

SAPO: Self-Adaptive Process Optimization Makes Small Reasoners Stronger

Kaiyuan Chen, Guangmin Zheng, Jin Wang*, Xiaobing Zhou, Xuejie Zhang

School of Information Science and Engineering, Yunnan University, Kunming, China
chenkaiyuan@stu.ynu.edu.cn, gmzheng@mail.ynu.edu.cn, {wangjin, zhouxb, xjzhang}@ynu.edu.cn

Abstract

Existing self-evolution methods overlook the influence of fine-grained reasoning steps, which leads to the reasoner-verifier gap. The computational inefficiency of Monte Carlo (MC) process supervision further exacerbates the difficulty in mitigating the gap. Motivated by the Error-Related Negativity (ERN), which the reasoner can localize error following incorrect decisions, guiding rapid adjustments, we propose a Self-Adaptive Process Optimization (SAPO) method for self-improvement in Small Language Models (SLMs). SAPO adaptively and efficiently introduces process supervision signals by actively minimizing the reasoner-verifier gap rather than relying on inefficient MC estimations. Extensive experiments demonstrate that the proposed method outperforms most existing self-evolution methods on two challenging task types: mathematics and code. Additionally, to further investigate SAPO’s impact on verifier performance, this work introduces two new benchmarks for process reward models in both mathematical and coding tasks.

Introduction

Recent advances (Achiam et al. 2023; Guo et al. 2025) in Large Language Models (LLMs) highlight their superiority in complex multi-step planning, particularly with Chain-of-Thought (CoT) (Wei et al. 2022; Kojima et al. 2022). Despite outperforming SLMs ($\leq 2B$), LLMs’ high computational and storage costs are prohibitive. Thus, developing efficient but strong SLMs for mobile devices is promising (Magister et al. 2023; Chen, Wang, and Zhang 2024).

Previous approaches (Pang et al. 2024; Jiao et al. 2024; Guan et al. 2025) have applied the self-evolutionary concept to LLMs to enhance reasoning performance, offering greater flexibility and efficiency. However, due to the poor instruction-following ability and reasoning performance, the general reward evaluation methods used on LLMs (Zheng et al. 2023; Chen et al. 2024; Liu et al. 2025b) are unsuitable for SLMs. In contrast, current classification-based discriminative methods are an efficient reward allocation approach better suited for the self-improvement of SLMs, typically a system composed of a reasoner and a verifier (Hosseini et al. 2024; Chen, Wang, and Zhang 2025), as shown in Figure 2.

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

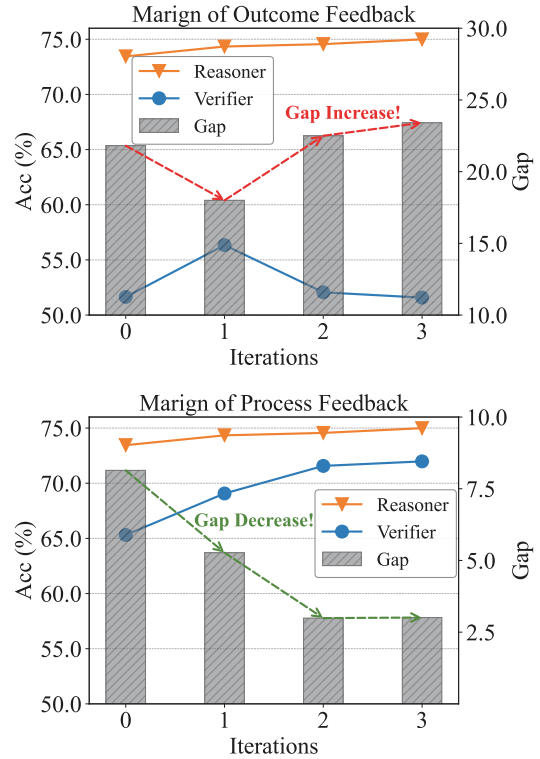


Figure 1: The performance and gap dynamics of the online reasoner and verifier across iterations in the step correction prediction of GSM_Process. The selected baseline model is Qwen-2.5-0.5B, and the reasoner estimates step correctness using a MC approach.

Nevertheless, most existing self-evolution methods still overlook feedback on fine-grained reasoning steps and instead opt for outcome rewards (Uesato et al. 2022; Lightman et al. 2024) as a more efficient form of supervision (Yu et al. 2025; Guo et al. 2025), which inevitably leaves room for reward hacking (Liu et al. 2025a; Yu et al. 2025). Thus, Monte Carlo-based process supervision transfers the posterior estimation of reasoning to the verifier model through step-level annotations, delivering superior performance (Jiao et al. 2024; Guan et al. 2025; Chen, Wang, and Zhang 2025).

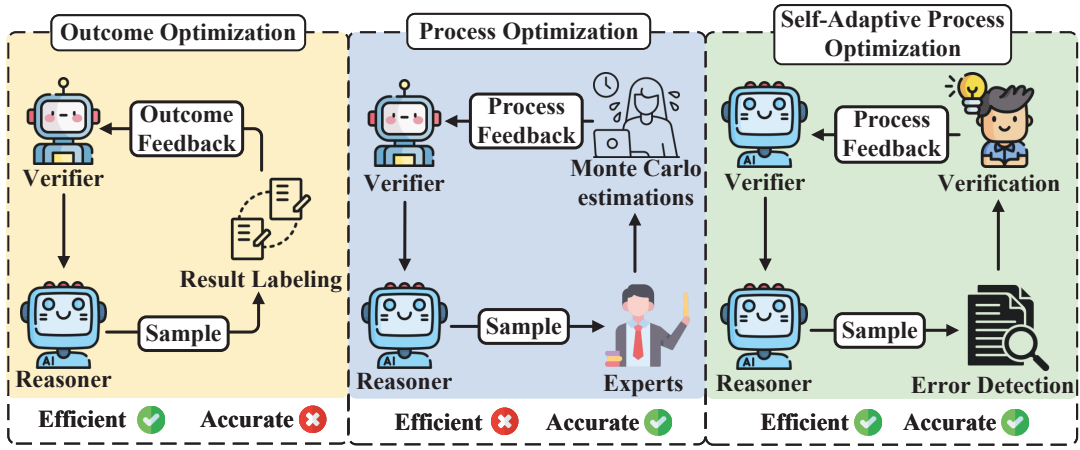


Figure 2: A comparison between the Self-Adaptive Process Optimization (SAPO) and the previous self-evolution framework. SAPO actively detects potential *first error positions* to determine which steps need to be verified, rather than performing step-by-step rollout estimation.

To better understand the limitations of self-evolution methods guided by process-level feedback, this work investigates the bias issues arising within the *exploration-exploitation* paradigm formed by the interaction between the reasoner and the verifier. Figure 1 (top) shows that, under multiple rounds of self-iteration, the absence of process supervision signals leads to an increasing reasoner-verifier gap. This gap undermines the verifier’s ability to assess the quality of reasoning paths accurately. In contrast, Figure 1 (bottom) shows that introducing online process supervision effectively reduces the reasoner-verifier gap.

However, the Monte Carlo-based process-supervised verifier (Wang et al. 2024; Jiao et al. 2024; Guan et al. 2025) is computationally expensive and highly inefficient. It poses a major bottleneck to scalability and practical application, as shown in Figure 2. If each problem takes 8 steps and has 10 reasoning trajectories, labeling 10k problems requires 800,000 rollouts. Thus, achieving collaborative optimization between the inference engine and the verifier without significantly increasing supervisory costs is a key challenge in the current self-evolution methods.

Therefore, we propose a novel Self-Adaptive Process Optimization (SAPO). This method enhances the efficiency of process supervision by introducing necessary step labels through localized error detection and correction. This new idea is motivated by the Error-Related Negativity (ERN) (Yeung, Botvinick, and Cohen 2004), whereby humans can spontaneously detect and localize errors shortly after making incorrect decisions, thereby guiding subsequent cognitive adjustments and behavioral corrections.

Refer to Figure 2 for a high-level understanding of SAPO. During the construction of process supervision, potential *first error locations* (Uesato et al. 2022; Lightman et al. 2024) are identified through the changes and ordinality in reward scores. Then, the verifier is adaptively adjusted to minimize the discrepancies at these locations with the rollout estimates by the reasoner. The contributions of this work can be summarized as follows:

- 1) This study proposes SAPO, a self-evolution paradigm that balances efficiency and performance by replacing inefficient step-by-step verification with first error detection and posterior estimation.
- 2) To better evaluate the performance of different verifiers, this study introduces two datasets, GSM_Process and MBPP_Process, for reasoning step correctness verification in mathematics and code, respectively, based on first-error annotation. The Self-Adaptive PRM (SAPRM) also outperforms other reward models for process verification.
- 3) Extensive experiments demonstrate that SAPO outperforms existing state-of-the-art self-evolution methods, e.g., GRPO (Shao et al. 2024), on two challenging categories of multi-step reasoning tasks: mathematics and code. Moreover, SAPO demonstrates superior performance in addressing the bias between reasoning and verification.

Preliminary

Problem Definition

Given a question q , a reasoning model M is to obtain its final answer a through a certain reasoning trajectory τ . Based on the Markov Decision Process (MDP) (Hao et al. 2023; Li and Li 2024), a reasoning trajectory can be defined as follows:

$$s_{j+1} \sim M(s_{(0:j)}, q|\theta) \quad (1)$$

$$\tau = \{ \langle s_0, c_0 \rangle, \langle s_1, c_1 \rangle, \dots, \langle s_n, c_n \rangle \} \quad (2)$$

where θ is the weight of the model M , s_i is one of the immediate steps and $s_{(1:j)}$ represents an unfinished trajectory ending at s_j . Additionally, each step can be further assigned a corresponding correctness label $c \in \{0, 1\}$.

Process Reward Supervision

An independent label can be assigned to each step, which can be manual. Recent works (Wang et al. 2024; Jiao et al.

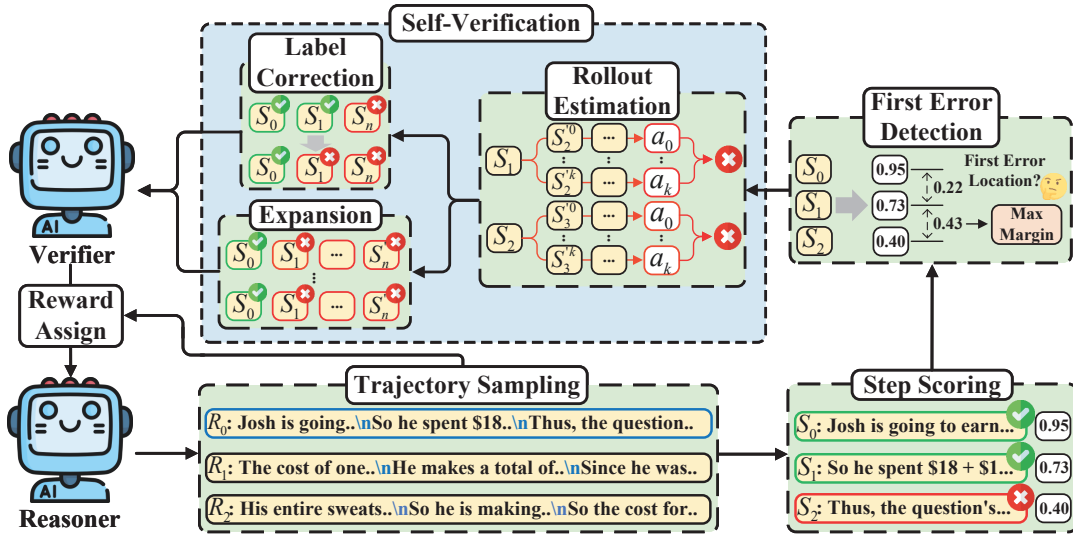


Figure 3: Overall framework diagram of Self-Adaptive Process Optimization (SAPO). The method adopts a self-iterative framework where the verifier pre-assigns step-level scores, error detection locates the first likely error, and the reasoner revisits it for posterior estimation. The corrected reasoning step labels then supervise the verifier, enabling the reasoner to self-optimize under more accurate process rewards.

2024) have adopted an MC estimation approach for automated labeling. Employing multiple rollouts starting from a certain step s_j to estimate its correctness c_j , as follows:

$$\text{Rollout}(M, q, s_{(0:j)}) = \{M(q, s_{(0:j)}^k | \theta)\}_{k=1}^T \quad (3)$$

$$= \{(s_0^k, \dots, s_j^k, s_{j+1}^{k'}, \dots, s_m^{k'}, a_k)\}_{k=1}^T$$

$$A = \{a_1, a_2, \dots, a_T\} \quad (4)$$

$$c_j = \begin{cases} 1, & \text{if } \exists a_k \in A, a_k = a^* \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

where T is the number of the sampling, and m is the length of the completed trajectory. However, MC estimation is always computationally inefficient. Thus, a classification-based Process Reward Model (PRM) can be trained using Mean Square Error (MSE) loss, as follows:

$$\mathcal{L}_{PRM} = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{M_i} (f(s_{(0:j)}^i; q) - c^k)^2 \quad (6)$$

where f denotes the predicted output of the classification head, N denotes the total number of questions q , and M_i represents the total steps in the i -th reasoning trajectory. Typically, the PRM (or verifier V) can be initialized from M .

Self-Adaptive Process Optimization

Reasoning Trajectory Sampling

For a given question q , diverse reasoning trajectories can be obtained by the reasoner M through high-temperature T sampling (Yuan et al. 2023b), as follows:

$$\text{Sample}(M, q) = \{\tau_i | \tau_i \sim M(q, T, \theta)\}_{i=1}^K \quad (7)$$

K denotes the number of sampled trajectories for a given problem. All sampled trajectories are deduplicated to ensure diversity, provided the minimum sampling count is met.

Self-Adaptive Process Supervision

First Error Detection. According to Equation (6), the verifier V can assign a reward score to any step s_j within a sampled trajectory τ , as follows:

$$\hat{c}_j = f(s_{(0:j)}; q) \in [0, 1] \quad (8)$$

according to previous works (Uesato et al. 2022; Lightman et al. 2024), the first error position strategy is sufficient to provide effective process supervision signals. As shown in Figure 3, the step with the maximum score difference Δ_j can be defined as the potential *first error position*:

$$\Delta_j = \hat{c}_j - \hat{c}_{j-1} \quad (9)$$

$$\hat{t} = \text{argmax}_{j \in \{1, 2, \dots, n-1\}} \Delta_j \quad (10)$$

where \hat{t} denotes the potential first error position predicted by the V . Thus, step-level labels for a reasoning trajectory can also be pre-assigned:

$$\tau = \{\langle s_0^w, c_0^w \rangle, \dots, \langle s_{\hat{t}}^w, c_{\hat{t}}^w \rangle, \langle s_{\hat{t}}^l, c_{\hat{t}}^l \rangle, \dots, \langle s_n^l, c_n^l \rangle\} \quad (11)$$

here, $c^w = 1$ and $c^l = 0$ indicate that the current step is correct and incorrect, respectively.

Self-Verification. The *first error position* predicted by the verifier is not unbiased, as shown in Figure 3. There may be a discrepancy between the predicted position \hat{t} and the gold position t : (a) If $c_{\hat{t}} = 1$ and $c_{\hat{t}+1} = 0$: $\hat{t} = t$; (b) If $c_{\hat{t}} = 1$ and $c_{\hat{t}+1} = 1$: $\hat{t} > t$; (c) If $c_{\hat{t}} = 0$ and $c_{\hat{t}+1} = 0$: $\hat{t} < t$.

According to the rollout estimation (Equations (3), (4), and (5)), each case requires only two rollouts, which is significantly fewer than the step-by-step rollout estimation.

Case (a) refers to the prediction being correct. For cases (b) and (c), the pre-assigned labels will be corrected as follows:

$$(b) \quad \tau = \{ \langle s_0^w, c_0^w \rangle, \dots, \langle s_{\hat{i}}^w, c_{\hat{i}}^w \rangle, \langle s_{\hat{i}}^l, c_{\hat{i}}^l \rangle, \dots, \langle s_n^l, c_n^l \rangle \} \quad (12)$$

$$(c) \quad \tau = \{ \langle s_0^w, c_0^w \rangle, \dots, \langle s_{\hat{i}}^l, c_{\hat{i}}^l \rangle, \langle s_{\hat{i}}^w, c_{\hat{i}}^w \rangle, \dots, \langle s_n^l, c_n^l \rangle \} \quad (13)$$

Expansion. Naively correcting a single-step error to reduce bias is suboptimal, as the corrected result may not generalize well to unseen cases. Fortunately, rollout inherently simulates diverse scenarios through sampling. Therefore, based on the *first error location*, the extended trajectories can be retained to enhance the generalization of verification:

- (i) If $c_j = 1$, it indicates that all extended trajectories with $s_{(1:j)}$ as a prefix and a correct final result are correct at every step.
- (ii) If $c_j = 0$, it indicates that any suffix $s_{(j+1:m)}$ starting from s_j leads to an incorrect final result being incorrect at every step.

Reasoning Optimization

The verifier trained with self-adaptive process supervision can assign overall reward scores to sampled reasoning trajectories and construct a preference dataset, as follows:

$$r(\tau) = \frac{\sum_{j=1}^m \hat{c}_j}{m} = \frac{\sum_{j=1}^m f(\hat{s}_{(0:j)}; q)}{m} \quad (14)$$

$$D_{pref} = \{q_i, \tau_i^w, \tau_i^l \mid r(\tau_i^w) - r(\tau_i^l) \geq \eta\} \quad (15)$$

where η denotes the threshold, while τ^w and τ^l represent the positive and negative samples, respectively. Given the preference data, SAPO employs the ORPO algorithm (Hong, Lee, and Thorne 2024) to achieve self-alignment of the reasoner M , as follows:

$$\mathcal{L}_{ORPO} = \mathbb{E}_{(q, \tau^w, \tau^l)} \left[\mathcal{L}_{SFT}(q, \tau^w) - \beta \left(\log \sigma \left(\frac{\text{odds}(\tau^w \mid q)}{\text{odds}(\tau^l \mid q)} \right) \right) \right] \quad (16)$$

$$\text{odds}(\tau \mid q) = \frac{P(\tau \mid q)}{1 - P(\tau \mid q)} \quad (17)$$

Iterative Self-Optimization

The proposed SAPO follows an iterative *exploration-exploitation* paradigm, where the reasoner is guided toward self-optimization by the verifier’s progressively refined process supervision signals. The framework can be illustrated in Algorithm 1.

As indicated, *Detect* and *Verify* follow the strategies outlined before, and each verifier V is initialized from the reasoner. Since the algorithm requires a trained verifier, we use binary estimation (Luo et al. 2024) (denoted as Ω) to perform the initial step labeling.

Algorithm 1: Procedure of SAPO

- 1: Initialize: A pretrained SLM M ; Original dataset $D_{\text{org}} = \{(q_i, \tau_i)\}_{i=1}^N$
 - 2: $M_0 \leftarrow \text{SFT}(M, D_{\text{org}})$
 - 3: $D_0^{\text{sample}} \leftarrow \text{Sample}(M_0, D_{\text{org}})$
 - 4: $D_0 \leftarrow D_0^{\text{sample}} \cup D_{\text{org}}$
 - 5: $D_{\text{step_label}} \leftarrow \Omega(M_0, \text{Subset}(D_0^{\text{sample}}))$
 - 6: $V_0 \leftarrow \text{SFT}(M_0, D_{\text{step_label}})$
 - 7: **for** $i = 1$ **to** T **do**
 - 8: $D_i^{\text{sample}} \leftarrow \text{Sample}(M_i, D_{i-1})$
 - 9: $D_i \leftarrow D_i^{\text{sample}} \cup D_{i-1}$
 - 10: $D_{\text{score}} \leftarrow \text{Score}(V_{i-1}, \text{Subset}(D_i^{\text{sample}}))$
 - 11: $D_{\text{step_detect}} \leftarrow \text{Detect}(D_{\text{score}})$
 - 12: $D_{\text{step_label}} \leftarrow \text{Verify}(M_i, D_{\text{step_detect}})$
 - 13: $V_i \leftarrow \text{SFT}(M_i, D_{\text{step_label}})$
 - 14: $D_{\text{perf}} \leftarrow \text{Score}(V_i, D_i)$
 - 15: $M_{i+1} \leftarrow \text{Align}(M, D_{\text{perf}})$
 - 16: **end for**
 - 17: **Return** The M_T after iteration.
-

Experiments

Benchmark

Two challenging multi-step reasoning benchmarks are selected to evaluate the model’s reasoning capability: GSM8K (Cobbe et al. 2021) for mathematical reasoning and MBPP (Austin et al. 2021) for code generation. MATH (Hendrycks et al. 2021) and HumanEval (Chen et al. 2021) are used to evaluate the Out-of-Domain (OOD) generalization on math and code tasks. Accuracy and Pass@1 (Chen et al. 2021) are evaluation metrics for math and code, respectively.

Previous studies used Best-of-N evaluation (Cobbe et al. 2021) for verifier models, but it is unreliable for PRM (Zheng et al. 2024a; Zhang et al. 2025). Instead, process-level verification provides a more dependable evaluation.

To enable step-wise verification in math and code tasks, we introduce two new benchmarks: GSM_Process and MBPP_Process. For each question in GSM8K and MBPP, different models generated two reasoning paths. GPT-4o (Achiam et al. 2023) was used to annotate the first incorrect step in each path, with results filtered for reliability. These benchmarks contain 3,786 and 1,499 examples, and evaluate verifiers based on step correction prediction accuracy.

Baselines

Except for Chain-of-Thought (CoT) (Wei et al. 2022) and Supervised Fine-Tuning (SFT), several strong self-evolution methods are also implemented as baselines, including: Rejection Sampling Fine-Tuning (RFT) (Yuan et al. 2023b), RFT combined with Direct Preference Optimization (Rafailov et al. 2023) (RFT+DPO), Reasoning Preference Optimization (RPO) (Pang et al. 2024), Group Relative Policy Optimization (Shao et al. 2024) after SFT (SFT+GRPO). Additionally, ORM (Lightman et al. 2024), OmegaPRM (Luo et al. 2024), and ShepherdPRM (Wang et al. 2024) are implemented as baselines for verifiers’ comparison.

Model	Qwen-2.5-0.5B		Llama-3.2-1B		Gemma-2-2B	
Method\Task	GSM8K	MATH(OOD)	GSM8K	MATH(OOD)	GSM8K	MATH(OOD)
CoT	28.51	25.97	5.31	3.65	19.86	16.22
SFT	34.19	24.84	22.14	2.60	39.12	20.95
RFT	37.83	27.35	26.31	<u>5.56</u>	45.34	22.14
RFT+DPO	8.26	12.45	23.88	4.06	41.02	15.38
Online-RFT	40.79	28.85	29.03	5.45	<u>48.67</u>	24.60
RPO	41.55	22.68	29.26	4.85	<u>45.41</u>	11.96
SFT+GRPO	46.24	34.53	26.46	5.92	44.65	24.36
SAPO-iter1	36.99	30.04	29.41	5.51	46.10	24.00
SAPO-iter2	38.14	31.48	<u>32.60</u>	5.92	48.07	<u>24.36</u>
SAPO-iter3	<u>41.62</u>	<u>31.72</u>	34.19	5.45	49.73	24.24
Method\Task	MBPP	HumanEval(OOD)	MBPP	HumanEval(OOD)	MBPP	HumanEval(OOD)
CoT	24.44	25.60	21.94	<u>19.51</u>	22.61	17.07
SFT	29.32	25.60	24.01	<u>19.51</u>	29.62	21.34
RFT	30.36	26.82	26.35	21.34	29.05	18.90
RFT+DPO	16.29	20.73	25.95	17.68	18.03	16.46
Online-RFT	34.00	28.04	28.25	15.85	34.20	20.73
RPO	<u>36.63</u>	32.31	28.45	15.24	34.93	18.29
SFT+GRPO	<u>35.20</u>	29.26	25.55	13.41	33.40	24.39
SAPO-iter1	33.83	<u>31.09</u>	25.35	18.90	32.76	22.56
SAPO-iter2	34.06	30.48	<u>28.82</u>	21.34	<u>35.17</u>	<u>23.78</u>
SAPO-iter3	36.67	30.48	28.92	21.34	35.43	24.39

Table 1: Comparative experimental results on multi-step reasoning tasks in both mathematics reasoning and code generation. Bold indicates the best performance, while underline represents the second best.

Setup

Qwen-2.5-0.5B (Yang et al. 2024), Llama-3.2-1B (Grattafiori et al. 2024), and Gemma-2-2B (Team et al. 2024) serve as evaluation backbones. For SAPO, distinct reasoning trajectories are sampled per question per iteration, using the same model as the verifier. Full fine-tuning is applied to Qwen-2.5-0.5B and Llama-3.2-1B, while QLoRA (Detrmers et al. 2023) fine-tunes Gemma-2-2B for efficiency. The publicly available Unsloth framework is used to accelerate training and inference, and all experiments are run on two NVIDIA 3090 GPUs.

Comparison of Different Self-Evolution Methods

Table 1 shows that SAPO consistently outperforms most other methods across different models and tasks, both in-domain and out-of-domain. Moreover, reasoning performance of SAPO can further improve with more iterations.

Applying DPO after RFT to align positive and negative sample results in worse performance, consistent with previous finding (Feng et al. 2024). And effective improvement can be achieved by emphasizing positive examples during the iterative process, as demonstrated by RPO (Pang et al. 2024) and ORPO (Hong, Lee, and Thorne 2024).

Although Qwen-2.5-0.5B trained with GRPO achieves the best performance on math tasks, this improvement is highly model-dependent. When the base model has weaker capabilities (e.g., LLaMA or Gemma) or when applied to code tasks, GRPO does not outperform other methods. Moreover, without SFT, training SLMs with GRPO struggles to converge to high-reward.

Comparison of Different Verifiers

Figure 4 presents the performance trends of SAPRM on mathematical and code tasks. With only local step corrections introduced, SAPRM achieves progressively improved process verification performance across iterations, often outperforming ORM, particularly on code verification. In contrast, Online ORM exhibits a clear performance ceiling—it’s verification accuracy tends to decline after reaching a peak. Moreover, SAPRM performs better with larger models.

To further investigate the reasoning-verification bias under different approaches, we also implemented two commonly used baselines, OmegaPRM and ShepherdPRM, for comparison, as shown in Figure 5. Additionally, we used step prediction accuracy based on MCE starting from each step of the reasoning process as an upper bound for each iteration. Figure 5 shows that SAPRM achieves lower bias on GSM.Process than other verifier models, despite only verifying and correcting local information.

Further Analysis

Analysis of Process Supervision Efficiency

As described before, Self-Adaptive Process Supervision (SAPS) enhances efficiency by identifying and verifying only the first potential error, thereby avoiding unnecessary step-by-step rollout. To validate the effectiveness of the proposed method, we compare SAPS with two current process labeling approaches, including Shepherd and Omega, on processing supervision signal annotation efficiency. The comparison uses FLOPs and elapsed time (in seconds) on




Figure 4: The performance evolution of SAPRM across different tasks (GSM_Process and MBPP_Process) and models over iterations. The online Outcome Reward Model (ORM) is used as a baseline for comparison.




Figure 5: Multi-round iterative performance comparison of different reward models on GSM_Process. Monte Carlo Estimation (MCE) serves as the upper bound of verification performance achievable by the reasoner, and Qwen-2.5-0.5B is used as the base model.

different tasks. The evaluation was conducted using Qwen-2.5-0.5B on a single NVIDIA RTX 3090 GPU. For the code tasks, compilation time was also taken into account.

Figure 6 shows that, compared to Shepherd, which relies on step-wise rollout for verification, SAP achieves a 2–3x improvement in FLOPs and time consumption. Moreover, SAPS outperforms Omega, which uses binary estimation to locate the predicted first error position.

Analysis of Self-Verification Errors

Self-verification error rate reflects the consistency between the reasoner and the verifier. Additionally, the efficiency of SAPO also depends on the number of mismatches or errors that occur when the reasoner is used to verify the labels pre-assigned by the verifier. Thus, it is worth exploring the relationship between the reasoner and the verifier across different iterations. Specially, Pre-assigned labels are considered correct in only three cases: (1) correct result with all correct steps, (2) incorrect result with all incorrect steps, or (3) incorrect result with correctly locating the first error.

To avoid interference from varying samples, we uniformly adopt the reasoning trajectories used by three different models from GSM_Process and MBPP_Process for self-verification error evaluation. The analysis is conducted to determine whether the reasoner and verifier are aligned.

Table 2 shows that self-verification performs best when the reasoner and verifier are from the same round (e.g., $V3-R3$), while mismatched pairs lead to higher error rates, highlighting the importance of online synchronization. Notably, although $V1-R1$ is also synchronized, its higher error rate compared to $V3-R3$ suggests that iteration helps close the gap between the verifier and reasoner.

Ablation Study Analysis

In the ablation study, *w/o PF* removes process feedback from verifier updates, assigning step labels only by final result; *w/o RM* excludes the verifier, using only result correctness to label samples; *w/o DV* replaces verifier-guided correction with random step correction; *w/o EP* disables verification expansion. Table 3 shows that removing any compo-




Figure 6: Comparison of efficiency in process labeling methods. The figure illustrates the average FLOPs and time (in seconds) required to label a single sample under different process supervision methods.

Model	Qwen-2.5-0.5B	Llama-3.2-1B	Gemma-2-2B
Task	GSM8K		
<i>VI-R1</i>	58.59	60.34	54.51
<i>VI-R3</i>	58.66	61.49	56.48
<i>V3-R3</i>	54.63	54.46	51.36
Task	MBPP		
<i>VI-R1</i>	67.48	78.22	67.48
<i>VI-R3</i>	68.85	76.96	67.31
<i>V3-R3</i>	67.08	76.21	63.29

Table 2: The self-verification error rates (%) obtained by the different reasoner and verifier. *VI-R3* indicates that the verifier from round 1 is used to pre-assign labels, while the reasoner from round 3 is used to perform self-verification.

ment leads to varying degrees of performance degradation, demonstrating their necessity. Moreover, *w/o PF* and *w/o RM* have distinct impacts despite both relying solely on outcome feedback. *w/o PF* leads to a more significant overall performance degradation due to its effect on the PRM.

Related Work

Self-Evolution of Language Reasoning. Reasoning models have long been a central focus of research (Guo et al. 2025; Shen et al. 2025a,b; Zheng et al. 2024b, 2025; Yuan et al. 2023a, 2025). And self-evolution is central to recent reasoning methods. STaR (Zelikman et al. 2022) and RFT (Yuan et al. 2023b) apply supervised learning to model-generated traces. RPO (Pang et al. 2024) and V-STaR (Hosseini et al. 2024) both use offline reinforcement (e.g., DPO) to align positive and negative samples over multiple rounds, while V-STaR further integrates a verifier to internalize this alignment. Current R1-like self-reflective reasoning (Liu et al. 2025a; Yu et al. 2025) that centers around algorithms like GRPO (Shao et al. 2024) or PPO (Schulman et al. 2017) can also be categorized as self-evolution. Specifically, these methods iteratively adjust the model through online sampling to produce higher-reward outputs.

While self-evolution approaches have improved pre-trained models’ reasoning, they rely mainly on outcome

Method\Model	Qwen-2.5-0.5B	Llama-3.2-1B
SAPO	41.62	34.19
<i>w/o PF</i>	39.65	32.37
<i>w/o DV</i>	40.86	31.69
<i>w/o RM</i>	40.71	32.75
<i>w/o EP</i>	40.33	33.73

Table 3: Ablation study on GSM8K. *PF* denotes process feedback, *DV* stands for error detection and self-verification, *RM* represents the reward model (or verifier), and *EP* refers to expansion during verification.

feedback, which leaves room for reward hacking during alignment (Liu et al. 2025a).

Process-guided Long-Chain Reasoning. Compared to outcome-level supervision, step-level process supervision has gained attention for offering denser feedback signals. Early studies (Uesato et al. 2022; Lightman et al. 2024) manually labeled the correctness of individual reasoning steps to train a PRM. Recent works (Wang et al. 2024; Luo et al. 2024; Jiao et al. 2024) have focused on automating step verification using Monte Carlo-based rollout estimation.

Research based on process supervision generally falls into two categories: (1) training-free methods (Hao et al. 2023; Guan et al. 2025) using Monte Carlo Tree Search (MCTS) for test-time performance scaling, and (2) self-evolutionary approaches (Jiao et al. 2024; Chen, Wang, and Zhang 2025) that improve reasoning via offline PRMs in an *exploration-exploitation* framework. However, these self-evolution methods still face inefficient training and weak alignment between rewards and reasoning steps.

Conclusion

This paper proposed the Self-Adaptive Process Optimization (SAPO) method to improve multi-step reasoning performance in SLMs. Specifically, the method gradually enhances reasoning performance by adapting to the bias between the reasoner and the verifier. Extensive experimental results show that the proposed method outperforms existing self-evolution methods and is more efficient in step labeling than previous approaches.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant Nos. 61966038 and 62266051, and the Postgraduate Practice and Innovation Foundation of Yunnan University under Grant No. ZC-242410094. We would like to thank the anonymous reviewers for their constructive comments.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; Michalewski, H.; Dohan, D.; Jiang, E.; Cai, C.; Terry, M.; Le, Q.; et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Chen, K.; Wang, J.; and Zhang, X. 2024. Mathematical reasoning via multi-step self questioning and answering for small language models. In *CCF International Conference on Natural Language Processing and Chinese Computing*, 81–93. Springer.
- Chen, K.; Wang, J.; and Zhang, X. 2025. Learning to reason via self-iterative process feedback for small language models. In *Proceedings of the 31st International Conference on Computational Linguistics*, 3027–3042.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; Pinto, H. P. D. O.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Chen, Z.; Deng, Y.; Yuan, H.; Ji, K.; and Gu, Q. 2024. Self-play fine-tuning converts weak language models to strong language models. *arXiv preprint arXiv:2401.01335*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Feng, D.; Qin, B.; Huang, C.; Zhang, Z.; and Lei, W. 2024. Towards analyzing and understanding the limitations of dpo: A theoretical perspective. *arXiv preprint arXiv:2404.04626*.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Guan, X.; Zhang, L. L.; Liu, Y.; Shang, N.; Sun, Y.; Zhu, Y.; Yang, F.; and Yang, M. 2025. rStar-Math: Small LLMs Can Master Math Reasoning with Self-Evolved Deep Thinking. *arXiv preprint arXiv:2501.04519*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hao, S.; Gu, Y.; Ma, H.; Hong, J.; Wang, Z.; Wang, D.; and Hu, Z. 2023. Reasoning with Language Model is Planning with World Model. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 8154–8173. Singapore: Association for Computational Linguistics.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Hong, J.; Lee, N.; and Thorne, J. 2024. Orpo: Monolithic preference optimization without reference model. *arXiv preprint arXiv:2403.07691*.
- Hosseini, A.; Yuan, X.; Malkin, N.; Courville, A. C.; Sordoni, A.; and Agarwal, R. 2024. V-STaR: Training Verifiers for Self-Taught Reasoners. *CoRR*, abs/2402.06457.
- Jiao, F.; Qin, C.; Liu, Z.; Chen, N. F.; and Joty, S. 2024. Learning Planning-based Reasoning by Trajectories Collection and Process Reward Synthesizing. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 334–350. Miami, Florida, USA: Association for Computational Linguistics.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213.
- Li, W.; and Li, Y. 2024. Process reward model with q-value rankings. *arXiv preprint arXiv:2410.11287*.
- Lightman, H.; Kosaraju, V.; Burda, Y.; Edwards, H.; Baker, B.; Lee, T.; Leike, J.; Schulman, J.; Sutskever, I.; and Cobbe, K. 2024. Let’s Verify Step by Step. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.
- Liu, Z.; Chen, C.; Li, W.; Qi, P.; Pang, T.; Du, C.; Lee, W. S.; and Lin, M. 2025a. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*.
- Liu, Z.; Wang, P.; Xu, R.; Ma, S.; Ruan, C.; Li, P.; Liu, Y.; and Wu, Y. 2025b. Inference-time scaling for generalist reward modeling. *arXiv preprint arXiv:2504.02495*.
- Luo, L.; Liu, Y.; Liu, R.; Phatale, S.; Guo, M.; Lara, H.; Li, Y.; Shu, L.; Zhu, Y.; Meng, L.; et al. 2024. Improve mathematical reasoning in language models by automated process supervision. *arXiv preprint arXiv:2406.06592*.
- Magister, L. C.; Mallinson, J.; Adámek, J.; Malmi, E.; and Severyn, A. 2023. Teaching Small Language Models to Reason. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 1773–1781. Association for Computational Linguistics.

- Pang, R. Y.; Yuan, W.; He, H.; Cho, K.; Sukhbaatar, S.; and Weston, J. 2024. Iterative reasoning preference optimization. *Advances in Neural Information Processing Systems*, 37: 116617–116637.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shen, T.; Cambria, E.; Wang, J.; Cai, Y.; and Zhang, X. 2025a. Insight at the right spot: Provide decisive subgraph information to Graph LLM with reinforcement learning. *Information Fusion*, 117: 102860.
- Shen, T.; Mao, R.; Wang, J.; Zhang, X.; and Cambria, E. 2025b. Flow-guided direct preference optimization for knowledge graph reasoning with trees. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1165–1175.
- Team, G.; Riviere, M.; Pathak, S.; Sessa, P. G.; Hardin, C.; Bhupatiraju, S.; Hussenot, L.; Mesnard, T.; Shahriari, B.; Ramé, A.; et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Uesato, J.; Kushman, N.; Kumar, R.; Song, F.; Siegel, N.; Wang, L.; Creswell, A.; Irving, G.; and Higgins, I. 2022. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*.
- Wang, P.; Li, L.; Shao, Z.; Xu, R.; Dai, D.; Li, Y.; Chen, D.; Wu, Y.; and Sui, Z. 2024. Math-Shepherd: Verify and Reinforce LLMs Step-by-step without Human Annotations. In Ku, L.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, 9426–9439. Association for Computational Linguistics.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2024. Qwen2.5 Technical Report. *CoRR*, abs/2412.15115.
- Yeung, N.; Botvinick, M. M.; and Cohen, J. D. 2004. The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychological review*, 111(4): 931.
- Yu, Q.; Zhang, Z.; Zhu, R.; Yuan, Y.; Zuo, X.; Yue, Y.; Fan, T.; Liu, G.; Liu, L.; Liu, X.; et al. 2025. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Yuan, L.; Cai, Y.; Shen, X.; Li, Q.; Huang, Q.; Deng, Z.; and Wang, T. 2025. Collaborative Multi-LoRA Experts with Achievement-based Multi-Tasks Loss for Unified Multimodal Information Extraction. *arXiv preprint arXiv:2505.06303*.
- Yuan, L.; Cai, Y.; Wang, J.; and Li, Q. 2023a. Joint multimodal entity-relation extraction based on edge-enhanced graph alignment network and word-pair relation tagging. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 11051–11059.
- Yuan, Z.; Yuan, H.; Li, C.; Dong, G.; Lu, K.; Tan, C.; Zhou, C.; and Zhou, J. 2023b. Scaling relationship on learning mathematical reasoning with large language models. *arXiv preprint arXiv:2308.01825*.
- Zelikman, E.; Wu, Y.; Mu, J.; and Goodman, N. 2022. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35: 15476–15488.
- Zhang, Z.; Zheng, C.; Wu, Y.; Zhang, B.; Lin, R.; Yu, B.; Liu, D.; Zhou, J.; and Lin, J. 2025. The lessons of developing process reward models in mathematical reasoning. *arXiv preprint arXiv:2501.07301*.
- Zheng, C.; Zhang, Z.; Zhang, B.; Lin, R.; Lu, K.; Yu, B.; Liu, D.; Zhou, J.; and Lin, J. 2024a. Processbench: Identifying process errors in mathematical reasoning. *arXiv preprint arXiv:2412.06559*.
- Zheng, G.; Kong, J.; Wang, J.; and Zhang, X. 2025. Enhanced Multimodal Chain-of-Thought with Visual Self-Contrastive Distillation. In *2025 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6.
- Zheng, G.; Wang, J.; Zhou, X.; and Zhang, X. 2024b. Enhancing Semantics in Multimodal Chain of Thought via Soft Negative Sampling. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 6059–6076.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36: 46595–46623.