# **REVISE:** LEARNING TO REFINE AT TEST-TIME VIA INTRINSIC SELF-VERIFICATION

**Hyunseok Lee<sup>1\*</sup> Seunghyuk Oh<sup>1\*</sup> Jaehyung Kim<sup>2</sup> Jinwoo Shin<sup>1</sup> Jihoon Tack<sup>1</sup>** <sup>1</sup>KAIST, <sup>2</sup>Yonsei University

### Abstract

Self-awareness, i.e., the ability to assess and correct one's generation, is a fundamental aspect of human intelligence, making its replication in large language models (LLMs) an important yet challenging task. Previous works tackle this by employing extensive reinforcement learning or relying on large external verifiers. In this work, we propose Refine via Intrinsic Self-Verification (ReVISE), an efficient and effective framework that enables LLMs to self-correct their outputs through self-verification. The core idea of ReVISE is to enable LLMs to verify their reasoning processes and continually rethink reasoning trajectories based on its verification. To implement this efficiently, we introduce a structured curriculum based on preference learning. Specifically, as ReVISE involves two challenging tasks (i.e., self-verification and reasoning correction), we tackle each task sequentially using curriculum learning, collecting both failed and successful reasoning paths to construct preference pairs for efficient training. During inference, our approach enjoys natural test-time scaling by integrating self-verification and correction capabilities, further enhanced by our proposed confidence-aware decoding mechanism. Our experiments on various reasoning tasks demonstrate that ReVISE achieves efficient self-correction and significantly improves the reasoning performance of LLMs.

### 1 INTRODUCTION

Large language models (LLMs) have demonstrated remarkable success across diverse domains, such as coding assistants (Zhang et al., 2024b), search engines (Xiong et al., 2024), and personal AI assistants (Sajja et al., 2024), progressively advancing toward human-like logical reasoning capabilities (Amirizaniani et al., 2024). However, tasks requiring rigorous System 2 thinking—such as complex reasoning (Jaech et al., 2024), iterative trial-and-error (Song et al., 2024), and dynamic planning (Xie & Zou, 2024)—remain highly challenging (Lowe, 2024; Cai et al., 2024). A key difficulty in LLM reasoning is that errors in early steps can accumulate over time, leading to substantial inaccuracies (LeCun, 2022), while the models' intrinsic ability to detect and rectify such self-generated errors—often framed as a form of self-awareness—remains insufficient. This issue is further exacerbated by the autoregressive nature of LLMs, which constrains their ability to revisit and revise prior steps (Bachmann & Nagarajan, 2024).

To tackle this issue, recent approaches have emphasized verification (or correction) of LLM-generated reasoning trajectories as a crucial mechanism (Zhang et al., 2024a; Madaan et al., 2023). For instance, some methods utilize external large-scale verifiers to iteratively validate outputs and trigger regeneration (Luo et al., 2024). However, the reliance on expensive external models introduces computational inefficiencies. Alternatively, reinforcement learning (RL)-based techniques have shown promise in improving reasoning accuracy by optimizing reward signals based on ground-truth correctness, enabling self-correction (Kumar et al., 2024). However, RL is a complex and often unstable procedure (Mnih et al., 2015; Rafailov et al., 2023), and it does not explicitly model the verification of intermediate reasoning steps, making it difficult to assess whether a model is confident in its current trajectory or prone to deviating toward incorrect conclusions, which may limit interpretability and adaptability in complex reasoning tasks.

<sup>\*</sup>equal contribution

This raises a key question: Can LLMs be equipped with an internal mechanism to explicitly verify their own reasoning and correct potential errors based on their verification?

**Contribution.** We propose **Re**fine Via Intrinsic **SE**If-Verification (ReVISE), a novel and effective self-correction framework for LLM reasoning using self-verification. The core idea of ReVISE is to enable LLMs to assess their reasoning process and refine reasoning trajectories based on self-verification. Specifically, we introduce a special token, which outputs whether to stop the generation or revise the reasoning trajectory. To train the model to utilize this token effectively, we design a two-stage curriculum to simplify the learning of two challenging tasks—self-verification and self-correction—by breaking them into separate training stages. Here, both stages employ preference learning, allowing the model to learn these tasks efficiently without heavy computational overhead. In the first stage, we collect pairs of correct and incorrect reasoning trajectories (i.e., positive and negative samples for preference learning) based on output correctness to develop the model's self-verification ability. In the second stage, we generate new preference pairs for self-correction by constructing positive samples where a correct reasoning path follows an incorrect one, and negative samples where an incorrect reasoning path follows a correct one.

Furthermore, we introduce an inference-time scaling strategy for ReVISE that leverages self-verification to enhance performance. First, as ReVISE inherently verifies and refines reasoning paths when it detects incorrect outputs, it naturally benefits from increased test-time computation. Additionally, we propose a novel test-time sampling scheme that incorporates self-verification confidence (i.e., the confidence in deciding whether to terminate generation). Specifically, we integrate this confidence into existing test-time sampling methods by adjusting the sampling score based on the predicted confidence, leading to more reliable output.

We demonstrated the effectiveness of ReVISE through evaluations on multiple reasoning datasets in the mathematical domain. Notably, ReVISE enhances reasoning performance beyond prior methods, improving accuracy from  $27.1 \rightarrow 31.1\%$  on GSM8K (Maj@3) (Cobbe et al., 2021) with Llama3 1B (Dubey et al., 2024) and from  $33.2 \rightarrow 36.0\%$  on MATH (Maj@3) (Hendrycks et al., 2021) with Llama3 8B. Furthermore, our experimental results show that ReVISE consistently improves accuracy without relying on external feedback mechanisms, which often degrade performance on complex reasoning tasks. For instance, unlike approaches such as Refine (Madaan et al., 2023), which struggle when combined with existing models on complex tasks, ReVISE achieves these gains purely through self-verification and self-correction. Finally, we show that the proposed sampling scheme is more efficient than other sampling strategies when applied to models trained with ReVISE, further enhancing the performance.

# 2 RELATED WORK

In this section, we provide a comprehensive review of related works, including reasoning, test-time scaling, and self-improvement for LLMs.

**LLM reasoning.** LLMs have made significant progress in reasoning through techniques such as Chain-of-Thought (CoT) prompting, fine-tuning, and self-improvement. CoT prompting, introduced by (Wei et al., 2022) and expanded by (Kojima et al., 2022) enables models to break down complex problems into intermediate steps, improving performance and interpretability. Structured reasoning methods, including self-consistency (Wang et al., 2022) and Tree-of-Thought (ToT) (Yao et al., 2024), enhance multi-step problem-solving by exploring various reasoning paths. Huang et al. (2022) have demonstrated self-improvement through iterative feedback, refining their outputs over time. Ensuring the reliability of reasoning approaches such as Reflexion (Shinn et al., 2024) and Self-Refine (Madaan et al., 2023) introduce iterative feedback loops, while verification techniques like step-by-step validation (Lightman et al., 2023) help maintain consistency and reduce errors. Unlike prior approaches, ReVISE learns self-verification during training, reducing train-test discrepancy and enabling more natural verification at inference.

**Test-time scaling for LLMs.** Recent works explored that scaling test-time computation, such as best-of-N sampling, can be even better than scaling train-time computation for performance (Snell et al., 2024). Specifically, test-time scaling strategies improve LLM performance by generating numerous candidate outputs and selecting the best. To enhance decision-making, external verifiers are often employed to evaluate and refine these outputs (Liang et al., 2024). Moreover, Kumar et al.



Figure 1: **Overview of ReVISE. Left:** ReVISE is a self-verifying and self-correcting reasoning framework. It first generates an initial answer, verifies its correctness, and decides whether to stop or refine. If the model generates the [refine] token, it refines the initial reasoning. **Right:** The structured curriculum-based training pipeline of ReVISE. In the first stage, the model learns self-verification by selecting between [eos] and [refine]. In the second stage, it learns to correct reasoning mistakes using golden data.

(2024); Qu et al. (2024) applied extensive reinforcement learning to overcome the efficiencies and dependence on the verifier's performance. In safety research, backtracking methods have introduced reset tokens to correct unsafe responses (Zhang et al., 2024c). While they focus on reducing the likelihood of unsafe outputs with limited second attempts to refuse answers, our approach targets complex reasoning tasks enabled by self-correction through an explicit verification process and two-stage curricula.

**Self-improvement for LLMs.** Self-training methods enable LLMs to refine themselves using their own outputs. Supervised fine-tuning (SFT) (Brown et al., 2020b) trains on human-annotated data but lacks self-correction (Huang et al., 2023). Rejection fine-tuning (RFT) (Yuan et al., 2023) improves robustness by filtering low-quality responses but discards useful learning signals. STaR (Zelikman et al., 2022) iteratively fine-tunes models on self-generated solutions but struggles with compounding errors due to the absence of explicit verification. V-STaR (Hosseini et al., 2024) extends STaR by jointly training a verifier alongside the generator, leveraging both correct and incorrect responses to improve self-assessment, though it still depends on large-scale self-generated data. However, discovering high-quality solutions remains a challenge, as (Luong et al., 2024) shows that RL-based fine-tuning is ineffective without supervised initialization. Kim et al. (2024) argue that LLMs struggle with self-correction. Our approach integrates both generation and verification, leveraging correct and incorrect responses for more effective self-improvement.

## 3 LEARNING TO REFINE AT TEST-TIME VIA INTRINSIC SELF-VERIFICATION

In this section, we present Refine via Intrinsic Self-Verification (ReVISE), an LLM reasoning framework that self-verifies and refines the reasoning trajectory based on the verification. We first introduce the problem of interest and a special token coined [refine], which is used for refining the LLM's generation (in Section 3.1). Then, we present the core training method, namely the two-stage curricula (in Section 3.2) and the test-time inference strategy (in Section 3.3). The overview of ReVISE is depicted in Figure 1.

#### 3.1 PROBLEM SETUP: LEARNING TO VERIFY AND REFINE

We describe the problem setup of our interest, i.e., self-verification and refinement. Given an input x, the initial output  $y_{\text{init}}$  is sampled from the LLM  $\mathcal{M}$ , i.e.,  $y_{\text{init}} \sim \mathcal{M}(\cdot|x)$ , where the reasoning path is included in  $y_{\text{init}}$ . The goal is to train an LLM that verifies the correctness of  $y_{\text{init}}$  and decides whether to terminate generation or continue generating by refining its reasoning. To this end, we introduce a special token [refine] that determines whether to proceed with refinement. Specifically, given  $y_{\text{init}}$ , the model verifies its correctness by predicting  $v \sim \mathcal{M}(\cdot|y_{\text{init}}, x)$ , where  $v \in \{[eos], [refine]\}$ , allowing it to either terminate generation by predicting [eos] or continue generating by refining its reasoning by outputting [refine]. If refinement is needed, the model generates a revised response  $y_{\text{refined}} \sim \mathcal{M}(\cdot|[\text{refine}], y_{\text{init}}, x)$ , completing the correction cycle. Note that this modeling has distinct advantages as one can access the model's verification confidence of v.

### 3.2 REVISE: REFINE VIA INTRINSIC SELF-VERIFICATION

We first describe our core training pipeline of ReVISE, namely the structured curriculum based on online preference learning. As ReVISE involves two challenging tasks (i.e., self-verification and refinement), we propose two-stage curricula. In the first stage, we train the LLM to intrinsically self-verify its generation by predicting the [eos] or [refine] tokens. Then, at the second stage, we continually train this LLM to correct the generation when the output reasoning is wrong. For efficient and stable training, we employ preference optimization (i.e., learning from preference-based positive and negative pairs) based on our proposed preference data collection strategy. This allows us to perform structured preference learning without relying on reinforcement learning (RL), which can be computationally extensive and unstable (Rafailov et al., 2023).

Stage 1: Learning to verify self-generations. Given an initial LLM  $\mathcal{M}_0$  and a supervised fine-tuning dataset  $\mathcal{D} = \{(x_i, y_i)\}_i$  consisting of input-label pairs (including reasoning traces), our goal is to construct preference pairs for training  $\mathcal{M}_0$ . Specifically, for each input x, we generate a positive output  $y^+$  and a negative output  $y^-$ . To achieve this, we first sample multiple responses from  $\mathcal{M}_0$ . This allows us to obtain both correct reasoning outputs  $y_{\text{correct}}$  and incorrect ones  $y_{\text{wrong}}$ , which are identified using the ground-truth answer y. Using these outputs, we construct a preference dataset by distinguishing two cases: (i) when the model generates the correct answer  $y_{\text{correct}}$ , predicting [eos] is preferred over [refine], and (ii) vice versa for incorrect output  $y_{\text{wrong}}$ , we define the preference triplets  $(x, y^+, y^-)$  as:

$$\begin{cases} (x, \hat{y} \oplus [\texttt{eos}], \hat{y} \oplus [\texttt{refine}]), & \text{if } \hat{y} = y_{\texttt{correct}} \\ (x \oplus \hat{y}, [\texttt{refine}], [\texttt{eos}]), & \text{if } \hat{y} = y_{\texttt{wrong}} \end{cases}$$

where  $\oplus$  is the concatenation operator. Based on the proposed collection strategy, we generate a preference dataset  $\mathcal{D}_{verify}$  for training the intrinsic verification of the LLM.

To this end, we jointly optimize the supervised fine-tuning loss with the direct preference optimization (DPO; Rafailov et al., 2023) loss. Specifically, for a given preference dataset D, the SFT and DPO preference losses are defined as:

$$\begin{split} \mathcal{L}_{\text{SFT}}(\mathcal{D}) &:= -\mathbb{E}_{(x,y^+)\sim\mathcal{D}}\log\mathcal{M}(y^+|x)\\ \mathcal{L}_{\text{Pref}}(\mathcal{D}) &:= -\mathbb{E}_{(x,y^+,y^-)\sim\mathcal{D}}\Big[\sigma\big(r(x,y^+) - r(x,y^-)\big)\Big]\\ \text{where} \quad r(x,y) = \beta\log\frac{\mathcal{M}(y\mid x)}{\mathcal{M}_0(y\mid x)}, \end{split}$$

where  $\beta \in \mathbb{R}^+$  is hyper-parameter controlling proximity to the base model  $\mathcal{M}_0$  and  $\sigma$  is the logistic function. It is worth noting that SFT loss only focuses on minimizing the negative log-likelihood of the positive output. Then, our training objective for self-verification is as:

$$\mathcal{L}_{\text{verify}} := \mathcal{L}_{\text{SFT}}(\mathcal{D}_{\text{verify}}) + \lambda \, \mathcal{L}_{\text{Pref}}(\mathcal{D}_{\text{verify}}) \tag{1}$$

where  $\lambda \in \mathbb{R}^+$  is a loss balancing hyperparameter. Here, we denote the initial model  $\mathcal{M}_0$  trained with  $\mathcal{L}_{verify}$  as  $\mathcal{M}_1$ , which is the output model of first curricula.

Table 1: Accuracy (%) for ReVISE (Ours) and other baselines, including Few-shot CoT, SFT, RFT, STAR<sup>+</sup> trained models. We consider two math reasoning benchmarks, GSM8K (Cobbe et al., 2021) and MATH-500 (Lightman et al., 2023). MATH-500 is a subset of the original MATH benchmark (Hendrycks et al., 2021). Maj@K indicates that majority voting for K samples, exceptionally ReVISE used own weighted majority voting system as described in Section 3.3. The **bold** indicated the best result within the group.

	Llama-3.2-1B				Llama-3.1-8B				
	GSM8K			MATH-500			MATH-500		
Methods	Maj@1	Maj@3	Maj@5	Maj@1	Maj@3	Maj@5	Maj@1	Maj@3	Maj@5
Few-shot CoT	5.7	6.8	7.2	3.0	2.6	3.2	23.4	22.2	23.2
SFT (Brown et al., 2020a)	22.1	23.5	26.4	10.4	10.6	11.4	27.8	31.0	33.2
RFT (Yuan et al., 2023)	26.2	26.8	28.6	12.6	12.4	12.8	30.8	33.2	35.6
STaR <sup>+</sup> (Zelikman et al., 2022)	26.2	27.1	29.9	11.4	13.1	13.4	30.4	31.8	32.8
ReVISE (Ours)	28.1	31.3	32.8	13.4	14.0	14.8	33.6	36.0	37.6

Stage 2: Learning to correct self-generations. We now describe how to train ReVISE to acquire another core ability: self-correction. Similar to self-verification, we perform preference learning using the same loss function. To this end, we aim to construct a new preference dataset, denoted as  $\mathcal{D}_{correct}$ . The core idea consists of two main components. First, the curriculum learning: we utilize outputs generated by the model  $\mathcal{M}_1$  and initialize stage 2 training from  $\mathcal{M}_1$ . Second, to learn how to correct incorrect outputs, we repurpose the wrong reasoning paths  $y_{wrong}$  used in stage 1 to construct the dataset.

Concretely, we consider two possible cases: whether the initial response is correct  $y_{\text{correct}}$  or incorrect  $y_{\text{wrong}}$ . If the initial response is correct  $y_{\text{correct}}$ , we construct preference data as same as stage 1, i.e., discouraging the generation of [refine] and encouraging [eos]. The key case is when the initial response is incorrect  $y_{\text{wrong}}$ . In this case, we need to have a positive preference sample that refines the incorrect reasoning  $y_{\text{wrong}}$  with the correct reasoning. To achieve this, we concatenate the ground-truth label y to the response. Formally, the preference pairs are defined as:

$$\begin{cases} (x, \hat{y} \oplus [\texttt{eos}], \hat{y} \oplus [\texttt{refine}]), & \text{if } \hat{y} = y_{\texttt{correct}} \\ (x \oplus \hat{y}, [\texttt{refine}] \oplus y, [\texttt{eos}]), & \text{if } \hat{y} = y_{\texttt{wrong}} \end{cases}$$

where y is the ground-truth label. Using the self-correction preference dataset  $\mathcal{D}_{correct}$ , we train the final model  $\mathcal{M}_2$  from  $\mathcal{M}_1$  with the following correction loss:

$$\mathcal{L}_{\text{correct}} := \mathcal{L}_{\text{SFT}}(\mathcal{D}_{\text{correct}}) + \lambda \, \mathcal{L}_{\text{Pref}}(\mathcal{D}_{\text{correct}}). \tag{2}$$

It is worth noting that stage 2 explicitly defines when and how refinements should be applied, preventing overgeneration and improving response accuracy. By distinguishing between necessary and unnecessary refinements, the model ensures efficient self-correction while simulating multi-step reasoning for complex scenarios.

Furthermore, our dataset collection strategy shares similarities with recent backtracking methods in that incorrect initial generations are utilized to create negative pairs (Zhang et al., 2024c). We also observe that leveraging past failure trajectories aids in ultimately achieving successful reasoning. In this regard, we believe that applying ReVISE to safety-critical applications, akin to backtracking, is an interesting future direction, where our proposed curriculum learning and explicit self-verification stage can contribute to developing safer models.

#### 3.3 VERIFICATION CONFIDENCE-AWARE SAMPLING

We propose an inference method for models trained with ReVISE. The key idea is to calibrate the standard sampling-based scoring approach using the self-verification confidence. Specifically, we apply this method to majority voting, where N samples are generated, and the most frequent prediction is selected. Unlike conventional approaches, our method explicitly accesses the self-verification confidence, as our model not only generates an answer but also determines its correctness by producing either an [eos] or [refine] token. This allows us to directly obtain the probability associated with self-verification, enabling confidence-weighted aggregation for more reliable predictions.

Concretely, given an input x, we generate N candidate answers  $\mathcal{Y} = \{y_1, y_2, \dots, y_N\}$  from the LLM at stage 2, denoted as  $\mathcal{M}$  for simplicity, where each  $y_i$  is sampled as  $y_i \sim \mathcal{M}(\cdot|x)$ . To refine the selection process, we leverage the softmax probability of the verification (i.e., the probability of [eos] token), denoted as follows:

$$c_i = \mathcal{M}([\mathsf{eos}]|y_i, x),$$

as a confidence score. Instead of selecting the most frequent prediction, we accumulate these scores by summing the confidence values of identical answers, leading to the final prediction as follows:

$$y^* = \arg \max_{y \in \mathcal{Y}} \sum_{i: y_i = y} c_i.$$

This approach calibrates the traditional majority voting method by weighting predictions based on their model-derived confidence, showing effective scaling at test time.

### 4 EXPERIMENTS

We provide an empirical evaluation of ReVISE by investigating the following questions:

- Can ReVISE enhance reasoning performance? (Table 1)
- Can ReVISE improve performance further through test-time compute scaling? (Figures 5, 8)
- Can ReVISE do self-verify and self-refine? (Figures 3, 7)

**Training setup.** For the main experiment, we train ReVISE on Llama-3 models with 1B and 8B parameters, which are not instruction-tuned. We avoid using instruction-tuned models to prevent potential bias from exposure to the gold data of the tasks (Wang et al., 2024). For this reason, the models were first supervised fine-tuned using the labeled dataset, followed by fine-tuning with each respective method. For GSM8K (Cobbe et al., 2021), we train ReVISE using the original training split. For MATH (Hendrycks et al., 2021), we train ReVISE using a 50k subset of MetaMath (Yu et al., 2024), an augmented version of MATH, and use a 3k subset for the validation set, respectively. Here, MetaMath was employed to mitigate the performance degradation caused by the limited size of the original MATH.

**Baselines.** We compare our method against several baseline approaches: Supervised Fine-Tuning (SFT), RFT(Yuan et al., 2023), and STaR<sup>+</sup>. In RFT, fine-tuning is performed on supervised fine-tuning data  $\mathcal{D}$  and correctly generated samples selected from k completions for each input in the training set by a tuned model. Like RFT, STaR (Zelikman et al., 2022) trains on correctly generated samples, including self-generated rationales given a hint (rationalization). However, unlike RFT, STaR iteratively refines this process without relying on  $\mathcal{D}$ . Since both ReVISE and RFT utilize ground truth data  $\mathcal{D}$ , we introduce an extended version of STaR that incorporates SFT data as a baseline, referred to as STaR<sup>+</sup>. Essentially, STaR<sup>+</sup> functions as a multi-iteration variant of RFT with rationalization. We run STaR<sup>+</sup> for three iterations, sampling k completions per iteration (GSM8K: k = 10, MATH: k = 4, GSM240K: k = 1) with a temperature of 0.7 for both RFT and STaR<sup>+</sup>. We initialize a model for each iteration from  $\mathcal{M}_0$  that is supervised fine-tuned with  $\mathcal{D}$  at each iteration for STaR<sup>+</sup> to prevent overfitting.

**Evaluation setup.** We mainly report Majority Voting at N (Maj@N) as a sampling-based metric, exceptionally ReVISE used verification confidence-aware majority voting as described in Section 3.3 (unless otherwise specified). We evaluate ReVISE and baselines on GSM8K (Cobbe et al., 2021) and MATH-500 (Hendrycks et al., 2021), a widely used evaluation benchmark subset of MATH.

#### 4.1 MAIN RESULTS

We first present the main result by comparing the math problem-solving performance with other baselines. Here, we mainly compare with various fine-tuning schemes and show that ReVISE can be jointly used with a test-time scaling method to improve the performance further.

As shown in Table 1, we present the reasoning and test-time scaling performance of ReVISE compared to reasoning-enhancing fine-tuning baselines. Overall, ReVISE significantly and consistently outperforms all prior fine-tuning methods. It is worth noting that for both GSM8K and MATH-500, ReVISE achieves the highest Maj@1, indicating that ReVISE is already strong without the proposed



Figure 2: Test-time scaling comparison between ReVISE (Ours) and baselines, including SFT, RFT, STAR<sup>+</sup>, and majority voting for ReVISE (Ours (Simple Maj.)) at sampling sizes  $N \in \{1, 2, 3, 4, 8\}$ . (a) Results for Llama-3.2-1B on the GSM8K dataset. (b) Results for Llama-3.2-8B on the MATH dataset. ReVISE consistently outperforms baselines across all sample sizes and datasets.

sampling scheme. For instance, ReVISE attains 33.6% for Maj@1, significantly outperforming SFT (30.4%) and few-shot CoT (23.4%) on MATH-500 with Llama-3.1-8B. In addition, with the proposed confidence-aware majority voting, ReVISE marked a 4.0% gain after refinement and consistently outperforms other baselines, under 5 sampled answers. These results demonstrate that ReVISE enhances problem-solving accuracy and improves self-sampling abilities.

### 4.2 INFERENCE SCALABILITY OF REVISE

In this section, we evaluate the inference scalability of ReVISE. To this end, we visualize how the test-time scaling improves as one samples more candidates. Specifically, we conduct experiments using our method with different sample sizes  $N \in \{2, 3, 4, 8\}$  and compare with results of other baselines using majority voting. As shown in Figure 2, ReVISE achieves significant and consistent gain in all setups. For instance, ReVISE shows a large gap with the strongest baseline RFT, showing 3.3% of improvement in MATH-500 at N = 8. Furthermore, our method even benefits under limited number of samples (N = 2), while majority voting does not show improvement. This is because majority voting does not use confidence and, hence, can not benefit from small samples (e.g., if all predictions are disjoint, the majority voting does not work). Finally, ReVISE shows scalable improvements in all model configurations, ranging from relatively small 1B models to large 8B models. Notably, ReVISE achieves a significant performance gain in the 8B model, suggesting strong generalization capabilities.

#### 4.3 ADDITIONAL ANALYSIS AND ABLATION

In this section, we provide a detailed analysis of ReVISE to validate the effect of each proposed component. Unless otherwise specified, we use a Llama-3.2-1B trained on GSM8K across all methods throughout this section.

**Effectiveness of curriculum learning.** We validate the effectiveness of the proposed curriculum learning (in Figure 3). To this end, we train two types of models. First, the model trains without curriculum by optimizing SFT  $\mathcal{L}_{SFT}$  and preference  $\mathcal{L}_{Pref}$  loss by using all preference dataset at once, i.e.,  $\mathcal{D}_{correct}$ . Second, we train the model only using the first verification loss, i.e.,  $\mathcal{D}_{verify}$  (note that self-verification already enables the model to generate the answer but does not know how to correct the generation). As shown in Figure 3a, the curriculum is indeed showing a significant improvement over no curriculum baseline (even the model has used the same preference dataset); two-stage curricula improve the performance from 22.6% to 28.1%.

To further investigate this phenomenon, we evaluate the self-verification accuracy of each method, which measures the model's ability to predict whether its own output is correct. In Figure 3b, we report the verification accuracy in terms of the Area Under the Receiver Operating Characteristic Curve (AUROC) for three models. Notably, the model without curriculum learning achieves an



Figure 3: Ablation study on curriculum learning in the aspect of (a) accuracy (%) and (b) AUROC (%). The experiments are conducted using Llama-3.1-1B on the GSM8K dataset. The comparison includes a model trained without curriculum learning (w/o Cur.), trained for only stage 1 (Stage 1), and trained using the full two-stage curriculum learning approach (ReVISE) (Stage 2). (a) Accuracy improves with curriculum learning by mitigating conflicts between competing objectives during early training stages. (b) AUROC results demonstrate enhanced classification performance of corrected and incorrect responses and effective transfer from Stage 1 to the final ReVISE model.

Table 2: Results on the GSM8K benchmark for ReVISE and baselines trained on Llama-3.2-1B Instruct. Except for ReVISE, all methods underperform compared to the zero-shot CoT baselines.

Methods	GSM8K	GSM240K
Zero-shot CoT	48.6	48.6
SFT	41.9	54.8
RFT	44.0	50.9
<b>ReVISE (Ours)</b>	52.3	59.4

AUROC of 71%, while two-stage curriculum learning improves this to 76%. This suggests that curriculum learning enhances self-verification, allowing the model to refine its predictions based on more reliable verification signals. However, we observe that training at stage 2 slightly degrades verification accuracy, indicating that the self-correction task  $\mathcal{D}_{correct}$  is particularly challenging and may lead to catastrophic forgetting (McCloskey & Cohen, 1989). Exploring optimization strategies that improve self-verification and self-correction without compromising overall performance remains an interesting direction for future work.

Effectiveness of preference learning. The role of DPO loss in ReVISE is to guide the model to prefer refining when the initial attempt is incorrect and terminating otherwise. Additionally, in our DPO objective, we applied SFT loss to the chosen sequence as introduced in Liu et al. (2024) which applied SFT loss to the selected sequence,  $\mathcal{L}_{Ours} := \mathcal{L}_{SFT}(\mathcal{D}) + \lambda \mathcal{L}_{Pref}(\mathcal{D})$ , where  $\lambda$  is a constant. For ReVISE, such SFT has two roles: to encourage the model to generate correct answers and learn special tokens. At the beginning of the training, the model had zero-initialized knowledge about special tokens, so SFT accelerated the model's use of special tokens.



Figure 4: Ablation study on DPO loss, evaluated on the GSM8K benchmark. Removing DPO loss significantly reduces accuracy.

Specifically, ablation experiments without the DPO loss—where only the SFT loss is utilized—in Figure 4 show that ReVISE without DPO demonstrates significantly lower performance -10.3% compared to the full-trained ReVISE. This indicates that the DPO loss is critical in ReVISE for effectively guiding the refinement process.

**ReVISE on instruction-tuned models.** While we have primally focused on pretrained model and initialized  $\mathcal{M}_0$  with the given supervised fine-tuning dataset  $\mathcal{D}$  due to the possible data contamination, we also have conducted on experiment on Llama-3.2-1B-Instruct, an instruction-tuned model.



Figure 5: Inference-time scaling comparison between ReVISE ( $\mathcal{M}([eos])$  (Ours)) and inference metrics. For  $\mathcal{M}(Answer)$  and  $\mathcal{M}([eos])$  (Ours), we have done weighted majority voting. ReVISE consistently outperforms other inference metrics, and ReVISE's sampling method using weighted majority voting exceeds the performance of majority voting. Experiments are conducted using the Llama-3.2-1B model.

Interestingly, as shown in Table 2, all fine-tuning methods, except for ReVISE, underperform the zero-shot CoT baseline when training with GSM8K. This outcome aligns with the widely recognized challenge that fine-tuning instruction-tuned models often leads to catastrophic forgetting, hindering their ability to learn new information effectively. Meanwhile, ReVISE remains notably resistant to this issue. We hypothesize that this advantage stems from how ReVISE utilizes the gold label y—only incorporating it as a revised second-attempt completion rather than directly fine-tuning it. In contrast, baselines such as SFT, RFT, and STaR<sup>+</sup> rely on fine-tuning the base model on D, which becomes problematic when the target model's performance is already strong, as it struggles to gain further improvements from D.

To mitigate forgetting, we also trained the model using GSM240K, a subset of MetaMath dataset (Yu et al., 2024), which expands the original data about 30-fold by rephrasing questions and answers. As shown in Table 2, while training with GSM240K improved the performance of the SFT baseline, ReVISE still exhibited better performance. This result suggests that ReVISE can adapt to various data characteristics, even in heavily augmented settings.

Analysis on the self-verification confidence of **ReVISE.** We further analyze the confidence distribution in self-verification to assess whether the model's confidence is well aligned with actual correctness. To this end, we visualize the probability gap between [eos] and [refine], simply defined as  $\mathcal{M}([eos]) - \mathcal{M}([refine])$ . As shown in Figure 6, incorrect responses tend to have lower [eos] probabilities, whereas correct responses exhibit higher [eos] probabilities. This demonstrates the model's intrinsic ability to assess its own correctness. Moreover, these results suggest that confidence serves as a reliable metric for calibrating the sampling score, further validating the effectiveness of our confidence-aware sampling method.



Figure 6: Distribution histrogram of  $\mathcal{M}([\texttt{eos}]) - \mathcal{M}([\texttt{refine}])$ . 0 is the threshold of ReVISE trigger intrinsically refine or not. Experiments are conducted using the Llama-3.2-1B model.

Ablation study on confidence-aware sampling. We explore the impact of different score calibrations during inference by leveraging ReVISE's self-verification mechanism to enable test-time compute-scalable inference strategies (see 3.3). Specifically, we compare three scoring schemes for selecting N-generated samples: (1) weighted majority voting using  $\mathcal{M}([eos]|x)$ , (2) unweighted majority voting, and (3) scoring based on the model's predicted answer likelihood. These calibration methods govern both the selection of candidate answers and the evaluation of their validity.

As shown in Figure 5,  $\mathcal{M}([eos]|x)$ -based (Ours) score consistently outperforms alternatives across GSM8K benchmarks. For example, with eight sampled candidates,  $\mathcal{M}([eos]|x)$ -based scoring achieves an accuracy of 33.9%, compared to 33.2% (unweighted majority), and 32.7% (likelihood-



Figure 7: Analysis of refinement capability of ReVISE. We compare the accuracy (%) on GSM8K and MATH when using different decoding approaches. *First* stops at [refine], *Retry* re-generates responses, while ReVISE refines its initial reasoning. The results show that ReVISE improves accuracy, demonstrating its ability to refine rather than randomly re-generate responses. Experiments are conducted using the Llama-3.2-1B model.

based). The trend persists across all tested sampling budgets, suggesting strong compatibility with self-verification mechanisms. This consistent advantage implies  $\mathcal{M}([eos]|x)$  better aligns with the model's intrinsic verification capability to distinguish correct reasoning paths. We carefully hypothesize that  $\mathcal{M}([eos]|x)$  acts as a latent indicator of solution correctness, as premature termination often correlates with reasoning errors.

Analysis on the refinement. We demonstrate that ReVISE refines its answers based on the initial attempt rather than randomly generating a new completion. To evaluate this, we compare ReVISE with two baselines: *First* and *Retry*. *First* terminates decoding at the [refine] token, while *Retry* generates a new completion upon encountering [refine]. Specifically, *Retry* greedily decodes the first attempt, and if [refine] appears, it samples a new completion with a temperature of 0.7 following the prompt x. In contrast, both *First* and ReVISE greedily generate completions. As shown in Figure 7, ReVISE outperforms both *First* and *Retry*. This result highlights that ReVISE does not generate new responses arbitrarily but instead meaningfully refines and improves upon its initial answer.

### 5 DISCUSSION AND CONCLUSION

In this paper, we introduce Refine via Intrinsic Self-Verification (ReVISE), a novel framework that enables Large Language Models (LLMs) to perform self-verification and self-correction during inference. Through a structured curriculum learning approach, we demonstrated how LLMs can progressively learn to verify their reasoning and improve their outputs. Our results across various reasoning benchmarks show that ReVISE significantly improves first-attempt accuracy while maintaining efficiency. Furthermore, the self-verification mechanism and a confidence-aware decoding strategy enhance model performance without introducing additional computational overhead.

# REFERENCES

Maryam Amirizaniani, Elias Martin, Maryna Sivachenko, Afra Mashhadi, and Chirag Shah. Do llms exhibit human-like reasoning? evaluating theory of mind in llms for open-ended responses. *arXiv* preprint arXiv:2406.05659, 2024.

Gregor Bachmann and Vaishnavh Nagarajan. The pitfalls of next-token prediction. 2024.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. 2020a.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. 2020b.

- Huanqia Cai, Yijun Yang, and Zhifeng Li. System-2 mathematical reasoning via enriched instruction tuning. *arXiv preprint arXiv:2412.16964*, 2024.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. arXiv preprint arXiv:2110.14168, 2021.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. 2021.
- Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron Courville, Alessandro Sordoni, and Rishabh Agarwal. V-star: Training verifiers for self-taught reasoners. *arXiv preprint arXiv:*, 2024.
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*, 2022.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*, 2023.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. 2024.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. arXiv preprint arXiv:2412.16720, 2024.
- Sehoon Kim, Karttikeya Mangalam, Suhong Moon, Jitendra Malik, Michael W Mahoney, Amir Gholami, and Kurt Keutzer. Speculative decoding with big little decoder. 2024.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213, 2022.
- Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M Zhang, Kay McKinney, Disha Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, and Aleksandra Faust. Training language models to self-correct via reinforcement learning. arXiv preprint arXiv:, 2024.
- Yann LeCun. A path towards autonomous machine intelligence. Open Review, 2022.
- Zhenwen Liang, Ye Liu, Tong Niu, Xiangliang Zhang, Yingbo Zhou, and Semih Yavuz. Improving llm reasoning through scaling inference computation with collaborative verification. *arXiv preprint arXiv:2410.05318*, 2024.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *arXiv preprint arXiv:2305.20050*, 2023.
- Zhihan Liu, Miao Lu, Shenao Zhang, Boyi Liu, Hongyi Guo, Yingxiang Yang, Jose Blanchet, and Zhaoran Wang. Provably mitigating overoptimization in rlhf: Your sft loss is implicitly an adversarial regularizer. *arXiv preprint arXiv:2405.16436*, 2024.
- Scott C Lowe. System 2 reasoning capabilities are nigh. In *The First Workshop on System-2 Reasoning at Scale, NeurIPS*'24, 2024.
- Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, Lei Meng, Jiao Sun, et al. Improve mathematical reasoning in language models by automated process supervision. *arXiv preprint arXiv:2406.06592*, 2024.

- Trung Quoc Luong, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. Reft: Reasoning with reinforced fine-tuning. *arXiv preprint arXiv:2401.08967*, 2024.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. 2023.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *The Psychology of Learning and Motivation*, 1989.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518, 2015.
- Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching language model agents how to self-improve. *arXiv preprint arXiv*:, 2024.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. 2023.
- Ramteja Sajja, Yusuf Sermet, Muhammed Cikmaz, David Cwiertny, and Ibrahim Demir. Artificial intelligence-enabled intelligent assistant for personalized and adaptive learning in higher education. *Information*, 15(10):596, 2024.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
- Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. Trial and error: Exploration-based trajectory optimization for llm agents. *arXiv preprint arXiv:2403.02502*, 2024.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2024.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- Chengxing Xie and Difan Zou. A human-like reasoning framework for multi-phases planning task with large language models. *arXiv preprint arXiv:2405.18208*, 2024.
- Haoyi Xiong, Jiang Bian, Yuchen Li, Xuhong Li, Mengnan Du, Shuaiqiang Wang, Dawei Yin, and Sumi Helal. When search engine services meet large language models: visions and challenges. *IEEE Transactions on Services Computing*, 2024.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Longhui Yu, Weisen Jiang, Han Shi, YU Jincheng, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. 2024.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Keming Lu, Chuanqi Tan, Chang Zhou, and Jingren Zhou. Scaling relationship on learning mathematical reasoning with large language models. *arXiv preprint arXiv:2308.01825*, 2023.

- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. Star: Bootstrapping reasoning with reasoning. 2022.
- Di Zhang, Xiaoshui Huang, Dongzhan Zhou, Yuqiang Li, and Wanli Ouyang. Accessing gpt-4 level mathematical olympiad solutions via monte carlo tree self-refine with llama-3 8b. *arXiv preprint arXiv:2406.07394*, 2024a.
- Kechi Zhang, Jia Li, Ge Li, Xianjie Shi, and Zhi Jin. Codeagent: Enhancing code generation with tool-integrated agent systems for real-world repo-level coding challenges. *arXiv preprint arXiv:2401.07339*, 2024b.
- Yiming Zhang, Jianfeng Chi, Hailey Nguyen, Kartikeya Upasani, Daniel M Bikel, Jason Weston, and Eric Michael Smith. Backtracking improves generation safety. *arXiv preprint arXiv:2409.14586*, 2024c.

# A EXPERIMENTAL DETAILS

In this section, we describe the experimental details of Section 4, including ReVISE and the baselines.

**Dataset details.** In this section, we describe the dataset we used in training and evaluation. Also, explain how we generated the additional datasets.

- Grade School Math 8K (GSM8K). The GSM8K dataset (Cobbe et al., 2021) consists of 8,790 high-quality grade-school math word problems. We used the provided train and test splits, ensuring consistency across all experiments. The dataset serves as a benchmark for evaluating the arithmetic and reasoning capabilities of language models.
- **MATH.** The MATH dataset (Hendrycks et al., 2021) is a challenging collection of problems from high school mathematics competitions, covering diverse topics such as algebra, geometry, calculus, and statistics. We utilized the original train and test splits, which include approximately 12,500 problems. Due to the dataset's complexity, it effectively evaluates the model's ability to handle higher-level mathematical reasoning.
- **MetaMath.** MetaMath (Yu et al., 2024) is an augmented version of the MATH and GSM dataset, designed to address the challenges posed by the limited size of the original dataset. We selected a 50k subset of MetaMath for training and sampled 3k problems for the validation set. MetaMath includes additional examples generated using synthetic data augmentation techniques, such as problem paraphrasing and structural variations, to enhance diversity and improve generalization. This augmentation mitigates performance degradation associated with small datasets while maintaining the original problem difficulty and format.

**Training details of ReVISE** We use AdamW optimizer with a learning rate  $lr \in \{10^{-4}, 10^{-5}\}$  with 10% warm up and cosine decay and train it for one epoch. We trained with batch size 32 for fine-tuning and 64 for preference tuning . For the  $\lambda$  constant for SFT loss, we used  $\lambda = 0.1$ . During the training, for the data sampling phase, we sampled 10 times for each sample in GSM8K and 4 times for each sample in MATH.

**Training model details.** We mainly used the open-source Large Language Models (LLMs) from Llama-family. Specifically we used meta-llama/Llama-3.2-1B, meta-llama/Llama-3.1-8B which are not instruction-tuned and meta-llama/Llama-3.2-1B-Instruct, which is instruction-tuned. We used the model checkpoint from huggingface library.

**Evaluation details.** Used lm-eval-harness for greedy decoding experiments and used our code to evaluate models in sampling settings. Since the output depends on the evaluation batch size, we fixed the batch size to 128 for a fair comparison.

- **GSM8K.** Used the test split as a benchmark dataset.
- **MATH-500.** The MATH-500 dataset is a curated collection of 500 MATH dataset. For our experiments, we used Math-500 exclusively for evaluation purposes.

**Resource details.** For the main development we mainly use Intel(R) Xeon(R) Platinum 8480+ CPU @ 790MHz and a 8 NVIDIA H100 GPUs. Additionally, we used NVIDIA RTX4090 GPUs for evaluation.

### **Baseline details.**

- SFT We fine-tuned the model using a language modeling loss, exploring learning rates from  $1e^{-6}$  to  $1e^{-4}$ , with epochs ranging from 1 to 3 and a batch size of 32.
- **RFT** We sampled ten completions for GSM8K, one for GSM240K, and four for MATH-50K. The model was trained for one epoch on the collected dataset with a fixed learning rate of  $1e^{-5}$ .
- **STaR**<sup>+</sup> We sampled the same number of samples as in **RFT**. The outer loop was fixed to 3 for all datasets, with one epoch per outer loop. Rationalization was performed with a hint, where the answer was provided except for the rationale, which served as the hint. The learning rate was fixed at  $1e^{-5}$ .

# **B** ADDITIONAL RESULTS

### B.1 ITERATIVE REFINING SEQUENTIALLY AT TEST TIME.



(a) Llama-3.2-8B fine-tuned in MATH (b) Llama-3.2-1B-instruct evaluated at and evaluated at MATH-500 GSM8K

Figure 8: Accuracy improvements through iterative refinement. The plot shows the accuracy (%) at GSM8K and MATH-500 of ReVISE across multiple rounds of iterative refinement (1, 2, and 3 tries).

Although ReVISE is trained to refine its output in a single, we explore its potential for iterative refinement. Specifically, after generating the second attempt, we append it to the original prompt x and treat it as the first-attempt output. This allows the model to either output [eos] to terminate the sequence or generate a third attempt following the same process, effectively enabling multiple rounds of refinement. As shown in Figure 8, in MATH-500, ReVISE trained on Llama-3.1-8B's accuracy consistently improves as the model iteratively refines its responses up to 3 times.

This observation suggests the potential for training a model explicitly designed for multi-round iterative refinement to enhance the iterative refinement even more. By incorporating iterative refinement directly into the training process, the model could learn to self-correct more effectively across multiple rounds. We leave this direction as an exciting avenue for future work.

## B.2 REVISE RESULTS EXAMPLE

### B.2.1 REFINEMENT VIA INTRINSIC SELF-VERIFICATION EXAMPLE

### GSM8K

GSM8K test sample	ReVISE trained on Llama-3.2-1B
Question	Mrs. Tatiana owns a grocery store that sells different fruits and vegetables, which includes carrots. The price of carrots in the grocery store increases by 5% of the original price every year. What would be the price of carrots after three years if it was \$120 initially? (Round to the nearest integer) Let's think step by step. Put your final answer at the end with 'The answer is: .'
Initial Response	The price of carrots increases by $5/100 * $120 = $6$ . So the price of carrots after one year is $120+6 = $126$ . The price of carrots after two years is $126+6 = $132$ . The answer is: 132
Verify	[refine]
Refined Response	The price of carrots increases by $5/100*$120 = $6$ per year. The price after one year is $$120 + $6 = $126$ The price after two years is $$126 + $6 = $132$ . The price after three years is $$132 + $6 = $138$ . The answer is: 138
Target Answer	138

# MATH

MATH test sample	ReVISE trained on Llama-3.2-1B
Question	Ten treeks weigh as much as three squigs and one goolee. Two treeks and one goolee are equal in weight to one squig. The combined weight of how many treeks equals the weight of one squig? Let's think step by step. Put your final answer at the end with 'The answer is: .'
Initial Response	We can set up a system of equations to solve this problem. Let t be the weight of one treek, s be the weight of one squig, and g be the weight of one goolee. From the given information, we can write the following equations: $10t = 3s + g \ 2t + g = s$ We can solve this system of equations to find the weight of one treek in terms of the weight of one squig. Multiplying the second equation by 3, we get: $6t + 3g = 3s$ Subtracting the first equation from this equation, we get: $6t + 3g - 10t = 3s - (3s + g)$ Simplifying, we get: $-4t + 3g = -g$ Adding $4t$ to both sides, we get: $3g = 4t$ Dividing both sides by 3, we get: $g = \frac{4}{3}t$ Substituting this into the second equation, we get: $2t + \frac{4}{3}t = s$ Multiplying both sides by 3, we get: $10t = 3s$ Since we know that $10t = 3s + g$ , we can substitute $g = \frac{4}{3}t$ into this equation: $10t = 3s + \frac{4}{3}t$ Multiplying both sides by 3, we get: $30t = 9s + 4t$ Subtracting $4t$ from both sides, we get: $26t = 9s$ Dividing both sides by 9, we get: $\frac{26}{9}t = s$ So, the weight of one treek is equal to $\frac{9}{26}$ times the weight of one squig. Therefore, the combined weight of how many treeks equals the weight of one squig is $\frac{9}{26}$ .
Verify	[refine]
Refined Response	Let t, s, and g represent the weights of one treek, one squig, and one goolee, respectively. From the given information, we have the following equations: $10t = 3s + g$ (Equation 1) $2t + g = s$ (Equation 2) We want to find the weight of one treek in terms of the weight of one squig, so we need to solve for t in terms of s. We can start by solving Equation 2 for g in terms of t and s: $g = s - 2t$ Substituting this into Equation 1, we have: $10t = 3s + (s - 2t)$ Simplifying, we get: $10t = 4s - 2t$ Combining like terms, we have: $12t = 4s$ Dividing both sides by 4, we get: $3t = s$ Therefore, the weight of one treek is equal to $\frac{1}{3}$ of the weight of one squig. The answer is: $3$
Target Answer	3

# $B.2.2 \quad STOP \ \text{VIA INTRINSIC SELF-VERIFICATION EXAMPLE}$

# GSM8K

GSM8K test sample	ReVISE trained on Llama-3.2-1B
Question	Jen is planning to sell her root crops. She has 6 yams which can be sold at \$1.5 each, 10 sweet potatoes that cost \$2 each, and 4 carrots which cost \$1.25 each. If she sells everything, how much will she earn?
Initial Response	Jen earns $6 \ge 1.5 = 9$ for the yams. She earns $10 \ge 2 = 20$ for the sweet potatoes. She earns $4 \ge 1.25 = 5$ for the carrots. Therefore, Jen earns $9 + 20 + 5 = 34$ in all. The answer is: $34$
Verify	[eos]
Target Answer	34

### MATH

MATH test sample	ReVISE trained on Llama-3.2-1B
Question	Evaluate $(1 + 2i)6 - 3i$ . Let's think step by step. Put your final answer at the end with 'The answer is:.'
Initial Response	(1+2i)6 - 3i = 6 + 12i - 3i = 6 + 9i = 6 + 9i. The answer is:
	6+9i.
Verify	[eos]
Target Answer	6+9i

### C THE FEW-SHOT PROMPT USED FOR EVALUATION

Few-shot prompt for GSM8K:

```
Given the following problem, reason and give a final answer to the

→ problem.

Problem: There are 15 trees in the grove. Grove workers will plant trees
\hookrightarrow in the grove today. After they are done, there will be 21 trees. How
\rightarrow many trees did the grove workers plant today?
Your response should end with "The final answer is [answer]" where
\hookrightarrow [answer] is the response to the problem.
There are 15 trees originally. Then there were 21 trees after some more
 \rightarrow were planted. So there must have been 21 - 15 = 6. The final answer
 \hookrightarrow is 6
Given the following problem, reason and give a final answer to the
\rightarrow problem.
Problem: If there are 3 cars in the parking lot and 2 more cars arrive,
\rightarrow how many cars are in the parking lot?
Your response should end with "The final answer is [answer]" where
\hookrightarrow [answer] is the response to the problem.
There are originally 3 cars. 2 more cars arrive. 3 + 2 = 5. The final
 \hookrightarrow answer is 5
Given the following problem, reason and give a final answer to the
\rightarrow problem.
Problem: Leah had 32 chocolates and her sister had 42. If they ate 35,
\rightarrow how many pieces do they have left in total?
Your response should end with "The final answer is [answer]" where
\leftrightarrow [answer] is the response to the problem.
Originally, Leah had 32 chocolates. Her sister had 42. So in total they
 \rightarrow had 32 + 42 = 74. After eating 35, they had 74 - 35 = 39. The final
 \hookrightarrow answer is 39
Given the following problem, reason and give a final answer to the
\rightarrow problem.
Problem: Jason had 20 lollipops. He gave Denny some lollipops. Now Jason
\, \hookrightarrow \, has 12 lollipops. How many lollipops did Jason give to Denny?
Your response should end with "The final answer is [answer]" where
\leftrightarrow [answer] is the response to the problem.
Jason started with 20 lollipops. Then he had 12 after giving some to
 \rightarrow Denny. So he gave Denny 20 - 12 = 8. The final answer is 8
Given the following problem, reason and give a final answer to the
\rightarrow problem.
Problem: Shawn has five toys. For Christmas, he got two toys each from
\leftrightarrow his mom and dad. How many toys does he have now?
```

Your response should end with "The final answer is [answer]" where  $\hookrightarrow$  [answer] is the response to the problem. Shawn started with 5 toys. If he got 2 toys each from his mom and dad,  $\rightarrow$  then that is 4 more toys. 5 + 4 = 9. The final answer is 9 Given the following problem, reason and give a final answer to the  $\hookrightarrow$  problem. Problem: There were nine computers in the server room. Five more ↔ computers were installed each day, from monday to thursday. How many computers are now in the server room? Your response should end with "The final answer is [answer]" where  $\hookrightarrow$  [answer] is the response to the problem. There were originally 9 computers. For each of 4 days, 5 more computers  $\rightarrow$  were added. So 5 \* 4 = 20 computers were added. 9 + 20 is 29. The  $\hookrightarrow$  final answer is 29 Given the following problem, reason and give a final answer to the  $\rightarrow$  problem. Problem: Michael had 58 golf balls. On tuesday, he lost 23 golf balls.  $\hookrightarrow$  On wednesday, he lost 2 more. How many golf balls did he have at the  $\rightarrow$  end of wednesday? Your response should end with "The final answer is [answer]" where  $\hookrightarrow$  [answer] is the response to the problem. Michael started with 58 golf balls. After losing 23 on tuesday, he had  $\rightarrow$  58 - 23 = 35. After losing 2 more, he had 35 - 2 = 33 golf balls.  $\rightarrow$  The final answer is 33 Given the following problem, reason and give a final answer to the  $\rightarrow$  problem. Problem: Olivia has \$23. She bought five bagels for \$3 each. How much  $\rightarrow$  money does she have left? Your response should end with "The final answer is [answer]" where  $\hookrightarrow$  [answer] is the response to the problem. Olivia had 23 dollars. 5 bagels for 3 dollars each will be 5 x 3 = 15  $\leftrightarrow$  dollars. So she has 23 - 15 dollars left. 23 - 15 is 8. The final  $\rightarrow$  answer is 8 Given the following problem, reason and give a final answer to the  $\rightarrow$  problem. Problem: Janet's ducks lay 16 eggs per day. She eats three for breakfast  $\hookrightarrow$  every morning and bakes muffins for her friends every day with four.  $\rightarrow$  She sells the remainder at the farmers' market daily for \$2 per  $\, \hookrightarrow \,$  fresh duck egg. How much in dollars does she make every day at the  $\hookrightarrow$  farmers' market? Your response should end with "The final answer is [answer]" where  $\leftrightarrow$  [answer] is the response to the problem.

#### Few-shot prompt for MATH-500:

```
Problem: Terrell usually lifts two 20-pound weights 12 times. If he uses
\, \hookrightarrow \, two 15-pound weights instead, how many times must Terrell lift them
\rightarrow in order to lift the same total weight?
Answer: If Terrell lifts two 20-pound weights 12 times, he lifts a total
\rightarrow of (2 \cdot 12 \cdot 20 = 480 \) pounds of weight. If he lifts two
\hookrightarrow 15-pound weights instead for \( n \) times, he will lift a total of
\label{eq:constraint} \begin{array}{l} \hookrightarrow & \ (2 \ dot \ 15 \ dot \ n \ = \ 30n \ ) \ \text{pounds of weight. Equating this to} \\ \leftrightarrow & \ 480 \ \text{pounds, we can solve for } (n \ ) \ (30n \ = \ 480 \ ) \ \text{Rightarrow} \end{array}
\hookrightarrow \quad n = 480 / 30 = \boxed{16} \] The answer is \boxed{16}.
Problem: If the system of equations [ 6x - 4y = a, ]  by - 9x = b, ]
\, \hookrightarrow \, has a solution \( (x, y) \) where \( x \) and \( y \) are both
\rightarrow nonzero, find \( \frac{a}{b} \), assuming \( b \) is nonzero.
Answer: If we multiply the first equation by ( -\frac{3}{2} ), we
\rightarrow obtain \[ 6y - 9x = -\frac{3}{2}a. \] Since we also know that \( 6y
\rightarrow - 9x = b \), we have \[ -\frac{3}{2}a = b \Rightarrow \frac{a}{b} =
\rightarrow \boxed{-\frac{2}{3}}. \] The answer is \boxed{-\frac{2}{3}}.
Problem: If 2^{8}=4^{x}, what is the value of x?
Answer:
```