a set of reasoning trajectories naturally forms a tree-like structure where common prefixes branch into diverse continuations, it is feasible to conceptualize reasoning trajectories generation as a tree-structured search process (Yang et al., 2025c; Li et al., 2025c). **Within our tree structure, each node represents a (sub-) reasoning trajectory.**

Based on this, we propose an algorithm for constructing tool-integrated rollout trees. The core idea is to iteratively expand the tree by forking new sub-reasoning trajectories that involve tool-use. As shown in Fig. 1, the rollout tree evolves through successive expansions. Forking operations are performed at positions identified as most promising within the existing reasoning trajectories—determined jointly by entropy and the likelihood of eliciting tool-use behavior—fostering exploration that integrates tool-use into the reasoning.

Initialization. We define the root node of the rollout tree as the question q . Based on q , we first generate M trajectories in parallel using the policy model π_θ , i.e., $\mathbf{y} \sim \pi_\theta(\cdot | q)$. The rollout tree is initialized with the set of generated trajectories, $\mathcal{T}_0 = \{\mathbf{y}\}^M$. The generation stops when the model generates [EOS] token or reaches the preset maximum length. After the initialization, the procedure comprising following sequential steps iteratively runs N times for further expansion with tool-integrated trajectories.

Entropy-based forking position selection. We aim to expand the tree by forking new branches (sub-trajectories) from the existing tree. We propose selecting positions with the highest uncertainty for forking. This is because high-uncertainty positions enable broader exploration, whereas low-uncertainty ones typically inherit and reveal previously established lines of thought (Wang et al., 2025). We use entropy over vocabulary V to quantify the policy model’s uncertainty of current positions. Specifically, from the entire rollout tree, we select the top- K positions $\{t\}$ that immediately follow the conclusion of a clause and have the highest

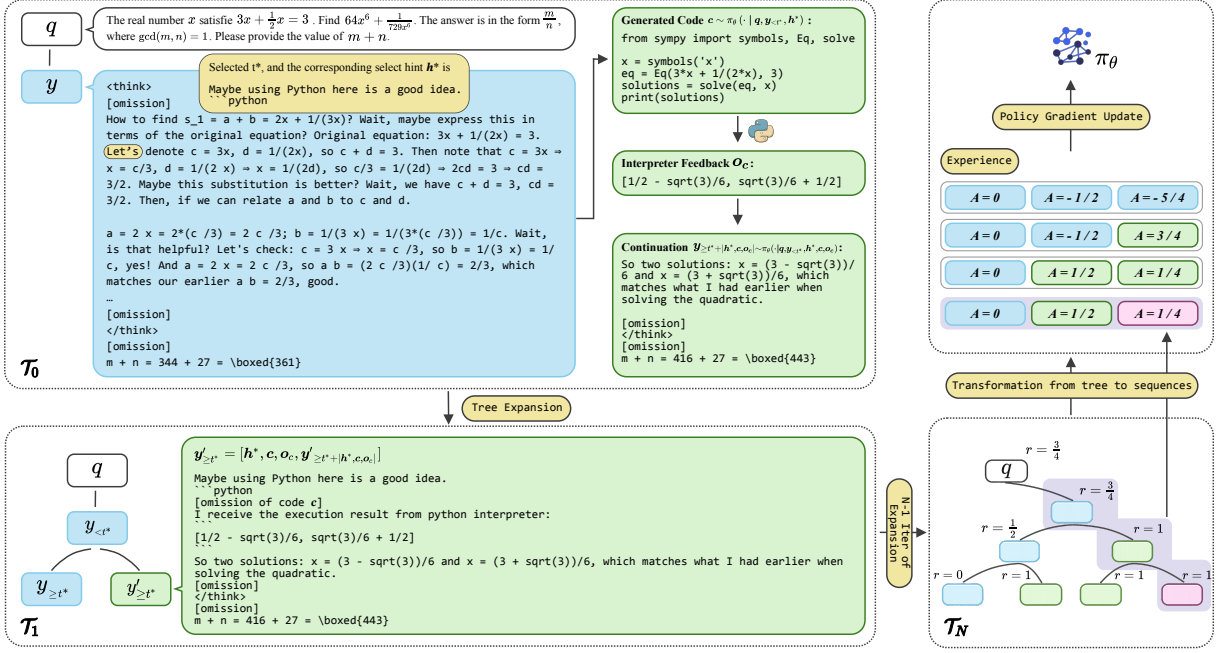


Figure 1: Overview of our approach on a running example with $M = 1$, $N = 3$. (i) $\mathcal{T}_0 \rightarrow \mathcal{T}_1 \rightarrow \mathcal{T}_N$: At each expansion step, a promising (t^*, h^*) is selected, where t^* is an intermediate position within the existing trajectories (e.g., y in \mathcal{T}_0), and h^* is the corresponding hint. Starting from the prefix $[y_{<t^*}, h^*]$, the policy model π_θ generates a Python snippet c ; the Python interpreter then returns the execution feedback o_c . π_θ subsequently completes the partial trajectory $[y_{<t^*}, h^*, c, o_c]$ by generating $y'_{>t^*} = [h^*, c, o_c, y'_{>t^*}]$. The new sub-trajectory $y_{>t^*}$ is plugged into the tree, while the original y is partitioned into $y_{<t^*}$ and $y_{>t^*}$, finally resulting in \mathcal{T}_1 . In the illustrated tool-use example, the code segment employing the sympy package is used to solve a function, replacing the original verbose natural language derivation. (ii) **Upper right:** Each leaf node, together with the nodes along the path from the root to that leaf, constitutes a complete reasoning trajectory. Tokens within a sub-trajectory inherit the advantage value associated with that sub-trajectory.

entropy:

$$\{t\} := \{t \in \mathcal{T}_n \mid \text{rank}_{H(t)} \leq K\},$$

$$H(t) = - \sum_v \pi_\theta(v \mid \mathbf{q}, \mathbf{y}_{<t}) \log \pi_\theta(v \mid \mathbf{q}, \mathbf{y}_{<t}).$$

(1)

We exclude positions near the end (within the final 20%) of a trajectory, since these are more associated with answer finalization than with exploration.

Eliciting tool-integrated trajectories. Once $\{t\}^K$ is determined, we complete the corresponding partial trajectory $y_{<t^*}$ by resuming the generation from a chosen t^* . Specifically, we concatenate a hint at t^* for eliciting tool-integrated sub-trajectories from the policy model, yielding novel reasoning trajectories. To this end, a set of tool-related hints $\{h\}$ is collected, as shown in Tab. 3, where each hint terminates with the opening delimiter of a code block to force the policy model to continue with a Python code snippet. This diverse set of hints is tailored to distinct tool-use scenarios. In order to identify a promising position-hint

pair, we sample¹ a single tuple (t^*, h^*) from the joint space $\{t\} \times \{h\}$, according to the conditional probability $\pi_\theta(h \mid \mathbf{q}, \mathbf{y}_{<t})$.² This strategy enables identifying the position t^* that has substantial potential of being succeeded by the tool-use hint h^* . With the tool-use hint h^* at position t^* , the policy model π_θ continues by generating a Python snippet c until the terminating delimiter of the code block:

$$c \sim \pi_\theta(\cdot \mid \mathbf{q}, \mathbf{y}_{<t^*}, h^*).$$

(2)

Upon completion, an external code sandbox with a Python interpreter receives and executes the code, and the interpreter feedback o_c is then returned. Subsequently, the policy model completes the trajectory until the model generates [EOS] token or

¹To mitigate the tendency of the policy model to repeatedly select a limited subset of specific hints, we employ sampling instead of the argmax operation here.

²Any (t, h) selected in previous iterations are excluded. For pairs that have been probed but not selected, we cache the $\pi_\theta(h \mid \mathbf{q}, \mathbf{y}_{<t})$ to reuse in subsequent iterations for saving computation. For this reason, to ensure efficient maintenance, the pairs are organized in a max-heap based on their entropy.

reaches the preset maximum length, given as

$$\mathbf{y}'_{\geq t^*+|h^*,c,o_c|} \sim \pi_\theta(\cdot | \mathbf{q}, \mathbf{y}_{<t^*}, h^*, c, o_c). \quad (3)$$

$\mathbf{y}'_{\geq t^*} = [h^*, c, o_c, \mathbf{y}'_{\geq t^*+|h^*,c,o_c|}]$ is a newly generated sub-trajectory, which, together with $\mathbf{y}_{<t}$, forms a complete trajectory that seamlessly blends real-time code execution and natural language. Note that, similar to the initialization, the procedure described by Eq. 2 and 3 runs in parallel M times for the selected (t^*, h^*) , yielding the set $\{\mathbf{y}'_{\geq t^*}\}^M$. The rollout tree is then updated as:

$$\mathcal{T}_{n+1} \leftarrow \mathcal{T}_n \cup \{\mathbf{y}'_{\geq t^*}\}^M. \quad (4)$$

We perform tool-integrated forking for N iteration after initializations, finally leading to the rollout tree \mathcal{T}_N with $M \times (N + 1)$ leaf nodes.

2.2 Reinforcement: Tree-based Advantage for Policy Optimization

In addition to enabling the elicitation of tool-integrated sub-trajectories, the rollout tree affords fine-grained process supervision for intermediate steps through Monte Carlo estimation, for encouraging the beneficial tool-use behaviors that facilitate superior reasoning performance. In this section, we describe how the completed rollout tree can be leveraged to derive signals over sub-trajectories for policy optimization.

Credit assignment for sub-trajectories. For each leaf node s_{leaf} in \mathcal{T} , an answer verifier takes as input the complete trajectory represented by the path from root node to s_{leaf} , and outputs a value $r(s_{leaf}) = 1$ if the answer is correct, or 0 otherwise. For each non-leaf node s_{non} , the value is estimated via Monte Carlo methods by aggregating values from its leaf descendants. Specifically, let $Leaves(s_{non})$ denote the set of all leaf nodes that are descendants of node s_{non} . The value $r(s_{non})$ is calculated as the average value over its leaf descendants:

$$r(s_{non}) = \frac{1}{|Leaves(s_{non})|} \sum_{s \in Leaves(s_{non})} r(s). \quad (5)$$

This propagates bottom-up, weighting all intermediate nodes. The value reflects the potential of the node s to lead to the correct answer.

Advantage determination. Our goal is to derive an advantage value for each node in the rollout tree, for reflecting the relative merit of the sub-trajectories represented by that node. This

advantage formulation enables reinforcement of tool-integrated sub-trajectories that demonstrate a comparatively higher likelihood of culminating in a correct answer, while discouraging those with lower potential. Motivated by this, we define the advantage of node s as:

$$A(s) = \underbrace{r(s) - r(s_{root})}_{\text{Global advantage}} + \underbrace{r(s) - r(p_s)}_{\text{Local advantage}}. \quad (6)$$

s_{root} is the root node of the rollout tree and thus $r(s_{root})$ can be interpreted as the global accuracy of the tree. p_s denotes the parent node of s . The global advantage of node s quantifies the extent to which s surpasses the tree-wide average performance. The local advantage of s measures the incremental improvement s brings relative to its parent. Crucially, a larger local advantage implies that s is more likely than its sibling nodes to produce a correct outcome. Specifically speaking, when a new forking sub-trajectory with appropriate tool-use behavior successfully arrives at the correct answer—while its natural language siblings fail—the proposed advantage assignment strategy confers a significant advantage upon this tool-integrated sub-trajectory.

Policy optimization. The set of reasoning trajectories $\{\mathbf{S}\}^{M \times (N+1)}$ can be extracted from the rollout tree \mathcal{T}_N by concatenating the sub-trajectories represented by the nodes along the path from the root to each of the $M \times (N + 1)$ leaf nodes. Tokens inherit the advantage associated with the sub-trajectory to which they belong. This process is exemplified as in the right-hand side of Fig. 1. We then adopt the on-policy Reinforce objective (Williams, 1992) without a KL penalty term³, and the loss function is given as:

$$\mathcal{J}(\theta) = \mathbb{E}_{q, \{\mathbf{S}\}^{M \times (N+1)} \sim \pi_{\theta_{old}}(\cdot | q)} \left[\frac{1}{\sum_{i=1}^{M \times (N+1)} |\mathbf{S}_i|} \sum_{i=1}^{M \times (N+1)} \sum_{j=1}^{|\mathbf{S}_i|} \frac{\pi_\theta(\mathbf{S}_{i,j} | \mathbf{q}, \mathbf{S}_{i,<j})}{\pi_{\theta_{old}}(\mathbf{S}_{i,j} | \mathbf{q}, \mathbf{S}_{i,<j})} A_{i,j} \right]. \quad (7)$$

$A_{i,j}$ denotes the advantage of the j -th token of the i -th trajectory. Following previous works (Li et al., 2025b; Lin and Xu, 2025), We mask out the interpreter feedback output from the loss computation to ensure training stability.

³Other algorithms like PPO (Schulman et al., 2017) and Reinforce++ (Hu et al., 2025) can also be applicable to our setting.

2.3 Inference

Tool-integrated rollout trees are used exclusively during training; at test time, we employ standard single-path generation. No hint guidance is enforced, allowing the model to generate freely while interacting solely with the Python interpreter. Notably, a significant proportion of tool-use triggers preceding code blocks do not match any predefined training hints (see Tabs. 8–10)⁴. This demonstrates that the model has learned the semantics of tool invocation rather than overfitting to the syntactic patterns of the hints.

3 Experiments

3.1 Implementation Details

Backbone models. The experiments are conducted with Qwen3 (Yang et al., 2025b), a family of LLM with native long CoT reasoning ability. We employ two model variants—Qwen3-8B and Qwen3-4B-Thinking-2507—to examine the generalizability of DART across different parameter scales. Compared to the former, the latter has fewer parameters, but has undergone more sufficient post-training and is proficient in complex reasoning.

Training data. Our approach benefits from a more accurate estimation of the performance gap between natural-language and tool-integrated reasoning, and employing challenging problems can more effectively reveal improvements in tool-integrated trajectories. For these reasons, we curate a set of challenging mathematics problems for training. Specifically, we first employ Qwen3-4B-Thinking-2507 to perform inference on all English training examples of the DAPO dataset (Yu et al., 2025), executing five independent runs per example. We filter out instances where the model’s accuracy exceeds 60%, finally retaining a set of 2,795 examples as training data.

Hyperparameters. By default, we set M, N, K to 2, 3, 6, respectively, except in § 5.3. This setting leads to a total of 8 rollout trajectories. We employ a maximum response length of 16,384 tokens and a sampling temperature of 1.0 for both rollout and testing. See more implementation details in App.A.

Environment. Our training is conducted on 8 GPUs using VeRL (Sheng et al., 2025), an RL training library for large language models. We

⁴We observed that during training (after approximately 50 steps), the model begins to invoke tools with novel tool triggers at positions other than where hints were inserted, further indicating independence from the predefined hint set.

implement an asynchronous code sandbox environment with SandboxFusion (Bytedance-Seed, 2025)⁵. This asynchronous setup accelerates RL training by enabling concurrent environment interactions across multiple threads, thereby maintaining high-throughput execution throughout the training.

3.2 Baselines & Evaluation

We compare our approach with a comprehensive set of baseline methods. Those baselines are divided into three categories. (1) **Direct RL** applies vanilla reinforcement learning without a cold-start stage. We use **DAPO** (Yu et al., 2025) to train a baseline for natural language reasoning. For TIR methods, we include **ToRL** (Li et al., 2025b) and **UTIR** (Lin and Xu, 2025). ToRL incentivizes TIR capability in LLMs via RL, by conditioning on the TIR-encouraging prompt and rewarding the answer correctness. UTIR additionally introduces an early-code reward to encourage early code invocation. (2) **SFT-then-RL**, before RL training, performs supervised fine-tuning on expert TIR trajectories to equip models with basic TIR abilities. For the re-implementation of **ReTool** (Feng et al., 2025), we finetune LLMs based on expert trajectories⁶, which consist of short CoT. After the SFT, an RL training process as ToRL is then applied. (3) **SFT on rejection sampling**, such as **START** (Li et al., 2025a), collects tool-integrated-long CoT through rejection sampling – using LLMs to complete long CoT after inserting tool-related hints, and retains those correct ones. The LLMs are then finetuned on the collected long CoT. Note that it is an offline method.

For a fair comparison, we use the same backbone LLMs and training data to re-produce all baselines above⁷. In addition to Qwen3-8B and Qwen3-4B-Thinking-2507, Qwen3-8B-Base and Qwen3-4B-Base are also employed to re-produce ToRL, denoted by ToRL-base, to study the performance gaps between LRMs and LLMs (without thinking) under TIR settings. Details of the implementation of all baselines can be found in App. B.

Our in-domain evaluation benchmarks are AIME24 and AIME25, two challenging math benchmarks. GPQA-Diamond (Rein et al., 2024)

⁵<https://bytedance.github.io/SandboxFusion/>

⁶The expert trajectories are annotated by ReTool, available in <https://huggingface.co/datasets/JoeYing/ReTool-SFT>.

⁷For ReTool, we use the official annotated trajectories for SFT and our training data for RL.

is used as an out-of-domain benchmark for the evaluation on scientific questions, including physics, chemistry, and biology. Pass@1 and Pass@8 are measured for each benchmark, and the average number of code-involutions is also recorded.

4 Main Results

Two general patterns can be observed in Tab. 1. (1) Thinking-based methods, even the ones without tool integration, are superior to tool-integrated but non-thinking counterparts, such as Text-DAPO vs. ToRL-base/ReTool, demonstrating the power of test-time scaling by long CoT (OpenAI, 2024; DeepSeek, 2025). (2) Under thinking mode, tool-integrated approaches outperform natural language-based counterparts, such as UTIR vs. Text-DAPO, corroborating the finding by Lin and Xu (2025) that tool-integrated-reasoning is able to expand LLMs’ capability boundary. These two findings collectively suggest that integrating tool-use into thinking trajectory (long CoTs) improves reasoning quality and leads to superior results.

Among tool-integrated-reasoning methods, our approach shows the best performance. START relies on SFT and functions as the offline variant of our approach. Compared to our approach, it has higher code-use frequencies but suffers significant performance drops. Our manual inspection reveals that it exhibits rigid tool-use behavior that contributes little to or even misleads reasoning (see Tab.11 for an example). The degradation of its reasoning ability can be attributed to SFT on a new task with limited data, which greatly distorts the model’s original distribution. In comparison, DART is trained with RL and this online training method can preserve the LLM’s prior knowledge and capabilities significantly better (Chen et al., 2025b; Shenfeld et al., 2025), enabling reasonable integration of tool-use and natural language reasoning. ToRL-thinking has virtually zero tool-use rates in its long CoT, as prompts fail to affect the underlying reasoning process. UTIR, another RL-based method, is driven by an explicit tool-invocation reward. Our manual inspection finds frequent void tool-use behavior (often with just a few lines of code, like a single print function) by this method, which is mainly due to reward hacking. In our approach, the rollout trajectories are injected with tool-use behavior by dynamically inserting hints and are rewarded solely based on their answer correctness, effectively avoiding reward hacking.

5 Analysis

We will investigate: 1) the extent to which tool-integrated reasoning benefits training; 2) the effectiveness of the tree-derived advantage; 3) the entropy-based selection of forking points; 4) the influence brought by the scale of the rollout trees. Throughout this section, we use Qwen3-4B-Thinking-2507 as the testbed.

5.1 Benefit of Tool-use for Reasoning

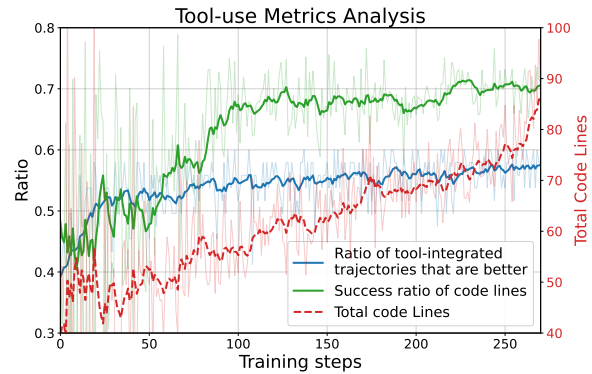


Figure 2: Metrics of tool-use during training.

A key step of our training method is to discover tool-integrated trajectories that can yield correct answers that the original natural language-based chains are unable to achieve. The blue line in Fig. 2 plots the dynamics of the proportion of training samples in which at least one tool-integrated trajectory yields the correct answer, while the natural language trajectories fail to do so. The curve starts at 0.36 at step 0—a non-negligible initial value—indicating that, with our hint-based guidance, the LLM inherently possesses a capacity to employ code tools for reasoning. As training progresses, the curve rises gradually before reaching a stable plateau at approximately 0.6 (for the vast majority of the remaining 0.4, both trajectory types either succeed or fail together).

Furthermore, the curves for the number of code lines (red) and the proportion of code lines that are successfully executed (green) also exhibit an upward trend. This suggests that, rewarded by the additional correctness afforded by tool-integrated trajectories, the model develops a stronger tendency and enhanced capability for tool use. These improvements, in turn, promote the discovery of tool-integrated trajectories. App. C shows analysis on tool-use by our approach during testing on mathematical questions.

Paradigm	Method	CoT Type	AIME24			AIME25			GPQA-D		
			Pass@1	Pass@8	#Invoc.	Pass@1	Pass@8	#Invoc.	Pass@1	Pass@8	#Invoc.
Qwen3-4B-Thinking-2507											
	Original checkpoint	pure-text long	52.22	71.93	0	46.39	65.00	0	65.98	83.32	0
RL	DAPO	pure-text long	69.17	85.67	0	61.53	84.09	0	65.53	82.82	0
	ToRL-Think	pure-text long	71.25	82.14	0	61.25	76.27	0	65.15	74.99	0
	ToRL-Base	tool-integrated short	40.49	50.73	3.12	31.81	42.24	2.99	31.67	43.65	1.01
	UTIR	tool-integrated long	70.56	87.61	10.08	63.89	77.96	8.12	66.08	81.90	2.11
SFT	Retool (only SFT)	tool-integrated short	17.36	34.14	2.12	13.47	23.42	1.89	20.12	42.06	1.12
	START	tool-integrated long	44.31	65.15	13.30	38.47	58.90	12.08	55.82	82.22	4.24
SFT+RL	Retool	tool-integrated short	23.33	35.46	1.85	20.13	36.36	1.35	13.26	20.95	0.98
RL	Ours	tool-integrated long	73.47	91.37	10.68	65.56	85.76	9.51	66.65	83.75	3.81
Qwen3-8B											
	Original checkpoint	pure-text long	64.17	79.28	0	52.08	69.22	0	58.88	79.53	0
RL	DAPO	pure-text long	66.11	81.44	0	56.11	76.98	0	59.41	80.29	0
	ToRL-Think	pure-text long	65.70	82.68	0	59.44	69.91	0	55.24	67.11	0
	ToRL-Base	tool-integrated short	42.64	54.95	3.86	25.83	37.28	2.98	41.56	56.62	1.06
	UTIR	tool-integrated long	69.12	84.59	11.97	61.11	74.96	10.22	59.70	80.71	2.74
SFT	Retool (only SFT)	tool-integrated short	22.78	43.46	2.00	19.67	31.18	2.07	21.23	46.32	1.16
	START	tool-integrated long	35.83	58.22	15.08	30.28	49.76	12.77	35.89	67.61	5.79
SFT+RL	Retool	tool-integrated short	25.41	46.73	1.67	30.00	43.51	1.20	18.27	36.04	0.92
RL	Ours	tool-integrated long	68.89	86.81	12.82	57.64	79.02	10.66	60.03	81.57	4.14

Table 1: Comparison of different methods on AIME24, AIME25, and GPQA-Diamond benchmarks. # *Invoc.* denotes the average number of code-invocations in short/long CoT on the test set. For Retool and ToRL-Base, which rely on short CoT, the value is measured in short CoT; for other methods, it is measured in long CoT. For ToRL-Think, tool-invocations occur only outside the thinking part of the responses.

5.2 Effectiveness of Tree-derived Process Advantage

To validate the benefit of the tree-derived process advantage, we implement a variant that omits process-based advantage. In this variant, analogous to GRPO (Zhihong et al., 2024), the advantage of all tokens within a trajectory is computed as the in-group normalization of the outcome rewards among all trajectories from the rollout tree, and other settings remain the same. As shown in Block *ii* of Tab. 2, while the variant without process-based advantage still surpasses most baselines in §4, it lags behind the default setting (tree-derived advantage) in both answer correctness and tool invocation frequency. Besides, we observe that the number of code lines of the default setting goes up more rapidly and eventually plateaus at a higher level than that of the variant (Fig. 3). Based on these observations, compared with the GRPO-style advantage assignment, where a uniform advantage is assigned to the entire sequence, our tree-derived process advantage is more effective at encouraging the model to invoke tools. This superiority arises because, under the tree-derived advantage scheme, tokens belonging to tool-integrated sub-trajectories that successfully lead to correct answers receive greater advantage values than their

sibling sub-trajectories, providing discriminative signals at the forking positions.

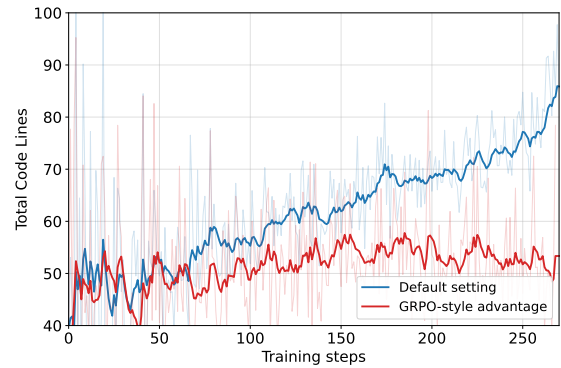


Figure 3: Effect of two advantage assignment schemes on the model’s tool-use behavior.

5.3 Influence of Tree Scale

The scale of the rollout tree plays an important role in the estimation of process advantage. Given that Monte Carlo-style estimation generally benefits from larger sample sizes, we investigate whether scaling up the tree can enhance our approach’s performance by yielding a more accurate estimate of the process advantage. We set M and N to 3 and 4 respectively. For a fair comparison, we keep the experience size for policy optimization consistent by randomly sampling 8 from a total of

	Pass@1	Pass@8	# Invoc.
<i>i. Default setting</i>			
AIME24	73.47	91.37	10.68
AIME25	65.56	85.76	9.51
GPQA-D	66.65	83.75	3.81
<i>ii. GRPO-style advantage</i>			
AIME24	70.56	88.16	6.81
AIME25	61.67	80.57	5.90
GPQA-D	64.34	80.94	3.11
<i>iii. Larger rollout Tree</i>			
AIME24	73.61	91.73	10.64
AIME25	66.81	87.02	10.11
GPQA-D	66.24	83.98	4.17
<i>iv. Random position forking</i>			
AIME24	66.81	84.07	9.83
AIME25	60.28	77.01	9.56
GPQA-D	61.19	82.95	3.22

Table 2: Comparisons across various variants of our approach. *i. Default setting*: the default settings of our approach (§ 3.1); *ii. GRPO-style advantage*: the advantage of tokens within a trajectory is computed as the in-group normalization of the outcome rewards (§ 5.2); *iii. Larger rollout Tree*: A scaled-up rollout tree with $M = 3$ and $N = 4$ is used for advantage estimation (§ 5.3); *iv. Random position forking*: random positions are selected for forking (§ 5.4). For *ii*, *iii*, and *iv*, unless stated otherwise, all other settings remain the same as *i*.

15 trajectories from the enlarged tree⁸. Note that the process advantage is still estimated by using the entire tree. The results in Block *iii* of Tab. 2 show that our approach with the larger rollout tree demonstrates higher pass rates across three benchmarks. However, constructing larger rollout trees will entail additional time overhead.

5.4 Effectiveness of Entropy-based Forking

Motivated by prior findings that high-entropy positions steer the model toward diverse reasoning pathways and facilitate exploration (Wang et al., 2025; Hou et al., 2025), our approach selects the positions with the highest entropy as the forking points to induce tool-integrated trajectories. To assess this selection strategy, we perform an ablation by randomly choosing positions from already generated trajectories, while holding other settings unchanged. The constraint in selection follows the default setting, ignoring tokens within the final 20% of a trajectory. As evidenced by the comparison between Blocks *i* and *iv* in Tab. 2, this variant suffers performance degradation across three benchmarks,

⁸The rollout trees with the default settings described in § 3.1 lead to 8 trajectories in total.

which substantiates the superiority of our strategy of selecting forking positions.

6 Related Works

Tool-Integrated Reasoning (TIR) has emerged as a pivotal paradigm to address the inherent limitations of Large Language Models (LLMs) in precise computation, symbolic manipulation, and complex algorithm execution, by augmenting LLMs with external tools such as code interpreters, search engines, and numerical libraries. Early advancements in TIR primarily relied on Supervised Fine-Tuning (SFT) to imprint tool-use patterns into models: Yang et al. (2024) synthesized data with tool-integrated responses for SFT, enabling Python interpreter integration and achieving superior performance on mathematical benchmarks; ToRA (Gou et al., 2024) and MathCoder (Wang et al., 2023) demonstrated the efficacy of code execution in enhancing mathematical reasoning accuracy, laying the foundation for tool-augmented symbolic computation. These SFT-based approaches, however, are constrained by predefined tool-use trajectories, limiting generalization to unseen tasks or adaptive strategy discovery. START (Li et al., 2025a) further innovated TIR with self-training pipelines: it activates latent tool-use capabilities by synthesizing high-quality training data through diverse hint patterns and rejection sampling. THOR (Chang et al., 2026) utilizes a multi-stage pipeline including data construction, SFT, and Hierarchical RL to integrate tool-use into long CoT. Beyond mathematical reasoning, TIR has been extended to various scenarios: Search-R1 (Jin et al., 2025) and WebSailor (Tongyi et al., 2025) integrated retrieval tools for multi-hop QA and web agent tasks respectively. Chai et al. (2025) orchestrated domain-specific tools for tasks such as automated scientific data analysis.

Recent research has shifted toward Reinforcement Learning (RL)-driven TIR, enabling models to autonomously explore optimal tool-use strategies via outcome feedback. ReTool (Feng et al., 2025) proposed a two-stage framework: cold-start SFT with code-augmented reasoning traces, followed by RL with interleaved real-time code execution. ToRL (Li et al., 2025b) and ZeroTIR (Mai et al., 2025) further scaled RL directly from base models without prior SFT, unlocking emergent behaviors of tool invocation for computational and analytical reasoning. Search-R1 (Jin et al., 2025) and WebSailor (Tongyi et al., 2025) adopted RL to

incentivize use of retrieval tools. Complementary to these, AutoTIR (Wei et al., 2025) designed a hybrid reward mechanism to balance tool integration with language capabilities, supporting autonomous selection among multiple tools (e.g., search engines, code interpreters) and achieving superior generalization across knowledge-intensive, mathematical, and instruction-following tasks. Lin and Xu (2025) provided a formal theoretical proof that TIR expands an LLM’s empirical and feasible support; they also proposed Advantage Shaping Policy Optimization to encourage tool-use behavior.

While prior works such as ToRL (Li et al., 2025b) and ZeroTIR (Mai et al., 2025) successfully incentivize tool-use behavior solely by RL training, their primary focus is based on short CoT. Our work specifically targets long CoT in Large Reasoning Models. THOR (Chang et al., 2026) and START (Li et al., 2025a) rely on complicated pipelines to encourage tool-integrated long CoT; in contrast, our approach achieves this through a streamlined single-stage RL process.

7 Conclusion

In this paper, we propose DART, a reinforcement learning framework that trains Large Reasoning Models to integrate tool-use into long Chain-of-Thought reasoning without relying on any tool-use annotation. By leveraging the proposed rollout tree mechanism, our approach enables the model to discover and reinforce beneficial tool-integrated trajectories. This method successfully preserves the model’s native reasoning capabilities while enhancing tool-use capabilities, achieving superior performance across complex reasoning tasks.

Limitations

In this work, we primarily focus on integrating Python interpreters for mathematical and scientific reasoning tasks (e.g., AIME, GPQA). While this serves as a robust testbed for complex logic, the potential of DART in other domains (such as fact-checking or creative writing) and with other types of tools (such as search engines or knowledge graph APIs) has yet to be fully explored. We believe extending our framework to support a wider range of tools is a promising direction for future research.

Although the proposed rollout-tree construction effectively bridges the gap between long chain-of-thought reasoning and tool-use, its iterative nature inevitably introduces additional latency

during the training phase when compared to reinforcement learning with standard rollout or supervised fine-tuning. In this work, we have implemented asynchronous rollout with the SGLang inference engine, which significantly alleviates the latency. Future work could explore more efficient tree construction strategies or asynchronous mechanisms to further optimize the training throughput.

Ethics Statement

A central component of our proposed framework, DART, involves the autonomous generation and execution of Python code by Large Language Models. While this significantly enhances reasoning capabilities, it introduces potential risks associated with executing arbitrary code. In our experiments, all code execution was strictly confined within a stateless, non-interactive, and isolated sandbox environment (SandboxFusion) to prevent any unauthorized access to the host system or external networks. We strongly emphasize that any real-world deployment of similar tool-integrated reasoning models must implement rigorous sandboxing and security protocols to mitigate the risks of malicious code execution or unintended side effects.

References

- Bytedance-Seed. 2025. [FullStack Bench: Evaluating LLMs as Full Stack Coders](#). *Preprint*, arXiv:2412.00535.
- Jingyi Chai, Shuo Tang, Rui Ye, Yuwen Du, Xinyu Zhu, Mengcheng Zhou, Yanfeng Wang, Weinan E, Yuzhi Zhang, Linfeng Zhang, and Siheng Chen. 2025. [Sci-master: Towards general-purpose scientific ai agents, part i. x-master as foundation: Can we lead on humanity’s last exam?](#) *Preprint*, arXiv:2507.05241.
- Qikai Chang, Zhenrong Zhang, Pengfei Hu, Jun Du, Jiefeng Ma, Yicheng Pan, Jianshu Zhang, Quan Liu, and Jianqing Gao. 2026. [THOR: Tool-integrated hierarchical optimization via RL for mathematical reasoning](#). In *The Fourteenth International Conference on Learning Representations*.
- Chen Chen, Xinlong Hao, Weiwen Liu, Xu Huang, Xingshan Zeng, Shuai Yu, Dexun Li, Yuefeng Huang, Xiangcheng Liu, Wang Xinzhi, and Wu Liu. 2025a. [ACEBench: A comprehensive evaluation of LLM tool usage](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 12970–12998, Suzhou, China. Association for Computational Linguistics.
- Howard Chen, Noam Razin, Karthik Narasimhan, and Danqi Chen. 2025b. [Retaining by doing: The role](#)

- of on-policy data in mitigating forgetting. *Preprint*, arXiv:2510.18874.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2023. [Program of Thoughts Prompting: Disentangling Computation from Reasoning for Numerical Reasoning Tasks](#). *Preprint*, arXiv:2211.12588.
- DeepSeek. 2025. [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#). *Preprint*, arXiv:2501.12948.
- Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. 2025. [ReTool: Reinforcement Learning for Strategic Tool Use in LLMs](#). *arXiv preprint arXiv:2504.11536*.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [PAL: Program-aided Language Models](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10764–10799. PMLR.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujia Yang, Minlie Huang, Nan Duan, and Weizhu Chen. 2024. [ToRA: A Tool-Integrated Reasoning Agent for Mathematical Problem Solving](#). In *The Twelfth International Conference on Learning Representations*.
- Zhenyu* Hou, Ziniu* Hu, Yujiang* Li, Rui* Lu, Jie Tang, and Yuxiao Dong. 2025. [TreeR1: Llm reinforcement learning with on-policy tree search](#). In *ACL*.
- Jian Hu, Jason Klein Liu, and Wei Shen. 2025. [Reinforce++: An efficient rlhf algorithm with robustness to both prompt and reward models](#). *arXiv preprint arXiv:2501.03262*.
- Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Serkan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. [Search-r1: Training LLMs to reason and leverage search engines with reinforcement learning](#). In *Second Conference on Language Modeling*.
- Chengpeng Li, Mingfeng Xue, Zhenru Zhang, Jiayi Yang, Beichen Zhang, Bowen Yu, Binyuan Hui, Junyang Lin, Xiang Wang, and Dayiheng Liu. 2025a. [START: Self-taught reasoner with tools](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 13523–13564, Suzhou, China. Association for Computational Linguistics.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025b. [Torl: Scaling tool-integrated rl](#). *arXiv preprint arXiv:2503.23383*.
- Yizhi Li, Qingshui Gu, Zhofutu Wen, Ziniu Li, Tianshun Xing, Shuyue Guo, Tianyu Zheng, Xin Zhou, Xingwei Qu, Wangchunshu Zhou, Zheng Zhang, Wei Shen, Qian Liu, Chenghua Lin, Jian Yang, Ge Zhang, and Wenhao Huang. 2025c. [Treepo: Bridging the gap of policy optimization and efficacy and inference efficiency with heuristic tree-based modeling](#). *Preprint*, arXiv:2508.17445.
- Heng Lin and Zhongwen Xu. 2025. [Understanding tool-integrated reasoning](#). In *The 5th Workshop on Mathematical Reasoning and AI at NeurIPS 2025*.
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Xinji Mai, Haotian Xu, Zhong-Zhi Li, Weinong Wang, Jian Hu, Yingying Zhang, Wenqiang Zhang, and 1 others. 2025. [Agent rl scaling law: Agent rl with spontaneous code execution for mathematical problem solving](#). *arXiv preprint arXiv:2505.07773*.
- OpenAI. 2024. [Openai o1 system card](#). *Preprint*, arXiv:2412.16720.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2024. [GPQA: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). *Advances in Neural Information Processing Systems*, 36:68539–68551.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *arXiv preprint arXiv:1707.06347*.
- Idan Shenfeld, Jyothish Pari, and Pulkit Agrawal. 2025. [R1’s razor: Why online reinforcement learning forgets less](#). *Preprint*, arXiv:2509.04259.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. [HybridFlow: A Flexible and Efficient RLHF Framework](#). In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys ’25, page 1279–1297. ACM.
- Tongyi, Baixuan Li, Bo Zhang, Dingchu Zhang, Fei Huang, Guangyu Li, Guoxin Chen, Huifeng Yin, Jialong Wu, Jingren Zhou, and 1 others. 2025. [Tongyi deepresearch technical report](#). *arXiv preprint arXiv:2510.24701*.
- Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi Song, Mingjie Zhan, and Hongsheng Li. 2023. [MathCoder: Seamless Code Integration in LLMs for Enhanced Mathematical Reasoning](#). *Preprint*, arXiv:2310.03731.

- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xiong-Hui Chen, Jianxin Yang, Zhenru Zhang, Yuqiong Liu, An Yang, Andrew Zhao, Yang Yue, Shiji Song, Bowen Yu, Gao Huang, and Junyang Lin. 2025. [Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for LLM reasoning](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Yifan Wei, Xiaoyan Yu, Yixuan Weng, Tengfei Pan, Angsheng Li, and Li Du. 2025. AutoTIR: Autonomous tools integrated reasoning via reinforcement learning. *arXiv preprint arXiv:2507.21836*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.
- Zhenghai Xue, Longtao Zheng, Qian Liu, Yingru Li, Xiaosen Zheng, Zejun Ma, and Bo An. 2025. [Simpletir: End-to-end reinforcement learning for multi-turn tool-integrated reasoning](#). *Preprint*, arXiv:2509.02479.
- 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. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025b. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. [Qwen2.5-math technical report: Toward mathematical expert model via self-improvement](#). *Preprint*, arXiv:2409.12122.
- Zhicheng Yang, Zhijiang Guo, Yinya Huang, Xiaodan Liang, Yiwei Wang, and Jing Tang. 2025c. [Treerpo: Tree relative policy optimization](#). *Preprint*, arXiv:2506.05183.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gao-hong Liu, Lingjun Liu, and 1 others. 2025. DAPO: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Tianhua Zhang, Jiaxin Ge, Hongyin Luo, Yung-Sung Chuang, Mingye Gao, Yuan Gong, Yoon Kim, Xixin Wu, Helen Meng, and James Glass. 2024. [Natural language embedded programs for hybrid language symbolic reasoning](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4131–4155, Mexico City, Mexico. Association for Computational Linguistics.
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. 2024a. [Sglang: Efficient execution of structured language model programs](#). *Preprint*, arXiv:2312.07104.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024b. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.
- Shao Zhihong, Wang Peiyi, Zhu Qihao, Xu Runxin, Song Junxiao, Zhang Mingchuan, Li Y.K., Wu Y., and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#).

A Implementation Details of Our Approach

A.1 Prompt Template

The prompt template for our approach is shown below. {Question} will be replaced with specific questions during training and inference.

```
<|im_start|>user
Solve the following problem. You can use \boxed to return
your answer. The last part of your response should be:
\boxed{'The final answer goes here.'}

{Question}
<|im_end|>
<|im_start|>assistant
```

A.2 Hints

We use a collection of hints for code-use guidance compiled by (Li et al., 2025a). Each hint targets a distinct application scenario, such as complex calculations or self-reflection. During rollout, each hint is probed at positions with the highest entropy to identify a promising position–hint pair, whereupon the hint is inserted at the selected position (§2.1).

Complex calculations

I can use Python to perform complex calculations for this problem.\n```python

Self-reflection

I can use Python to check if my approach is correct and refine it, if necessary.\n```python

Check logic

maybe Python can assist in ensuring our logical deductions are sound.\n```python

Alternative method

I can use Python to explore an alternative method for solving this problem.\n```python

General

maybe using python here is a good idea.\n```python

Deeper think

I can think more deeply about this problem through python tools.\n```python

Table 3: The collection of hints, curated by Li et al. (2025a).

A.3 Training Hyperparameters

Tab.4 lists the hyperparameters for RL training. The rollout runs in the SGLang (Zheng et al., 2024a) inference engine under asynchronous settings to speed up the training.

A.4 Interaction with Environment

We employ SandboxFusion to build a sandbox integrated with a Python interpreter. This environment executes Python code generated by the model

Hyperparameters	
Training batch size	32
Optimizer	AdamW (Loshchilov and Hutter, 2017)
Learning rate	1e-7
Warmup step	0
Gradient accumulation step	1
Learning rate scheduler	Linear
KL coefficient	0
Rollout temperature	1
Rollout maximum length	16,384
Rollout number	8 by default, depending on rollout trees
Total epoch	5

Table 4: The hyperparameters used in our RL training.

during the reasoning process and returns the corresponding results. To manage output length, any overlong result is truncated to a maximum of 512 tokens, preserving the final segment.

B Implementation Details of Baselines

B.1 Text-DAPO

Text-DAPO uses the same prompt and training hyperparameters as our approach (A). In addition, it adopts dynamic sampling to enhance training stability: prompt groups whose accuracy is exactly 1 or 0 are filtered out. The overlong penalty factor is set to 1.0.

B.2 ToRL

We re-use the prompt template in the official code, as shown below. All other settings for RL strictly follow those of our approach (App.A.3).

```
<|im_start|>system
A conversation between User and Assistant. The user asks a
question, and the Assistant solves it.
<|im_end|>
<|im_start|>user
Please integrate natural language reasoning with programs
to solve the problem above, and put your final
answer within \boxed{}.

{Question}

<|im_end|>
<|im_start|>assistant
```

B.3 ReTool

ReTool consists of sequential stages of supervised fine-tuning followed by reinforcement learning. For SFT, we finetune the models with the annotated tool-integrated short CoT, provided by Feng et al. (2025). The training set consists of 2,000 samples on math problems. We employ LlamaFactory (Zheng et al., 2024b) to perform SFT, with the training hyperparameters shown in Tab. 5.

In order to retain the competencies learned in SFT, the same prompt template from the SFT data

Hyperparameters	
Training batch size	16
Optimizer	AdamW
Learning rate	1e-4
Warmup ratio	0.1
Gradient accumulation step	3
Learning rate scheduler	Cosine
Maximum length	16,384
Total epoch	3

Table 5: The hyperparameters used in the SFT stage of ReTool and START.

is employed during reinforcement learning training, as illustrated below. All other settings for RL strictly follow those of our approach (App.A.3).

```
<|im_start|>user
Solve the following problem step by step. You now have the
ability to selectively write executable Python code to
enhance your reasoning process. The Python code will be
executed by an external sandbox, and the output (wrapped
in `<interpreter>output_str</interpreter>`) can be
returned to aid your reasoning and help you arrive at the
final answer. The Python code should be complete scripts,
including necessary imports. \nEach code snippet is
wrapped with `<code>\n`python\ncode snippet\n`</code>
>`. \nThe last part of your response should be in the
following format:\n<answer>\n\n\boxed{{'The final answer
goes here.'}}\n</answer>\n\n*user question:*

{Question}

Remember to place the final answer in the last part using
the format: \n<answer>\n\n\boxed{{'The final answer goes
here.'}}
<|im_end|>
<|im_start|>assistant
```

B.4 START

The original work has not been open-sourced yet. To reproduce START (Li et al., 2025a), we follow the default settings reported in the original paper, but uses our own backbone models and training data (§ 3.1) to ensure a fair comparison. For collecting self-training data with tool-integrated long CoT, we use Qwen3-8B and Qwen3-4B-Thinking-2507⁹ to perform inference ten times for each question in the training set. Guided by the hints in Tab.3, the models generate reasoning-aiding code. We retain only those generated trajectories that both yield correct answers and involve tool use, ultimately obtaining 11,360 samples. Self-training on these data is conducted with the hyperparameter set listed in Tab. 5.

B.5 UTIR

We use the official prompt template as shown below.

```
<|im_start|>system
Solve the following problem step by step. You now have the
ability to selectively write executable Python code to
enhance your reasoning process. The Python code will be
executed by an external sandbox, and the output (wrapped
in `output_str`) can be returned to aid your reasoning
and help you arrive at the final answer. The Python code
should be complete scripts, including necessary imports.
Important: The sandbox is stateless and non-interactive;
thus, prior imports, definitions, and state do not persist
between executions and cannot be referenced.\nEach code
snippet is wrapped with `python\ncode snippet\n`. \nThe
last part of your response should be in the following
format:
```

```
<answer>
\n\boxed{{'The final answer goes here.'}}
</answer>
```

```
<|im_end|>
<|im_start|>user
{Question}
```

Remember to place the final answer in the last part using the format:

```
<answer>
\n\boxed{{'The final answer goes here.'}}
</answer>
```

```
<|im_end|>
<|im_start|>assistant
```

For early-code reward, we set δ to -2.0 and k to 0.7 for a conservative setting. We refer readers to Lin and Xu (2025) for further details of this method. All other settings for RL training follow those of our approach (App.A.3).

C Analysis

C.1 Latency Analysis during Training

we present a comparison of the average time spent per training step between Text-DAPO and our approach. During rollout, Text-DAPO samples 8 pure-text trajectories separately and in parallel, and our approach iteratively samples tool-integrated trajectories through iteratively constructing a rollout tree. All experiments are conducted with the default settings in §3.

Method	Average time spent per step (sec)
Text-DAPO	512.2
DART	877.8

Table 6: Latency comparison between Text-DAPO and DART during training.

During training, DART runs slower than Text-DAPO, mainly due to two reasons. First, for DART or other tool-integrated methods, each time a complete code snippet is generated, LLM generation is paused. After code execution, the execution results are concatenated with the existing output, and the LLM resumes generation. This process involves

⁹QwQ-32B-Preview was used in the original paper.

multiple prefill operations. Second, DART’s iterative rollout process (i.e., building rollout tree) introduces overhead compared with the parallel rollout strategy employed by Text-DAPO. Notably, rollout tree construction occurs only during the training stage; therefore, DART has the same level of inference latency as other tool-integrated methods.

C.2 Tool-use Analysis

Tab. 7 lists the Python libraries most frequently used by our models during AIME testing, which are primarily specialized in mathematical and scientific computing.

Package	Description
math	standard mathematical functions for arithmetic
numpy	calculation of multi-dimensional arrays and matrices
complex	mathematical functions specifically for complex numbers
networkx	creation, manipulation, and study of complex networks and graph structures
pulp	generation of linear programming problems and calling of external solvers
gmpy2	fast, multiple-precision arithmetic
pandas	data structures and data analysis tools
divisors	specialized utility for efficiently generating divisors of integers
primefactors	prime factorization of numbers
complex-number	complex number arithmetic
symbolic	symbolic mathematics and algebraic manipulation
sympy	symbolic mathematics and computer algebra
scipy	scientific computing, including optimization, integration, and statistics
shapely	manipulation and analysis of planar geometric objects

Table 7: Python libraries for mathematical and scientific computing used by our models during testing on AIME.

C.3 Case Study

We provide cases answered by DART (Tab. 8, 9, 10) and START (Tab. 11) respectively.

[Question] The 9 members of a baseball team went to an ice-cream parlor after their game. Each player had a single scoop cone of chocolate, vanilla, or strawberry ice cream. At least one player chose each flavor, and the number of players who chose chocolate was greater than the number of players who chose vanilla, which was greater than the number of players who chose strawberry. Let N be the number of different assignments of flavors to players that meet these conditions. Find the remainder when N is divided by 1000.

[Response]

<think>

...

Now, 2016 divided by 1000 is 2 with a remainder of 16. So the remainder is 16.

Wait, but I just have a tiny doubt because sometimes problems like this can have tricks, like if they said at least one player chose each flavor but maybe I misinterpreted that as $s \geq 1$, which I did, and the inequalities are strict, which I did. Let me check an external source in my mind: for example, if it was non-strict, but no, problem says greater than, so strict inequalities. Is there any chance the problem meant ordered by flavor counts but we need to consider permutations? But no—we fixed c as

chocolate, v as vanilla, s as strawberry, so each flavor has a specific count, so we don't have to multiply by permutations for flavor labels because the flavors are distinct and labeled (chocolate, vanilla, strawberry are different flavors). For example, if we were counting for unlabeled flavors, it would be different, but here each flavor is distinct: chocolate is a specific flavor, so assigning 6 to chocolate is a different count than assigning 6 to vanilla, but in our problem, we're just counting assignments where the numbers satisfy $c > v > s$ for chocolate $>$ vanilla $>$ strawberry, so we don't have to permute the counts among flavors—each flavor is fixed in its role (chocolate is the one with most, vanilla medium, strawberry least). Another check:

suppose we had a smaller case to see if our method is correct. But maybe too time-consuming. Alternatively, let's think: is 2016 the right total? Wait, maybe I can use Python to verify the calculations:

```python

```
Calculate multinomial coefficients for each valid (c, v, s) triple
from math import factorial

Case 1: (6, 2, 1)
case1 = factorial(9) // (factorial(6) * factorial(2) * factorial(1))

Case 2: (5, 3, 1)
case2 = factorial(9) // (factorial(5) * factorial(3) * factorial(1))

Case 3: (4, 3, 2)
case3 = factorial(9) // (factorial(4) * factorial(3) * factorial(2))

N = case1 + case2 + case3
remainder = N % 1000

print(f"Case 1: {case1}, Case 2: {case2}, Case 3: {case3}, N: {N}, Remainder: {remainder}")
```
```

I receive the execution result from python interpreter:

...

```
Case 1: 252, Case 2: 504, Case 3: 1260, N: 2016, Remainder: 16
```

```

Yes, Python confirms it. So my initial calculation is correct. The remainder when 2016 is divided by 1000 is 16. I think I was overworried because sometimes problems can be tricky, but in this case, all steps check out—found all valid distributions using integer partitions with strict inequalities, verified each multinomial coefficient with multiple methods, cross-checked with Python (which is a good way to avoid arithmetic errors), and the total is definitely 2016. Dividing by 1000 gives a remainder of 16.

So the answer should be 16.

</think>

...

---

Table 8: A case by DART (Qwen3-4B-Thinking): in the long CoT, a python snippet involving `MATH.FACTORIAL` is used to perform calculation for self-verification.

---

**[Question]** The set of points in 3-dimensional coordinate space that lie in the plane  $x + y + z = 75$  whose coordinates satisfy the inequalities

$$x - yz < y - zx < z - xy$$

forms three disjoint convex regions. Exactly one of those regions has finite area. The area of this finite region can be expressed in the form  $a\sqrt{b}$ , where  $a$  and  $b$  are positive integers and  $b$  is not divisible by the square of any prime. Find  $a + b$ .

---

**[Response]**

<think>

...

Wait, the problem says exactly one region has finite area. In a plane with four lines, the maximum number of bounded regions is 4 (if they form a convex quadrilateral with diagonals), but here we have six intersection points, so it's a convex polygon with six vertices? Wait, no—four lines in general position form a quadrilateral (4 sides, 4 vertices), but here four lines: L1, L2, L3, L4—do they form a convex quadrilateral or more? Let's use the Python shell to check for intersections.

```
```python
from sympy import symbols, Eq, solve

x, y = symbols('x,y')

# Define all four boundary lines (in xy-plane)
L1 = Eq(x, y) # x = y
L2 = Eq(x + y, 76) # x + y = 76
L3 = Eq(x + 2*y, 75) # x + 2y = 75
L4 = Eq(x, -1) # x = -1

# Find all intersections (6 points)
intersections = []
intersections.append(solve((L1, L2), (x, y)))
intersections.append(solve((L1, L3), (x, y)))
intersections.append(solve((L1, L4), (x, y)))
intersections.append(solve((L2, L3), (x, y)))
intersections.append(solve((L2, L4), (x, y)))
intersections.append(solve((L3, L4), (x, y)))

print(intersections)
```
```

I receive the execution result from python interpreter:

...

```
[[{x: 38, y: 38}, {x: 25, y: 25}, {x: -1, y: -1}, {x: 77, y: -1}, {x: -1, y: 77}, {x: -1, y: 38}]]
```

...

This gives the six intersection points in the xy-plane:

1. (38, 38)
2. (25, 25)
3. (-1, -1)
4. (77, -1)
5. (-1, 77)
6. (-1, 38)

...

</think>

...

---

Table 9: A case by DART (Qwen3-4B-Thinking): in the long CoT, a Python snippet utilizes SYMPY to symbolically solve for the intersection points of boundary lines.

---

**[Question]** There are  $8! = 40320$  eight-digit positive integers that use each of the digits 1, 2, 3, 4, 5, 6, 7, 8 exactly once. Let  $N$  be the number of these integers that are divisible by 22. Find the difference between  $N$  and 2025.

---

**[Response]**

<think>

...

I can use Python to define a function that counts the number of even digits in each subset.

```
```python
def count_evens_in_subsets(subsets):
    evens = [2, 4, 6, 8]
    total_evens = 0
    for subset in subsets:
        # Subset is a tuple, convert to list for easy access
        subset_list = list(subset)
        count = sum(1 for digit in subset_list if digit in evens)
        total_evens += count
    return total_evens

total_evens = count_evens_in_subsets(subsets)
print(total_evens)
```
```

I get the error message from python interpreter:

...

```
NameError: name 'subsets' is not defined
```

...

Let me correct that.

```
```python
from itertools import combinations

def find_subsets():
    digits = list(range(1, 9))
    subsets = []
    for subset in combinations(digits, 4):
        if sum(subset) == 18:
            subsets.append(subset)
    return subsets

subsets = find_subsets()
evens = [2, 4, 6, 8]

def count_evens_in_subsets(subsets, evens):
    total_evens = 0
    for subset in subsets:
        subset_list = list(subset)
        count = sum(1 for digit in subset_list if digit in evens)
        total_evens += count
    return total_evens

total_evens = count_evens_in_subsets(subsets, evens)
print(total_evens)
```
```

I receive the execution result from python interpreter:

...

```
16
```

...

Perfect, so the code confirms it: each of the 8 subsets has 2 even digits, so total count of even digits across all subsets is 16. ...

</think>

...

---

Table 10: A case by DART (Qwen3-4B-Thinking): the code initially generated by DART contained a bug. After receiving feedback from the interpreter, it was tried again and then written correctly

---

**[Question]** Every morning Aya goes for a 9-kilometer-long walk and stops at a coffee shop afterwards. When she walks at a constant speed of  $s$  kilometers per hour, the walk takes her 4 hours, including  $t$  minutes spent in the coffee shop. When she walks  $s + 2$  kilometers per hour, the walk takes her 2 hours and 24 minutes, including  $t$  minutes spent in the coffee shop. Suppose Aya walks at  $s + \frac{1}{2}$  kilometers per hour. Find the number of minutes the walk takes her, including the  $t$  minutes spent in the coffee shop.

---

**[Response]**

<think>

...

Now, I can use Python to check if my approach is correct.

```
```python
```

```
# Let's set up the equations in minutes to be thorough
```

```
# First scenario: v = s, total time T1 = 240 minutes
```

```
# Second scenario: v = s + 2, total time T2 = 144 minutes
```

```
# So:
```

```
# 540/s + t = 240 --> equation 1
```

```
# 540/(s + 2) + t = 144 --> equation 2
```

```
# Subtract equation 2 from equation 1: 540/s - 540/(s + 2) = 240 - 144 = 96
```

```
'''
```

This code does not return any output, perhaps because there is no print statement.

...

</think>

...

Table 11: A case by START (Qwen3-4B-Thinking): in the long CoT, the generated Python code did not return any valid information.

Algorithm 1 DART: Rollout Tree Construction and Advantage Estimation

Require: Question q ; Policy model π_θ ; Set of tool-use hints $\{\mathbf{h}\}$; Tree width M (chains per fork); Max expansion depth N ; Top- K entropy positions.

Ensure: Set of trajectories with advantages $\{(S, A(S))\}$

```
1: // Phase 1: Rollout Tree Construction
2: Initialize tree  $\mathcal{T}_0$  with root node  $q$ 
3: Generate  $M$  initial reasoning chains:  $\mathcal{Y}_0 = \{\mathbf{y}_1, \dots, \mathbf{y}_M\} \sim \pi_\theta(\cdot|q)$ 
4:  $\mathcal{T}_0 \leftarrow \mathcal{T}_0 \cup \mathcal{Y}_0$ 
5: for  $n = 1$  to  $N$  do
6:   Step 1.1: Entropy-based Forking Position Selection
7:   Select candidate positions  $\{t\}^K$  with highest entropy
8:   Step 1.2: Hint Selection & Branching
9:   Sample position-hint  $\{t^*, h^*\} \sim \pi_\theta(h|q, \mathbf{y}_{<t^*})$  from  $\{\mathbf{h}\}$ 
10:  Generate code snippet:  $c \sim \pi_\theta(\cdot | q, \mathbf{y}_{<t^*}, h^*)$  (Eq. 2)
11:  Execute code in sandbox:  $o_c \leftarrow \text{PythonInterpreter}(c)$ 
12:  Generate continuation:  $\mathbf{y}' \sim \pi_\theta(\cdot | q, \mathbf{y}_{<t^*}, h^*, c, o_c)$  (Eq. 3)
13:  Form new sub-trajectory:  $\mathbf{y}_{new} = [h^*, c, o_c, \mathbf{y}']$ 
14:  Update tree:  $\mathcal{T}_n \leftarrow \mathcal{T}_{n-1} \cup \{\text{all } \mathbf{y}_{new}\}$ 
15: end for
16: Let  $\mathcal{T}_{final} = \mathcal{T}_N$  be the completed rollout tree
17: // Phase 2: Tree-based Advantage Estimation
18: Step 2.1: Leaf Evaluation
19: for each leaf node  $s_{leaf} \in \mathcal{T}_{final}$  do
20:   Verify answer correctness:  $r(s_{leaf}) \leftarrow \mathbb{I}(\text{answer is correct})$ 
21: end for
22: Step 2.2: Value Propagation (Bottom-up)
23: for each non-leaf node  $s$  in  $\mathcal{T}_{final}$  (from leaves to root) do
24:    $r(s) \leftarrow \frac{1}{|\text{Leaves}(s)|} \sum_{l \in \text{Leaves}(s)} r(l)$  ▷ Avg correctness of descendants
25: end for
26: Step 2.3: Advantage Calculation
27: Extract all nodes/sub-chains  $\{S\}$  from  $\mathcal{T}_{final}$ 
28: for each node  $s \in \{S\}$  do
29:   Global Advantage:  $A_{global} = r(s) - r(s_{root})$ 
30:   Local Advantage:  $A_{local} = r(s) - r(\text{parent}(s))$ 
31:   Total Advantage:  $A(s) \leftarrow A_{global} + A_{local}$  ▷ Reflects both global and incremental merit
32: end for
33: return Trajectories and token-level advantages  $\{(S, A(S))\}$  for optimization
```
