

FlipTS: Failure-to-Success Mathematical Self-Refinement via Trajectory-Aware Sampling in End-to-End Online RL

Yaqian Hao*, Yifan Hao*, Yankai Liu*, Fangning Chao, Zhaojun Cui,
Haiyu Zhang, Huan Bai, Fanyu Meng, Chao Deng, Junlan Feng[†]

JIUTIAN Research, China Mobile
Beijing, China

{haoyaqian, haoyifan, liuyankai, fengjunlan}@it.chinamobile.com

Abstract

Self-Refinement has emerged as a promising paradigm for improving Large Language Model (LLM) performance by iteratively refining its response at inference time. Existing training methods typically treat refinement as a decoupled task and rely on pre-generated responses, which detaches training from the model’s own reasoning trajectories and limits its performance. In this paper, we propose **FlipTS** (Failure-to-Success Learning for self-refinement with Potential-guided Trajectory-aware Sampling), an end-to-end reinforcement learning (RL) framework that optimizes the entire self-refinement loop online. We identify a key bottleneck in this setting: the sparsity of informative *Failure-to-Success* (F2S) refinement behaviors. **FlipTS** addresses this by utilizing a non-stationary Bayesian model to estimate data potential and applying trajectory-aware sampling to enrich training with valuable refinement signals. Experiments on Qwen3-4B-Instruct-2507, Qwen3-1.7B and LLaMA-3.1-8B-Instruct across challenging mathematical benchmarks demonstrate that **FlipTS** consistently outperforms offline and online RL baselines. Notably, it exhibits robust generalization to factuality, scientific reasoning and legal reasoning tasks without domain-specific tuning. Our code is available¹.

1 Introduction

LLMs can improve their outputs at inference time through *Self-Refinement* (see Figure 1). *Self-Refinement* allows the model to iteratively improve its responses and correct errors, effectively exploiting additional computation at inference time (Madaan et al., 2023; Shinn et al., 2023; Khatri et al., 2025). However, its efficacy is constrained by the model’s limited verification and reasoning capabilities (Zhang et al., 2025). This motivates explicit training to strengthen self-correction.

¹<https://github.com/Circlet-infinity/FlipTS>




Figure 1: An example of *Self-Refinement*. The model diagnoses flaws in its initial answer and iteratively corrects them to achieve ground truth.

Recent advancements highlight the potential of outcome-based end-to-end RL (Guo et al., 2025; Wang et al., 2025; Xue et al.). Despite these successes, this paradigm remains underexplored for *Self-Refinement*, where the multi-turn process can be optimized as a unified trajectory. Existing approaches (Zeng et al., 2025; Huang et al., 2024; Hu et al., 2025) typically formulate refinement as an isolated, single-step task, relying predominantly on synthetic trajectories or offline generated responses for training. While straightforward, such methods suffer from a distribution mismatch between training data and actual inference trajectories. As a result, the learned policies may overfit to fixed refinement patterns and generalize poorly to spontaneous self-correction.

Applying end-to-end RL to *Self-Refinement* poses non-trivial challenges, primarily stemming from the sparsity of high-value behaviors. Specifically, *Failure-to-Success* (F2S) trajectories, where the model corrects an initially wrong response, are exceptionally rare under naive sampling. Instead, positive rollouts are predominantly com-

posed of *Success-to-Success* (S2S) instances, where the model generates a correct initial response and merely maintains it. While S2S trajectories provide positive reward signals, they lack explicit error-correction actions and are therefore less effective for learning refinement.

To address these challenges, we propose **FlipTS** (**F**ailure-to-**S**uccess **L**earning for self-refinement with **P**otential-guided **T**rajectory-aware **S**ampling), a novel framework that applies end-to-end RL to *Self-Refinement*. Our approach generates and optimizes complete multi-turn self-refinement trajectories online, enabling on-policy optimization and reducing the distribution shift.

Furthermore, we tackle the sparsity issue with a potential-guided trajectory-aware sampling strategy. We model the occurrence of successful correction for each query as a non-stationary Bernoulli process. Leveraging the Beta-Bernoulli conjugate prior for Bayesian estimation, we employ Thompson Sampling to actively prioritize queries with high refinement potential. **FlipTS** achieves superior performance while adhering to simple outcome-based rewards, effectively circumventing the need for manual reward engineering and mitigating the risk of reward hacking. Our method and subsequent analysis underscore the pivotal role of rollout distribution in RL training: sufficient exposure to high-value trajectories serves as both a prerequisite and a powerful catalyst for the emergence of advanced behaviors.

To evaluate the efficacy of **FlipTS**, we conduct extensive training and evaluate models across challenging benchmarks: AIME25, HMMT (Feb 2025), BRUMO25, and CMIMC25 (Balunović et al., 2025). Experimental results demonstrate that **FlipTS** consistently outperforms competitive baselines, including step-level RL trained on offline data and trajectory-level RL with dynamic filtering. Our work offers three key contributions:

1. We present an end-to-end RL framework for multi-turn *Self-Refinement*, enabling models to evolve intrinsic correction capabilities with complete on-policy trajectories rather than offline curated data.
2. We identify the efficiency bottleneck arising from the sparsity of high-value behaviors and propose a potential-guided Thompson Sampling strategy to address it.
3. We demonstrate that our method enhances per-

formance on challenging mathematical benchmarks, while significantly improving out-of-distribution generalization on factuality, scientific reasoning and legal reasoning evaluations.

2 Related Work

Self-Correction and Self-Refinement in LLMs.

LLMs can boost their reasoning skills simply by critiquing and refining their own outputs, no extra training required. Early approaches primarily focused on inference-time prompting strategies. (Madaan et al., 2023) demonstrated that LLMs can improve their answers through iterative feedback loops without updating model parameters. Moreover, (Wang et al., 2024; Shridhar et al., 2024) further revealed that *Self-Refinement* can serve as an intrinsic mechanism for continuous reasoning enhancement and model self-evolution. However, studies like (Huang et al., 2023; Liu et al., 2025) suggest that LLMs often struggle to self-correct intrinsic reasoning errors without specific training, particularly in mathematical reasoning.

Online vs. Offline RL Learning for Reasoning.

Existing training-based *Self-Refinement* methods are predominantly *offline*: they rely on static datasets synthesized before training and cannot adapt to the model’s evolving error patterns. One line of work uses SFT on teacher-generated refinement trajectories, e.g., (Hu et al., 2025; Zhang et al., 2024) trains on math problems with refined answers produced by a stronger model. Another employs offline RL, optimizing aggregation or selection policies over pre-collected multi-solution traces, as in (Lee et al., 2024; Zeng et al., 2025). Consequently, these methods suffer from distribution shift and tend to imitate fixed, off-policy behaviors rather than correcting their own on-policy errors. In contrast, our fully online RL framework updates exclusively from on-policy trajectories, closing the *Self-Refinement* learning loop and enabling the model to initiate corrections autonomously.

Data Efficiency and Trajectory Selection.

A key challenge in online RL for reasoning is the scarcity of informative feedback signals. Existing methods struggle to adapt to the evolving policy: Kumar et al. (2024) attempts to mitigate sparsity via static, offline pre-screening, which inevitably leads to a growing data-policy mismatch due to distributional shifts as training progresses. Conversely,

while methods like DAPO (Yu et al., 2025) perform online trajectory filtering, they require prohibitively expensive large-scale rollouts. To efficiently navigate this trade-off between stability and cost, we introduce an online Bayesian mechanism, inspired by Qu et al. (2025), that dynamically estimates the informativeness of trajectories. By applying Thompson Sampling, our approach actively prioritizes *Failure-to-Success* trajectories, as these are most likely to yield a high learning signal. This maximizes information gain and sample efficiency without incurring significant inference overhead.

3 Methodology

3.1 End-to-End RL for Self-Refinement

We treat the iterative *Self-Refinement* process as a multi-turn generation task. For a given question q , the process initiates by generating an initial draft solution. In subsequent turns, the model conditions on the original question and the previous solution to produce a refined response. This cycle repeats for a fixed number of turns T , yielding a trajectory of reasoning steps where each iteration aims to improve the correctness of the preceding output.

MDP Formulation. We formalize this process as a Markov Decision Process (MDP). A complete episode is represented as a trajectory τ containing a sequence of states, actions, and immediate rewards:

$$\tau = ((s_1, a_1, r_1), (s_2, a_2, r_2), \dots, (s_T, a_T, r_T)).$$

Here, the state and action spaces are defined dynamically based on the turn t . At the initial turn ($t = 1$), the state is the question itself, $s_1 = q$. For refinement turns ($t > 1$), the state encapsulates the context of the previous attempt, defined as $s_t = (q, a_{t-1})$. The action a_t corresponds to the generated solution at step t . Crucially, r_t denotes the immediate reward received after action a_t , which evaluates the quality of the current step.

RLVR with Trajectory-level GRPO. Following the spirit of Reinforcement Learning with Verifiable Rewards (RLVR), we define the immediate reward r_t based on the objective correctness of the solution. Formally,

$$r_t = \begin{cases} 1, & \text{if } a_t \text{ is correct,} \\ -1, & \text{otherwise.} \end{cases}$$

Specifically, we extract the answer within `\boxed` from a_t and evaluate it against the ground truth

using the `math_verify`² toolkit of version 0.7.0. The trajectory reward $R(\tau)$ is defined as the reward of the final output, i.e., $R(\tau) = r_T$.

We employ Group Relative Policy Optimization (GRPO) (Shao et al., 2024) for RL training. For each question q , a group of G trajectories $\tau_i = \{(s_t^{(i)}, a_t^{(i)}, r_t^{(i)})\}_{t=1}^T$ is sampled from an old policy π_{old} . The advantage for the i -th trajectory is computed by normalizing the rewards within the group:

$$A_i = \frac{R(\tau_i) - \text{mean}(\{R(\tau_j)\}_{j=1}^G)}{\text{std}(\{R(\tau_j)\}_{j=1}^G)}.$$

Then the trajectory-level GRPO objective is formulated as follows, omitting the KL divergence constraint:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{\tau_i\} \sim \pi_{\text{old}}(\cdot|q)} \left[\frac{1}{GT} \sum_{i=1}^G \sum_{t=1}^T \frac{1}{|a_t^{(i)}|} \sum_{l=1}^{|a_t^{(i)}|} \min\{\rho_{t,l}^{(i)} A_i, \text{clip}(\rho_{t,l}^{(i)}, 1 \pm \epsilon) A_i\} \right], \quad (1)$$

where ϵ is the clip ratio and $\rho_{t,l}^{(i)}$ is the importance sampling ratio for l -th token in solution $a_t^{(i)}$,

$$\rho_{t,l}^{(i)} = \frac{\pi_{\theta}(a_{t,l}^{(i)} | s_t^{(i)}, a_{t,<l}^{(i)})}{\pi_{\text{old}}(a_{t,l}^{(i)} | s_t^{(i)}, a_{t,<l}^{(i)})}.$$

3.2 Sparsity of High-value Behaviors

Applying end-to-end RL to complex processes is limited by the intrinsic sparsity of high-value behaviors: desirable capabilities such as error recovery are rare, depriving the optimizer of gradient signals needed to reinforce targeted skills.

This bottleneck is especially acute in iterative *Self-Refinement*, where *Failure-to-Success* (F2S) instances are extremely rare under natural sampling, leaving few positive examples for learning error correction. Two factors compound this sparsity: LLMs’ strong *self-consistency bias* (Xu et al., 2024) favoring initial responses over revision, and *misaligned problem difficulty*, as tasks are often too easy or too hard for refinement to help. As shown in Figure 3, with a horizon of $T = 2$, valid corrections appear in only about 5% of trajectories, severely diluting training signal and hampering policy optimization.

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




Figure 2: Framework of the Proposed Method.




Figure 3: Proportion of four transition types during first 50 training steps of Trajectory GRPO.

To test the generality of this sparsity, we measure the four transition types under natural rollouts for all three models at early (1–50), middle (126–175), and late (251–300) training stages (12,800 trajectories each; Table 1). *F2S* remains rare across models and stages (0.55–8.79%, 4.54% average), while *F2F* and *S2S* dominate. Its trend differs by model: Llama-3.1-8B-Instruct drops sharply after early training (4.04% → 0.55–0.78%), Qwen3-1.7B (Thinking) gradually rises (5.57% → 8.79%), and Qwen3-4B-Instruct stays around 5%. Thus, the sparsity that motivates FlipTS persists throughout training.

3.3 Active Data Selection via Thompson Sampling

To address the sparsity of informative refinement signals, we propose an active data selection strat-

Table 1: Transition statistics (%) under natural rollouts across models and training stages. *F2S*/*F2F*/*S2S*/*S2F* denote transitions between the initial and the refined answer (F = incorrect, S = correct).

Model	Steps	<i>F2S</i>	<i>F2F</i>	<i>S2S</i>	<i>S2F</i>
Llama-3.1-8B Instruct	1–50	4.04	65.30	26.00	4.66
	126–175	0.55	64.77	34.53	0.14
	251–300	0.78	63.96	35.09	0.16
Qwen3-4B Instruct	1–50	4.97	17.79	76.01	1.23
	126–175	4.88	17.84	76.52	0.77
	251–300	5.16	15.69	78.77	0.39
Qwen3-1.7B (Thinking)	1–50	5.57	25.84	68.27	0.32
	126–175	6.16	27.11	66.64	0.09
	251–300	8.79	24.23	66.88	0.10

egy that prioritizes *Failure-to-Success* (*F2S*) trajectories, where an initially incorrect solution is corrected during refinement. For a refinement horizon of $T = 2$, these trajectories correspond to the reward pattern $r_1 = -1$ followed by $r_2 = 1$.

We formulate trajectory selection as a non-stationary multi-armed bandit (NS-MAB) problem, where each question q is treated as an arm and the reward reflects its likelihood of producing *F2S* behavior. Assuming *F2S* transitions follow a Bernoulli process with unknown probability $\xi_{k,q}$ at training step k , we model the uncertainty using a Beta-Bernoulli framework:

$$\xi_{k,q} \sim \text{Beta}(\alpha_{k,q}, \beta_{k,q}),$$

where $\alpha_{k,q}$ and $\beta_{k,q}$ denote the effective pseudo-

counts of observed successes and failures.

Prior Initialization and Posterior Update. To mitigate cold-start issues, we estimate an informative global prior from a warm-up phase of the first W training steps using uniformly sampled questions. For each question q_i , we compute its empirical $F2S$ probability $\hat{\xi}_i$ from collected rollouts. We then fit a Beta prior via the Method of Moments:

$$\alpha_0 = \mu_0 \kappa, \quad \beta_0 = (1 - \mu_0) \kappa, \quad (2)$$

where μ_0 and σ_0^2 are the sample mean and variance of $\{\hat{\xi}_i\}$, and $\kappa = \mu_0(1 - \mu_0)/\sigma_0^2 - 1$. All questions are initialized with $\alpha_{0,q} = \alpha_0$ and $\beta_{0,q} = \beta_0$.

At training step k , let $s_{k,q}$ and $f_{k,q}$ denote the counts of trajectories with and without $F2S$ transitions for question q . The posterior is updated as:

$$\begin{aligned} \alpha_{k+1,q} &= s_{k,q} + \lambda \alpha_{k,q}, \\ \beta_{k+1,q} &= f_{k,q} + \lambda \beta_{k,q}, \end{aligned} \quad (3)$$

where $\lambda \in (0, 1]$ is an exponential decay factor for handling non-stationarity. At training step k , we construct the training batch $\mathcal{B}_k$ by leveraging the posterior to identify questions whose estimated $F2S$ probabilities are closest to a target ξ_{target} . Using Thompson Sampling, we draw

$$\tilde{\xi}_{k,q} \sim \text{Beta}(\alpha_{k,q}, \beta_{k,q})$$

for each question q , and select the top- B questions minimizing the absolute deviation:

$$\mathcal{B}_k = \text{Top-}B_{q \in \mathcal{D}} \left(-|\tilde{\xi}_{k,q} - \xi_{\text{target}}| \right). \quad (4)$$

3.4 Scalable Exploration via a Dynamic Candidate Pool

Applying Thompson Sampling directly to a massive dataset $\mathcal{D}$ with $|\mathcal{D}| \gg B$ is inefficient, as large candidate spaces amplify statistical outliers and lead to stale posterior estimates due to sparse observations. To address this, we restrict Thompson Sampling to a dynamically maintained candidate pool $\mathcal{C} \subset \mathcal{D}$ with $B < M \ll |\mathcal{D}|$. After warm-up, the pool is initialized by uniformly sampling M questions from $\mathcal{D}$. During training, once the fraction of observed questions exceeds a threshold η , we refresh the pool by removing a fraction γ of questions whose posterior means deviate most from the target ξ_{target} , and replace them with newly sampled questions. This strategy improves exploration efficiency while mitigating posterior staleness and pool saturation. The complete workflow is provided in Appendix J.

4 Experiments

4.1 Experimental Settings

We implement **FlipTS** and all baselines using the ver1-agent framework (Feng et al., 2025). Experiments are conducted on **Qwen3-4B-Instruct-2507**, **Qwen3-1.7B** (Team, 2025), and **Llama3.1-8B-Instruct** (Dubey et al., 2024). All models are trained on a mixture of **DeepMath** (He et al., 2025) and the level-1/2 subset of **AM** (Ji et al., 2025).

For challenging mathematical reasoning, we evaluate Qwen models on **AIME25**, **HMMT (Feb 2025)**, **BRUMO25**, and **CMIMC25** (Balunović et al., 2025), reporting mean accuracy over 128 independently sampled trajectories, denoted as $\text{Acc}_{\text{avg}}@128$. For Llama models, we evaluate on **MATH500** (Lightman et al., 2024) with $\text{Acc}_{\text{avg}}@8$. To assess out-of-distribution generalization, we further evaluate on **TruthfulQA** (Lin et al., 2022), **SuperGPQA** (Du et al., 2025), and the **Gaokao-Physics** and **Law-Analytics** subsets from **AGIEval** (Zhong et al., 2023). Unless otherwise specified, OOD results are reported using $\text{Acc}_{\text{avg}}@128$. Detailed benchmark statistics are provided in Appendix C.

Training follows a strict on-policy setting without KL regularization. We use learning rate $1e-6$, batch size 16, group size 16, and trajectory length $T = 2$. For **FlipTS**, we set $\xi_{\text{target}} = 0.5$, pool size $M = 2000$, refresh threshold $\eta = 0.3$, and culling ratio $\gamma = 0.25$. These hyperparameters are chosen without exhaustive tuning; sensitivity analyses are provided in Appendix G. Additional implementation details, including model-specific decoding budgets and temperatures, are summarized in Appendix D.

4.2 Main Results and Comparative Analysis

We compare against three competitive baselines: (1) **Step-level GRPO**, which applies single-turn GRPO to refinement tasks with offline-generated initial solutions; (2) **Trajectory GRPO**, the end-to-end RL algorithm in Section 3.1 without active data selection or the dynamic candidate pool; and (3) **Dynamic-filter GRPO**, a Trajectory GRPO variant that dynamically filters trajectories with zero advantage. Unlike Dynamic-filter GRPO, which requires extra online rollouts to maintain the same effective batch size after filtering, FlipTS updates the candidate pool without additional rollout cost and incurs negligible memory overhead. More baseline details are given in Appendix E.

Table 2: $\text{Acc}_{\text{avg}}@128$ (%) of Qwen3-4B-Instruct under 10-turn *Self-Refinement* across four benchmarks. ‘Rollouts’ indicates the relative number of generated trajectories required to obtain one update during training.

Method	AIME25	HMMT (Feb 2025)	BRUMO25	CMIMC25	Average $\uparrow$	Rollouts $\downarrow$
w/o training	52.50	36.88	63.59	36.29	47.31	N/A
Step-level GRPO	55.45	37.86	63.10	37.75	48.29	1 $\times$
Trajectory GRPO	58.41	40.13	65.42	39.34	50.82	1 $\times$
Dynamic-filter GRPO	59.74	42.71	66.54	40.04	52.36	4 $\times$
FlipTS	62.76	42.73	67.11	40.08	53.17	1 $\times$

Table 2 summarizes the performance of different training strategies applied to Qwen3-4B-Instruct-2507 on four challenging mathematical reasoning benchmarks using 10-turn self-refinement. Step-level GRPO yields marginal gains, indicating limitations of offline-synthesized solutions. Trajectory GRPO improves via multi-step reasoning but is limited by uniform sampling. Dynamic-filter GRPO boosts performance via online filtering but needs 4 $\times$ trajectories per valid update. FlipTS attains the highest accuracy with Trajectory GRPO’s rollout budget (1 $\times$).




Figure 4: $\text{Acc-Avg}@128$ of FlipTS on AIME25, HMMT25, BRUMO25, and CMIMC25 across refinement turns.

Figure 4 shows the evaluation performance of **FlipTS** under iterative self-refinement. Although trained with a short horizon ($T = 2$), the model consistently improves over up to 10 refinement turns, suggesting that FlipTS effectively learns a transferable self-correction capability.

Compared with both offline baselines (e.g., Trajectory GRPO) and online filtering methods (e.g., Dynamic-filter GRPO), FlipTS achieves stronger refinement performance without additional rollout cost. These results demonstrate that targeted trajectory refinement is more effective and efficient than

uniform sampling or difficulty-based filtering.

4.3 Statistical Significance Analysis

To assess the statistical significance of the main results in Table 2, we run 100,000 problem-level paired-bootstrap replicates within each benchmark and equally average the four benchmark effects. Table 3 reports the differences and the 95% confidence intervals (CIs) of FlipTS against the two strongest baselines.

Table 3: Paired-bootstrap significance analysis (percentage points; 95% CI, 100,000 replicates).

Comparison	Diff	95% CI
FlipTS – Dynamic-filter GRPO	+0.81	[−1.05, 2.66]
FlipTS – Trajectory GRPO	+2.35	[0.54, 4.22]

The interval versus Trajectory GRPO excludes zero, indicating a significant improvement under the same rollout budget. The interval versus Dynamic-filter GRPO includes zero, so we treat their accuracies as statistically comparable; notably, FlipTS attains this accuracy without the approximately 4 $\times$ rollout cost per valid update incurred by Dynamic-filter GRPO.

4.4 Performance of the Thinking Model

Table 4 reports the refinement performance of Qwen3-1.7B (thinking mode) before and after iterative refinement. Since Thinking models typically generate substantially longer reasoning trajectories, we focus on the refinement gain brought by a single refinement stage and report the results after one refinement.

The untrained model exhibits a moderate intrinsic self-revision ability, improving by 1.87 points on average, from 30.12% to 31.99%. Trajectory GRPO increases the average gain to 3.01 points, corresponding to a 61% larger refinement gain than the untrained baseline. This confirms that

Table 4: $\text{Acc}_{\text{avg}}@128$ (%) of Qwen3-1.7B (thinking mode) before and after *Self-Refinement* across four benchmarks. “Init” denotes the initial response, while “Refine” denotes the response after 2-turn refinement. Δ_{avg} denotes the average refinement gain across four benchmarks.

Method	AIME25		HMMT (Feb 2025)		BRUMO25		CMIMC25		Refinement Gain Δ_{avg}
	Init	Refine	Init	Refine	Init	Refine	Init	Refine	
w/o training	35.34	36.09	20.73	21.74	46.09	49.64	18.32	20.47	+1.87
Trajectory GRPO	35.96	37.68	23.36	25.18	45.21	50.99	20.25	22.97	+3.01
FlipTS	34.90	38.33	23.88	27.19	44.97	53.28	19.12	23.07	+4.75

trajectory-level RL can strengthen refinement behavior. However, FlipTS further enlarges the gain to 4.75 points, which is 2.55 times the untrained gain and 1.58 times the Trajectory GRPO gain. By emphasizing failure-to-success trajectories, FlipTS provides more informative learning signals for correcting initially imperfect reasoning paths, resulting in stronger self-refinement gains.

4.5 Can FlipTS Rescue Models without Intrinsic Self-Refinement Ability?




Figure 5: Iterative self-refinement performance of LLaMA3.1-8B-Instruct across refinement turns. The vanilla model exhibits refinement collapse, while Trajectory GRPO reverses the degradation trend through trajectory-level reinforcement learning. FlipTS further improves both initialization and long-horizon refinement stability.

We further examine whether FlipTS can improve models with weak intrinsic self-refinement capability. As shown in Figure 5, vanilla LLaMA3.1-8B-Instruct exhibits refinement collapse: accuracy drops from 41.95% at initialization to 37.68% at Turn 10. This suggests that self-refinement may accumulate reasoning errors when the model lacks an explicit self-correction mechanism.

Trajectory GRPO reverses this trend, improving Turn-10 accuracy to 52.53%, which indicates that trajectory-level reinforcement learning can induce basic self-correction behavior. FlipTS fur-

ther strengthens this process through trajectory-aware sampling, reaching the best refinement performance of 55.48% and improving over vanilla refinement by 17.80 percentage points. These results show that prioritizing failure-to-success trajectories can transform unreliable refinement into stable improvement.

4.6 Mechanism Analysis of FlipTS

To better understand why FlipTS works, we analyze both its architectural components and its sampling objective. We first ablate the dynamic candidate pool and Thompson Sampling, and then compare the proposed *Failure-to-Success* (F2S) target with a *Final Success* (Final-S) target. The results are summarized in Table 5.

4.6.1 Ablation of Architectural Components

We consider three variants: *w/o Dynamic Pool*, which applies Thompson Sampling over the full training set; *w/o Thompson Sampling*, which keeps the dynamic pool but replaces F2S-targeted Thompson Sampling with random sampling; and *w/o Both*, which reduces the method to standard Trajectory GRPO.

The results show that both components are crucial to FlipTS. Removing either the dynamic candidate pool or Thompson Sampling substantially reduces the average score from 53.17 to 50.33 and 50.68, respectively, making both variants comparable to Trajectory GRPO. This suggests that neither component alone provides a consistent advantage over standard trajectory-level RL; the gain of FlipTS comes from their joint effect.

Without the dynamic candidate pool, Thompson Sampling is applied over the full training set, where sparse observations, stale posterior estimates, and noisy high-variance samples can make selection ineffective, as discussed in Section 3.4 and Appendix B. Conversely, without Thompson Sampling, the dynamic pool merely restricts training to a smaller candidate space but does not actively

Table 5: Mechanism analysis of FlipTS. We study both component ablations and the effect of sampling targets.

Variant	AIME25	HMMT (Feb 2025)	BRUMO25	CMIMC25	Average ↑
<i>Component ablation</i>					
FlipTS (F2S Target)	62.76	42.73	67.11	40.08	53.17
w/o Dynamic Pool	57.68	38.52	65.78	39.34	50.33
w/o Thompson Sampling	59.95	39.11	64.22	39.45	50.68
w/o Both (Trajectory GRPO)	58.41	40.13	65.42	39.34	50.82
<i>Sampling target ablation</i>					
Final-S Target	58.88	40.91	66.30	40.49	51.54

Table 6: OOD performance of three models under 10-turn self-refinement on four benchmarks.

Model	Method	TruthfulQA		Gaokao-Physics		SuperGPQA		Law-Analytics		Average	
		Init	Refine	Init	Refine	Init	Refine	Init	Refine	Init	Refine
Qwen3-4B-Instruct	w/o training	81.73	84.33	90.04	91.88	36.13	45.30	81.55	86.59	72.36	77.03
	Trajectory GRPO	80.80	84.99	93.04	93.73	44.07	47.00	84.09	88.52	75.50	78.56
	FlipTS	84.05	86.08	93.84	94.65	44.19	47.46	85.47	89.56	76.89	79.44
LLaMA-3.1-8B	w/o training	50.25	59.73	37.57	38.86	21.34	21.54	13.97	18.91	30.78	34.76
	FlipTS	62.56	62.55	43.61	43.39	25.19	26.24	17.12	17.93	37.12	37.53
Qwen3-1.7B-Thinking	w/o training	67.88	69.53	77.87	78.09	31.67	32.40	66.14	68.75	60.89	62.19
	FlipTS	67.23	69.19	77.73	79.24	31.40	32.95	68.80	72.34	61.29	63.43

identify questions with suitable *F2S* likelihoods. Thus, Thompson Sampling discovers informative self-correction examples, while the dynamic pool makes this discovery reliable and scalable under large-scale non-stationary RL training.

4.6.2 Which Trajectories Drive Effective Self-Refinement?




Figure 6: Training dynamics of *F2S* and *Final-S* sampling. (a) *F2S* Rate: proportion of *F2S* trajectories per batch. (b) Effective *F2S* Rate: proportion of *F2S* trajectories with non-zero advantage. (c) Final Success Rate: proportion of trajectories achieving the correct final answer. (d) Non-zero Advantage Rate: proportion of trajectories with non-zero advantage.

and MoPPS (Qu et al., 2025) typically select questions based on the *Final Success* rate, aiming to filter out trivial samples with zero advantage while retaining questions of moderate difficulty. To examine whether final correctness or explicit correction behavior is more useful for refinement learning, we compare FlipTS with a *Final-S*-targeted variant. As shown in Table 5, replacing the *F2S* target with *Final-S* reduces the average accuracy from 53.17 to 51.54, indicating that successful self-correction provides a more informative signal than final correctness alone.

Figure 6 further explains this gap. The *F2S*-guided strategy maintains a substantially higher proportion of *F2S* trajectories, both in raw counts and among non-zero-advantage samples, increasing the density of explicit correction signals with negligible additional overhead. Although the *Final-S* strategy tends to select harder prompts with lower final accuracy, both strategies yield comparable non-zero-advantage rates. This suggests that simply increasing task difficulty or filtering zero-advantage samples is insufficient; effective self-refinement training benefits more from enriching the training distribution with successful correction behaviors.

Prior methods such as DAPO (Yu et al., 2025)

Table 7: Disentangling initial-answer quality and refinement ability on Qwen3-4B-Instruct-2507 ($\text{Acc}_{\text{avg}}@128, \%$). (a) Turn-1 accuracy with each model’s own initial answer; (b) Turn-10 accuracy refining the *same* initial answers from the untrained model. Δ_{avg} is the average gain over the untrained initial answers.

Method	AIME25	HMMT (Feb 2025)	BRUMO25	CMIMC25	Average	Δ_{avg}
<i>(a) Turn-1 accuracy (own initial answers)</i>						
w/o training (Turn-1)	47.24	29.79	57.42	29.84	41.07	–
Trajectory GRPO (Turn-1)	51.59	31.98	59.90	32.99	44.11	+3.04
FlipTS (Turn-1)	54.04	32.58	61.28	33.32	45.31	+4.24
<i>(b) Turn-10 refinement from fixed initial answers</i>						
Fixed Init (untrained, Turn-1)	47.24	29.79	57.42	29.84	41.07	–
w/o training (Turn-10)	52.50	36.88	63.59	36.29	47.31	+6.24
Trajectory GRPO (Turn-10)	55.23	38.83	64.69	37.71	49.12	+8.05
FlipTS (Turn-10)	60.34	41.02	66.22	39.98	51.89	+10.82

4.7 Out-of-Distribution Performance

We assess whether refinement behavior learned from mathematical reasoning transfers to out-of-domain tasks (Table 6). For Qwen3-4B-Instruct and LLaMA-3.1-8B, FlipTS improves both initialization and 10-turn refinement on all benchmarks, with the largest gains on LLaMA-3.1-8B (e.g., +12.31 points on TruthfulQA initialization). For Qwen3-1.7B-Thinking, gains center on refinement: refined accuracy improves on three of four benchmarks and both averaged scores rise (60.89/62.19 $\rightarrow$ 61.29/63.43), while initialization stays comparable to or slightly below the untrained model.

For the two Instruct-tuned models, FlipTS improves TruthfulQA (Lin et al., 2022) and yields larger gains on STEM benchmarks (SuperGPQA (Du et al., 2025), Gaokao-Physics (Zhong et al., 2023)), suggesting better suppression of hallucinated or self-inconsistent reasoning and stronger multi-step inference; it also improves LawAnalytics (Zhong et al., 2023) for all models. Overall, FlipTS transfers broadly on the two Instruct-tuned models and mainly via refinement on Qwen3-1.7B-Thinking, indicating that learned refinement generalizes beyond mathematics.

4.8 Disentangling Initial Accuracy and Refinement Ability

The end-to-end protocol in Table 2 shows two capabilities: the quality of the initial answer and the ability to correct it. Moreover, the trajectory-level advantage reinforces all tokens of a successful trajectory, including an incorrect initial answer, which may gradually degrade first-turn accuracy. We disentangle the two on Qwen3-4B-Instruct-2507 in Table 7.

Initial answers. Training improves rather than degrades Turn-1 accuracy: the average rises from 41.07 to 44.11 under Trajectory GRPO and to 45.31 under FlipTS, with gains on all four benchmarks.

Conditional refinement. Panel (b) fixes the initial answers to those generated by the untrained model, so differences reflect pure refinement ability. FlipTS gains the most ($\Delta_{\text{avg}} = +10.82$, up to +13.10 on AIME25), ahead of Trajectory GRPO (+8.05) and the untrained model (+6.24), and leads on every benchmark.

Hence FlipTS improves both sources of end-to-end performance: more accurate initial answers, and, given identical initial answers, the strongest refinement.

5 Conclusion

This work introduces **FlipTS**, an end-to-end online RL framework that enhances iterative *Self-Refinement* in LLMs and addresses key limitations of offline methods. FlipTS combines Bayesian trajectory-value estimation with targeted Thompson Sampling to prioritize promising trajectories, effectively mitigating the sparsity of informative refinement behaviors. FlipTS reverses **refinement collapse**, where untrained models degrade over multiple self-refinement steps due to the lack of an internal verification mechanism. Through principled training and targeted trajectory sampling, FlipTS stabilizes iterative refinement, enabling the model to reliably correct its own errors and progressively improve reasoning. Results demonstrate that FlipTS achieves strong performance on challenging mathematical reasoning benchmarks and generalizes effectively to OOD tasks, highlighting its robust and transferable reasoning capabilities.

6 Limitations

Our analysis highlights the importance of rollout distribution in RL training, particularly in terms of the model’s exposure to high-value trajectories. In the present work, training is restricted to a short horizon of $T = 2$, consisting of an initial response followed by a single refinement step, primarily due to limitations in available computational resources. This setting allows us to observe basic self-refinement behavior. However, how extending RL training to longer, multi-turn rollouts would influence trajectory coverage or reasoning dynamics remains a worthwhile direction for investigation. Systematic investigation of such settings may help clarify under what conditions rollout length affects the emergence and utilization of high-value trajectories.

Our current work focuses on enhancing the efficiency of sequential, iterative *Self-Refinement*. Moving forward, exploring parallel self-evolving strategies is crucial for scalability. Future research will investigate integrating our trajectory-aware sampling principles with parallel refinement techniques. In particular, we aim to develop policies that can concurrently generate and evaluate multiple candidate solutions or reasoning chains, while supporting efficient online RL training.

Furthermore, we address the bottleneck of sparse high-value trajectories through active online sample selection. While this strategy effectively prioritizes informative data, the optimization process still relies on relatively coarse trajectory-level rewards. In future work, we will integrate our sampling framework with finer-grained step-level reward shaping and dynamic advantage optimization to jointly improve online sample selection and credit assignment, thereby alleviating reward sparsity and enabling more robust self-correction in complex reasoning scenarios.

References

- Mislav Balunović, Jasper Dekoninck, Ivo Petrov, Nikola Jovanović, and Martin Vechev. 2025. Matharena: Evaluating llms on uncontaminated math competitions. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmark*.
- Xinrun Du, Yifan Yao, Kaijing Ma, Bingli Wang, Tianyu Zheng, King Zhu, Minghao Liu, Yiming Liang, Xiaolong Jin, Zhenlin Wei, and 1 others. 2025. Supergpqa: Scaling llm evaluation across 285 graduate disciplines. *arXiv preprint arXiv:2502.14739*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv-2407.
- Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. Group-in-group policy optimization for llm agent training. *arXiv preprint arXiv:2505.10978*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, and 1 others. 2025. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv preprint arXiv:2504.11456*.
- Ting-Yao Hu, Hema Swetha Koppula, Hadi Pouransari, Cem Koc, Oncel Tuzel, and Raviteja Vemulapalli. 2025. Learning from self critique and refinement for faithful llm summarization. *arXiv preprint arXiv:2512.05387*.
- Audrey Huang, Adam Block, Dylan J Foster, Dhruv Rohatgi, Cyril Zhang, Max Simchowitz, Jordan T Ash, and Akshay Krishnamurthy. 2024. Self-improvement in language models: The sharpening mechanism. *arXiv preprint arXiv:2412.01951*.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xanyuan Peng, and Denny. 2023. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*.
- Yunjie Ji, Shuo Zhao, Xiaohao Tian, and 1 others. 2025. How difficulty-aware staged reinforcement learning enhances llms’ reasoning capabilities: A preliminary experimental study. *arXiv preprint arXiv:2504.00829*.
- Devvrit Khatri, Lovish Madaan, Rishabh Tiwari, Rachit Bansal, Sai Surya Duvvuri, Manzil Zaheer, Inderjit S Dhillon, David Brandfonbrener, and Rishabh Agarwal. 2025. The art of scaling reinforcement learning compute for llms. *arXiv preprint arXiv:2510.13786*.
- Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, and 1 others. 2024. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*.
- Kyungjae Lee, Dasol Hwang, Sunghyun Park, Youngsoo Jang, and Moontae Lee. 2024. Reinforcement learning from reflective feedback (rlrf): Aligning and improving llms via fine-grained self-reflection. *arXiv preprint arXiv:2403.14238*.

- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. [Let’s verify step by step](#). In *The Twelfth International Conference on Learning Representations*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 3214–3252.
- Liping Liu, Chunhong Zhang, Likang Wu, Chuang Zhao, Zheng Hu, Ming He, and Jianping Fan. 2025. Instruct-of-reflection: Enhancing large language models iterative reflection capabilities via dynamic-meta instruction. *arXiv preprint arXiv:2503.00902*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, and 1 others. 2023. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594.
- Yun Qu, Qi Wang, Yixiu Mao, Vincent Tao Hu, Björn Ommer, and Xiangyang Ji. 2025. Can prompt difficulty be online predicted for accelerating rl finetuning of reasoning models? *arXiv preprint arXiv:2507.04632*.
- 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*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652.
- Kumar Shridhar, Koustuv Sinha, Andrew Cohen, Tianlu Wang, Ping Yu, Ramakanth Pasunuru, Mrinmaya Sachan, Jason Weston, and Asli Celikyilmaz. 2024. The art of llm refinement: Ask, refine, and trust. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5872–5883.
- Qwen Team. 2025. [Qwen3 technical report](#). Preprint, arXiv:2505.09388.
- Yifei Wang, Yuyang Wu, Zeming Wei, Stefanie Jegelka, and Yisen Wang. 2024. A theoretical understanding of self-correction through in-context alignment. *Advances in Neural Information Processing Systems*, 37:89869–89912.
- Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Xing Jin, Kefan Yu, Minh Nhat Nguyen, Licheng Liu, and 1 others. 2025. Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*.
- Wenda Xu, Guanglei Zhu, Xuandong Zhao, Liangming Pan, Lei Li, and William Wang. 2024. [Pride and prejudice: LLM amplifies self-bias in self-refinement](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15474–15492, Bangkok, Thailand. Association for Computational Linguistics.
- Zhengkai Xue, Longtao Zheng, Qian Liu, Yingru Li, Xiaosen Zheng, Zejun MA, and Bo An. Simpletir: End-to-end reinforcement learning for multi-turn tool-integrated reasoning. In *NeurIPS 2025 Fourth Workshop on Deep Learning for Code*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, and 1 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Yongcheng Zeng, Xinyu Cui, Xuanfa Jin, Guoqing Liu, Zexu Sun, Dong Li, Ning Yang, Jianye Hao, Haifeng Zhang, and Jun Wang. 2025. Evolving llms’ self-refinement capability via synergistic training-inference optimization. *arXiv preprint arXiv:2502.05605*.
- Qingjie Zhang, Di Wang, Haoting Qian, Yiming Li, Tianwei Zhang, Minlie Huang, Ke Xu, Hewu Li, Liu Yan, and Han Qiu. 2025. [Understanding the dark side of LLMs’ intrinsic self-correction](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 27066–27101, Vienna, Austria. Association for Computational Linguistics.
- Xiaoying Zhang, Baolin Peng, Ye Tian, Jingyan Zhou, Lifeng Jin, Linfeng Song, Haitao Mi, and Helen Meng. 2024. Self-alignment for factuality: Mitigating hallucinations in llms via self-evaluation. *arXiv preprint arXiv:2402.09267*.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. [Agieval: A human-centric benchmark for evaluating foundation models](#). Preprint, arXiv:2304.06364.

A Notations and Symbols

Table 8 summarizes the primary symbols, variables, and parameters used in the **FlipTS** framework, categorized by trajectory dynamics, RL optimization, and adaptive sampling mechanism.

Notation	Description
<i>Problem & Trajectory Formulation</i>	
q, y	Input query and its corresponding ground-truth answer.
τ	Multi-turn <i>Self-Refinement</i> trajectory $\{(s_t, a_t, r_t)\}_{t=1}^T$.
T	Maximum number of turns per episode.
s_t, a_t, r_t	State, action (solution), and verifiable reward at turn t .
<i>RL Optimization & Dynamics</i>	
$\pi_\theta, \mathcal{D}$	Policy network and the global training dataset.
$R(\tau), A_i$	Trajectory-level return and group-relative advantage.
G	Group size (number of sampled trajectories per query).
$\mathcal{J}_{\text{GRPO}}$	Objective function for trajectory-level GRPO.
<i>Active Selection via Thompson Sampling</i>	
$\xi_{k,q}$	$F2S$ transition probability for question q at step k .
α, β	Shape parameters (pseudo-counts) of the Beta distribution.
λ	Decay factor addressing non-stationarity of the policy.
ξ_{target}	Target $F2S$ probability for prioritizing trajectories.
$\mathcal{C}, M$	Candidate pool and its capacity ($M = \mathcal{C} $).
$\mathcal{O}$	Set of observed questions within the candidate pool.
η	Saturation threshold for triggering a pool refresh.
γ	Fraction of candidates to be culled during replenishment.
S_q	Posterior mean used for candidate filtering.

Table 8: Summary of notations and definitions.

B Theoretical and Empirical Analyses of Thompson Sampling Collapse

We first prove the collapse of Thompson sampling on large-scale datasets theoretically, then validate its mechanism empirically.

B.1 Theoretical Analysis of Thompson Sampling Collapse

We provide a formal justification for the adoption of Dynamic Candidate Pool. We analyze the asymptotic behavior of Thompson Sampling when the dataset size N is much larger than the sampling batch B , and show that, despite the presence of informative samples, the selection mechanism effectively degenerates into uniform random sampling: a phenomenon we term *Thompson Sampling Collapse*.

Theorem 1 (Asymptotic Collapse of Thompson Sampling Precision). Let $\mathcal{D}$ be a dataset of cardinality N , partitioned into two disjoint sets:

- $\mathcal{O}$: The set of **Observed** samples with cardinality $|\mathcal{O}| = K \ll N$.
- $\mathcal{U}_N$: The set of **Unobserved** samples with cardinality $|\mathcal{U}_N| = N - K \approx N$.

Statement: As $N \rightarrow \infty$ with fixed batch size B , the expected number of Observed samples in the batch approaches zero, causing the batch precision to converge to the prior sparsity ρ .

Proof. 1. The Budget Constraint

The selection of the top- B samples defines a selection window

$$\mathcal{W}_N = [x^* - \delta_N, x^* + \delta_N].$$

The total count B comprises selected observed samples (M_{obs}) and unobserved samples (M_{un}). Using the Law of Large Numbers to approximate discrete counts with probability integrals:

$$B = \underbrace{K \int_{\mathcal{W}_N} f_{\text{post}}(x) dx}_{\mathbb{E}[M_{\text{obs}}]} + \underbrace{(N - K) \int_{\mathcal{W}_N} f_{\text{prior}}(x) dx}_{\mathbb{E}[M_{\text{un}}]} \quad (5)$$

2. Asymptotic Shrinkage of δ_N

As $N \rightarrow \infty$, the unobserved component $(N - K)$ dominates the budget equation. To keep the total sum finite (B), the integration interval δ_N must shrink to zero, i.e., $\delta_N \rightarrow 0$. By the Mean Value Theorem for small intervals, we can approximate:

$$B \approx 2 \delta_N \left[K \cdot f_{\text{post}}(x^*) + (N - K) \cdot f_{\text{prior}}(x^*) \right].$$

Since $N \gg K$, the contribution from the unobserved portion dominates:

$$B \approx 2 \delta_N (N - K) f_{\text{prior}}(x^*),$$

Thus, the batch density is governed almost entirely by the prior of the unobserved population.

3. Asymptotic Dilution of Informative Samples

We now determine the expected number of high-value observed samples that successfully enter the batch by substituting δ_N back into the expectation term:

$$\begin{aligned} \mathbb{E}[M_{\text{obs}}] &\approx K \cdot 2 \delta_N \cdot f_{\text{post}}(x^*) \\ &\approx \frac{K \cdot B}{N} \cdot \underbrace{\frac{f_{\text{post}}(x^*)}{f_{\text{prior}}(x^*)}}_{\text{Density Ratio}}. \end{aligned} \quad (6)$$

The observed samples typically exhibit **Posterior Concentration** (low variance) due to evidence updates, whereas unobserved samples follow a diffuse prior. This implies that the local density at the target is topically higher for observed samples:

$$f_{\text{post}}(x^*) \geq f_{\text{prior}}(x^*) > 0.$$

The density ratio remains a finite constant. As the dataset size scales ($N \rightarrow \infty$), the linear growth of N in the denominator overwhelms the density benefit:

$$\lim_{N \rightarrow \infty} \mathbb{E}[M_{\text{obs}}] = 0. \quad (7)$$

Therefore, the batch becomes statistically saturated with unobserved samples ($\mathcal{B} \approx \mathcal{B}_{\text{un}} \subset \mathcal{U}_N$). Since the unobserved set is a random subset inheriting the global signal sparsity, the final batch precision reverts to ρ . $\square$

Remark 1. This theoretical insight directly motivates our sampling design. In practice, datasets are large ($N \sim 10^5$ – 10^7), while batch sizes are small ($B \sim 2^4$ – 2^7). Under this regime ($N \gg B$), Theorem 1 shows that standard Thompson Sampling suffers from **asymptotic dilution of informative samples**, effectively degenerating into uniform random sampling and failing to leverage newly discovered high-value trajectories. To mitigate this, we introduce a **Candidate Pool** of manageable size (e.g., $M = 2000$), ensuring the ratio B/M remains sufficiently large. This preserves the concentration of the posterior density $f_{\text{post}}(x^*)$ around informative samples, thereby maintaining the efficacy of Thompson Sampling for active exploitation.

B.2 Empirical Mechanism of the Collapse

We complement this theoretical analysis with a diagnostic comparison across pool scales $M \in \{2\text{k (default)}, 20\text{k}, 100\text{k}, \text{full (302k)}\}$, examined over early (steps 1–100) and late (steps 401–550) training; the statistics are defined in Table 9. Larger pools monotonically reduce both the observed share and the batch $F2S$. Full-dataset Thompson sampling remains dominated by unseen problems: only 3.1–12.5% of each batch has been observed and the batch $F2S$ is 4.3–8.1%, versus 40.5–92.3% and 16.3–42.6% at $M = 2\text{k}$. Meanwhile, the late-stage observed-only $F2S$ is similar across pools (44.2–49.3%), indicating that the collapse stems from the scarcity of observed examples rather than uniformly poorer posterior estimates. Repetition alone does not resolve it: despite a much lower

usage p75, the 20k pool approaches 2k in batch $F2S$ and in late-stage validation scores (Figure 8). These observations corroborate Theorem 1: the dynamic candidate pool prevents over-exploration of the full dataset and preserves the density of $F2S$ training signals.

C Dataset Details and Statistics

In this section, we provide detailed documentation of the datasets used for training, validation, and benchmarking. The primary domains cover competition-level mathematics, scientific reasoning, and AI safety/trustworthiness.

C.1 Training Dataset

Our training corpus is constructed by uniformly mixing the following mathematical reasoning datasets: (1) **DeepMath** (He et al., 2025): Contains 103k samples; (2) **AM-Math-Difficulty-RL** (Ji et al., 2025): We utilize the subset of difficulty level-1 (98k) and level-2 (101k). The total size of the mixed training dataset is approximately 302k samples.

C.2 Evaluation Benchmarks

We employ a diverse set of benchmarks to evaluate mathematical capability and out-of-distribution (OOD) generalization. Table 10 summarizes the statistics, domains, and evaluation settings for each dataset. Note that for TruthfulQA, we adhere to the new, improved multiple-choice version³ recommended by the authors. For SuperGPQA, we use a fixed random subset of 100 questions for evaluation.

D Implementation Details

For training, we set the maximum output length to 8192 tokens for Instruct models and 32768 tokens for Thinking models. During evaluation, the maximum output length is set to 16384 tokens for Instruct models and 38912 tokens for Thinking models. We use temperature 1.0 for Instruct models and temperature 0.6 for Thinking models in both training-time sampling and evaluation.

For the Bayesian estimator in FlipTS, all experiments use a warm-up of $W = 10$ steps and an exponential decay of $\lambda = 0.5$, shared across all model families: W determines the number of initial uniformly sampled steps used to initialize the

³<https://www.alignmentforum.org/posts/Bunfwz6JNsNd44kgLT/new-improved-multiple-choice-truthfulqa>

Table 9: Diagnostic statistics of Thompson sampling across pool scales (early = steps 1–100, late = steps 401–550). Batch F2S: *F2S* fraction of the selected batch. Observed share: fraction of the batch observed in previous steps. Observed-only F2S: *F2S* rate within that observed subset, excluding batches without observed examples. Usage p75: 75th percentile of per-question usage.

Pool	Batch F2S (%) early / late	Observed share (%) early / late	Observed-only F2S (%) early / late	Usage p75
$M = 2k$ (default)	16.3 / 42.6	40.5 / 92.3	36.0 / 45.8	14.89 / 110.86
$M = 20k$	8.0 / 39.7	14.9 / 78.9	39.2 / 48.8	1.30 / 37.16
$M = 100k$	6.8 / 30.8	9.9 / 59.2	43.3 / 49.3	0.20 / 27.58
Full (302k)	4.3 / 8.1	3.1 / 12.5	44.6 / 44.2	0.00 / 0.15

Table 10: Statistics of evaluation datasets.

Category	Dataset	Size (Questions)	Evaluation Setting
Validation	AIME24 ³ , AIME25, HMMT (Feb 2025) (Balunović et al., 2025)	30 each	Acc _{avg} @8
Math Benchmark	AIME25, HMMT (Feb 2025), BRUMO25 (Balunović et al., 2025)	30 each	Acc _{avg} @128
	CMIMC25 (Balunović et al., 2025)	40	Acc _{avg} @128
	MATH500 (Lightman et al., 2024)	500	Acc _{avg} @8
Science	Gaokao-physics (Rein et al., 2024)	174	Acc _{avg} @8 / Acc _{avg} @128 ⁴
	SuperGPQA (Subset) (Du et al., 2025)	100	Acc _{avg} @8 / Acc _{avg} @128 ⁴
Law	Law-Analytics (Zhong et al., 2023)	240	Acc _{avg} @8 / Acc _{avg} @128 ⁴
Factuality	TruthfulQA (Lin et al., 2022)	817	Acc _{avg} @8 / Acc _{avg} @128 ⁴

³ <https://huggingface.co/datasets/AI-MO/aimo-validation-aime>

⁴ For OOD datasets, Qwen3-4B-Instruct uses Acc_{avg}@128; LLaMA-3.1-8B and Qwen3-1.7B-Thinking use Acc_{avg}@8.

Beta prior via the Method of Moments, and λ discounts outdated evidence in the posterior update (Eq. 3) to handle non-stationarity.

E Baseline Details

We detail the baselines and summarize their key differences and characteristics in Table 11.

Step-level GRPO. This baseline treats solution generation and refinement as decoupled single-turn tasks. We randomly sample 20,000 questions and use the original model to generate one initial solution for each. The training is conducted via single-turn GRPO on a dataset comprising two tasks: (1) generating a correct solution directly from the question, and (2) generating a refined solution given the pre-generated initial response as context. This baseline serves to isolate the benefits of end-to-end online optimization.

Trajectory GRPO. This baseline implements the end-to-end RL framework detailed in Section 3.1, optimizing self-refinement with complete multi-turn trajectories. Unlike our proposed method, it omits the Thompson Sampling mechanism and instead adheres to the original dataset distribution for rollout generation.

Dynamic-filter GRPO. This baseline integrates the dynamic filtering spirit of DAPO into the Tra-

jectory GRPO framework. Instead of following the original DAPO formulation strictly, it utilizes on-line rejection sampling to filter out trajectories with non-positive advantages. This ensures a fair comparison by applying DAPO’s adaptive sampling logic to the same trajectory-level search space as our method. However, it incurs significant compute overhead from oversampling and rejection.

F Computation Budget

Our experimental framework and algorithms are implemented as modifications to the open-source ver1-agent library, utilizing RL training for model optimization. All experiments were conducted on a dedicated compute node equipped with 8× NVIDIA GPUs (80GB VRAM), operating with CUDA 12.2. The computational cost for training to approximately 800 update steps is as follows:

- **Qwen3-4B-Instruct:** Completing 800 update steps requires approximately 3.6 days of continuous wall-clock time, largely due to the iterative generation and self-refinement cycles.
- **Llama-3.1-8B-Instruct:** Under identical configurations, the training process for 800 steps requires approximately 2.5 days.

The method-specific overhead of FlipTS is negligible: maintaining the two posterior parameters

Table 11: Comparison of baseline methodologies across core algorithmic dimensions. E2E RL: End-to-End RL optimization of self-refinement; Sample Pruning: Online filtering of zero-advantage trajectories.

Method	E2E RL Paradigm	Thompson Sampling	Sample Pruning	Compute Efficiency	Gradient Fidelity
Step-level GRPO	×	×	×	High	Moderate
Trajectory GRPO	✓	×	×	High	Moderate
Dynamic-filter GRPO	✓	×	✓	Low [†]	High
FlipTS (Ours)	✓	✓	×	High	High

[†] Dynamic-filter GRPO increases the rollout-to-update ratio through extra online sampling and rejection.

(α, β) per problem, drawing Beta samples over the candidate pool ($M = 2000$), selecting the training batch, and periodically refreshing and culling the pool are lightweight CPU operations that involve no additional model inference or gradient computation, so training cost is dominated by rollout generation and model updates. By contrast, the oversampling and rejection of Dynamic-filter GRPO require approximately $4\times$ more generated trajectories per valid update (Table 2): measured over the first 500 common training updates, it takes 1368.3 seconds per update on average versus 807.5 seconds for FlipTS ($1.69\times$ slower), of which rollout generation accounts for 1001.1 versus 495.6 seconds ($2.02\times$ slower). FlipTS therefore achieves comparable or better performance while substantially reducing both rollout generation and wall-clock training cost.

G Sensitivity Analysis of Hyperparameters

In this section, we conduct several experiments to investigate the impact of key hyperparameters and design choices within our dynamic candidate pool framework. To provide a comprehensive view of the training dynamics, we visualize the validation scores throughout the optimization process. The validation set comprises three challenging benchmarks: AIME24, AIME25, and HMMT (Feb 2025), evaluated under a self-refinement process with a horizon of $T = 2$. We report performance using the Average Accuracy over 8 sampled trajectories (denoted as $\text{Acc}_{\text{avg}}@8$).

G.1 Target Failure-to-Success Probability

We investigate the impact of the target *Failure-to-Success* probability ξ_{target} . We evaluate the model’s performance with $\xi_{\text{target}} \in \{0.2, 0.5, 0.8\}$, and the corresponding validation scores are presented in Figure 7. The setting of $\xi_{\text{target}} = 0.8$ consistently




Figure 7: Validation scores for target Failure-to-Success probability $\xi_{\text{target}} \in \{0.2, 0.5, 0.8\}$ over training steps.

yields superior performance compared to lower values of 0.5 and 0.2. This empirical evidence validates the hypothesis that exposing the model to a higher density of Failure-to-Success transitions is essential for mastering *Self-Refinement*. Thus, prioritizing these trajectories through the sampling mechanism serves as a crucial driver for performance improvements.

We note that all main results in this paper use the pre-specified setting $\xi_{\text{target}} = 0.5$; the $\xi_{\text{target}} = 0.8$ experiment was conducted afterward as a post-hoc sensitivity analysis and was not used to select or replace any main result. Moreover, 0.5 is not a special midpoint: since the natural *F2S* rate under uniform sampling is only about 5%, both settings enrich *F2S* trajectories substantially, with 0.8 being the more aggressive enrichment. Both settings therefore support, rather than alter, our central motivation that a denser supply of successful conditional-refinement (*F2S*) trajectories benefits end-to-end *Self-Refinement* RL.

G.2 Size of Candidate Pool M

To analyze the sensitivity of our method to the pool size M , we compare the validation scores of $M \in \{2000, 20000, 100000\}$ in Figure 8. The results indicate a trade-off between convergence




Figure 8: Validation scores for pool sizes $M = 2,000$ and $M = 20,000$ over training steps.

speed and final performance. A smaller pool ($M = 2,000$) accelerates early-stage training by encouraging rapid exploitation, whereas a moderately larger pool ($M = 20,000$) exhibits a slower initial learning curve but fosters more thorough exploration, ultimately yielding superior asymptotic performance. This suggests that M can be adjusted to balance computational efficiency against peak accuracy depending on the available budget. Notably, an excessively large pool ($M = 100,000$) leads to significant performance degradation. This observation empirically validates the necessity of the candidate pool mechanism itself, aligning with our discussion in Section 3.4 and corroborate the theoretical proofs provided in Appendix B.

H Does Refinement Depend on the Previous Candidate?

An $F2S$ trajectory marks a successful refinement rather than necessarily a local edit: the model may repair the previous attempt or, finding it unusable, re-solve the problem. In either case, refinement should rely on the previous candidate; if the model simply ignored it and re-attempted independently, removing the candidate would cause no loss. Table 12 evaluates this control on Qwen3-4B-Instruct-2507.

Without the candidate, none of the models improves on average; retaining the candidate instead yields average gains of +6.24, +8.05, and +10.82 under the same-initial-answer protocol of Table 7. The gains therefore stem from conditioning on the previous answer rather than from an extra independent attempt. We thus interpret $F2S$ as *successful conditional refinement*, i.e., repairing the prior solution or re-solving it upon recognizing it is unusable, rather than a local edit in every trajectory.

I Artifact Licenses and Usage Compliance

This work strictly complies with the licenses and terms of service for all utilized artifacts. Our usage is confined to academic, non-commercial research.

- **Permissive Licenses (MIT, Apache-2.0, CC-BY-4.0, ODC-By):** We adhere to the attribution and open-source terms for DeepMath, TruthfulQA, SuperGPQA, AIME24, and AGIEval, from which Gaokao-Physics and Law-Analytics are derived. For AGIEval, we follow its MIT license for the repository and comply with the licenses and terms of the original datasets as required by AGIEval. Our codebase utilizes `verl-agent` and `math_verify` under the Apache-2.0 license.
- **Non-Commercial Licenses (CC-BY-NC-4.0, CC-BY-NC-SA-4.0):** We utilize AM-Math-Diffculty-RL, AIME25, HMMT (Feb 2025), BRUMO25, and CMIMC25 strictly for non-commercial research purposes, ensuring no violation of the “NC” (Non-Commercial) or “SA” (Share-Alike) clauses.

Attribution is provided for all cited assets to maintain transparency and reproducibility in line with community standards.

Furthermore, we confirm that, to the best of our knowledge, none of the utilized datasets contain personally identifiable information (PII). All datasets are used solely for academic, non-commercial research and evaluation purposes, in accordance with their corresponding licenses and terms of use.

J Algorithm

Algorithm 1 summarizes our proposed method, **FlipTS**, which integrates dynamic Bayesian posterior estimation with targeted Thompson sampling to guide end-to-end RL for *Self-Refinement*. It maintains a candidate pool of questions and continuously updates belief parameters based on observed trajectory outcomes. By sampling trajectories whose estimated refinement potential is closest to a target probability ξ_{target} , **FlipTS** focuses training on the model’s zone of proximal development. A refresh-and-cull mechanism ensures pool diversity and mitigates posterior staleness under non-stationary learning dynamics.

Table 12: Control evaluation with the candidate answer removed from the refinement prompt ($\text{Acc}_{\text{avg}}@128$, %). Each cell reports *Initial / 2-turn refinement without candidate*; Δ_{avg} is the change of the averaged scores.

Model	AIME25	HMMT (Feb 2025)	BRUMO25	CMIMC25	Average	Δ_{avg}
w/o training	48.39 / 46.90	31.07 / 30.68	58.83 / 58.67	30.74 / 31.45	42.26 / 41.92	−0.34
Trajectory GRPO	52.32 / 52.08	33.31 / 33.02	61.28 / 59.77	33.83 / 33.54	45.18 / 44.60	−0.58
FlipTS	54.69 / 53.83	32.73 / 32.37	60.73 / 60.81	33.91 / 34.28	45.51 / 45.32	−0.19

K Self-Refinement Prompt Template

Our self-refinement process is structured as a multi-turn dialogue. The specific prompt configurations for each stage of the process are detailed below.

System Message (All Turns) For every turn in the conversation (both initial generation and subsequent refinements), the system message remains constant to enforce step-by-step reasoning and correct output formatting:

System:

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

Turn 1: Initial Generation In the first turn, the user message consists solely of the math problem. The model generates the initial candidate solution based on this input.

User:

{question}

Subsequent Turns: Refinement For all subsequent turns, the user message follows a structured refinement template. This template feeds the problem and the previous candidate solution (which may be flawed) back into the model, instructing it to critique and improve the answer.

User:

You are given a math problem and a potential solution.
This attempt may be flawed or partial.
Correct any errors and build upon its ideas to produce an improved, high-quality solution.
If the original attempt is unusable, generate a correct solution from scratch.
End with the final result in `\boxed{}`.

Problem:

{question}

Candidate solution (may contain mistakes):

-- Candidate --

{candidate_answer}

Now refine the candidate into an improved solution.

Provide clear reasoning and end with the final answer in `\boxed{}`.

Here, {question} is the original problem statement, and {candidate_answer} is the model’s output from the immediate previous turn.

L

Algorithm 1: FlipTS

Input: Dataset $\mathcal{D}$; Batch size B ; Decay factor λ ; Target probability ξ_{target} ; Warm-up steps W ; Pool size M ; Refresh threshold η ; Culling fraction γ ; Initial Policy π_{θ_0} .

Output: Optimized Policy π_{θ_N}

```
1 for  $k = 1, 2, \dots, N$  do
  // Thompson Sampling for Questions
2  if  $k \leq W$  then
3    | Sample batch  $\mathcal{B}_k$  uniformly from  $\mathcal{D}$ ;
4  end
5  else
6    | Sample a proxy probability  $\tilde{\xi}_{k,q} \sim \text{Beta}(\alpha_{k,q}, \beta_{k,q})$  for each  $q \in \mathcal{C}$ ;
7    | Select batch  $\mathcal{B}_k$  as the top- $B$  questions in  $\mathcal{C}$  minimizing  $|\tilde{\xi}_{k,q} - \xi_{target}|$  (Eq. 4);
8  end
  // RL Training (Trajectory-level GRPO)
9  Generate trajectories  $\{\tau_i\}_{i=1}^{B \times G}$  using  $\pi_{\theta_k}$  for questions in  $\mathcal{B}_k$ ;
10 Update  $\theta_k$  to  $\theta_{k+1}$  by maximizing  $\mathcal{J}_{GRPO}(\theta)$  (Eq. 1);
  // Bayesian Estimation
11 Count the number of trajectories with and without  $F2S$  transition  $s_q, f_q$ ; compute  $\hat{\xi}_q = s_q/G$ ;
12 if  $k = W$  then
13   | Compute prior parameters  $\alpha_0, \beta_0$  using the empirical probabilities  $\{\hat{\xi}_q\}$  from the first  $W$ 
   | steps (Eq. 2); initialize  $\alpha_{k,q} \leftarrow \alpha_0, \beta_{k,q} \leftarrow \beta_0, \forall q \in \mathcal{D}$ ;
14 end
15 if  $k > W$  then
16   | Update  $\alpha_{k+1,q} \leftarrow s_q + \lambda\alpha_{k,q}, \beta_{k+1,q} \leftarrow f_q + \lambda\beta_{k,q}$  for each  $q \in \mathcal{B}_k$  (Eq. 3);
17 end
  // Dynamic Pool Mechanism
18 if  $k = W$  then
19   | Initialize candidate pool  $\mathcal{C}$  by uniformly sampling  $M$  questions from  $\mathcal{D}$ ;
20   | Initialize observed questions in pool  $\mathcal{O} = \emptyset$ ;
21 end
22 if  $k > W$  then
23   | Update observed questions:  $\mathcal{O} \leftarrow \mathcal{O} \cup \mathcal{B}_k$ ;
24   | if  $|\mathcal{O}|/M \geq \eta$  then
25     | Calculate scores  $S_q = \alpha_{k,q}/(\alpha_{k,q} + \beta_{k,q})$  for  $q \in \mathcal{O}$ ;
26     | Identify cull set  $\mathcal{Q}_{rem} \subset \mathcal{O}$  with  $\lfloor \gamma M \rfloor$  questions with largest  $|S_q - \xi_{target}|$ ;
27     | Uniformly sample a new set  $\mathcal{Q}_{new}$  from  $\mathcal{D} \setminus \mathcal{C}$  with  $|\mathcal{Q}_{rem}| = |\mathcal{Q}_{new}|$ ;
28     | Update pool  $\mathcal{C} \leftarrow (\mathcal{C} \setminus \mathcal{Q}_{rem}) \cup \mathcal{Q}_{new}; \mathcal{O} \leftarrow \mathcal{O} \setminus \mathcal{Q}_{rem}$ 
29   | end
30 end
31 end
32 return  $\pi_{\theta_N}$ 
```
