

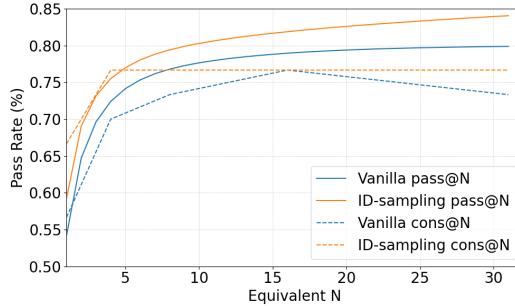
EFFICIENT TEST-TIME SCALING VIA ITERATIVE DEEP-ENING

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

Recent reasoning models, such as OpenAI’s O1 series, have demonstrated exceptional performance on complex reasoning tasks and revealed new test-time scaling laws. Inspired by this, many people have been studying how to train models to achieve effective self-evaluation and self-correction to further enable the scaling paradigm. However, less studied is how to efficiently scale test-time compute from a fixed model, and this remains a challenge. In this paper, we focus on whether LLMs can benefit from matching the pattern of correct responses. Specifically, we explore how systematically triggering a model’s self-correction mechanisms can improve performance on challenging reasoning tasks. To this end, we propose a novel iterative deepening sampling algorithm framework designed to enhance self-correction and generate higher-quality samples. Through extensive experiments on Math500, AIME, and GPQA-diamond benchmarks, we demonstrate that our method achieves a higher success rate on difficult tasks and provide detailed ablation studies to analyze its effectiveness across diverse settings.



037 Figure 1: Comparison of our proposed ID-sampling to vanilla sampling on Qwen3-8B on AIME-
038 25. The x-axis is the equivalent number of samples after considering the additional cost from ID-
039 sampling.

1 INTRODUCTION

044 Since ChatGPT, large language models (LLMs) have been a rapidly evolving domain that tries to
045 solve problems beyond traditional language tasks like summarization or question answering Chen
046 et al. (2023); Yao et al. (2023); Chen et al. (2024b;d). Significantly, the newly released OpenAI
047 O1 has introduced its new paradigm of test-time scaling, which shows strong capability in complex
048 problem-solving through its detailed reasoning steps before outputting the final answer Jaech et al.
049 (2024). Since then, many researchers have studied how to replicate success from an open-source
050 perspective and how to train models that are even better at efficiently solving problems that still
051 remain unsolvable by the current LLMs Huang et al. (2024); Zeng et al. (2024); DeepSeek-AI et al.
052 (2025). One key finding is that through reinforcement learning itself, LLM can spontaneously learn
053 to self-evaluate and self-correct from time to time. However, there is no clear conclusion on whether
self-evaluation is triggered often enough.

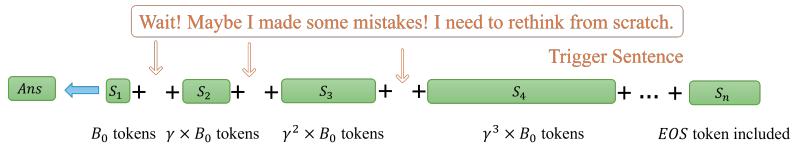


Figure 2: An illustration of our method, Iterative Deepening Sampling (ID-sampling).

On the other hand, while training is the primary focus recently, it remains uncertain whether one could more efficiently scale test-time compute from a fixed model without additional changes in training or fine-tuning. Moreover, a more efficient sampling algorithm can not only enhance inference-time efficiency but also facilitate the generation of high-quality synthetic data, which can be leveraged to train the next generation of models and evaluate their performance Guan et al. (2025).

Recent work has investigated the differing patterns of correct versus incorrect responses. Notably, while o1-like models tend to achieve higher performance when provided with greater computational resources Huang et al. (2024); DeepSeek-AI et al. (2025), other studies have found that longer responses can actually reduce accuracy, a phenomenon often referred to as overthinking Marjanović et al. (2025). Given these distinct response patterns within the same model, a fundamental question arises: can we improve the test-time performance of LLMs by guiding generation to mimic the patterns of correct responses?

In this paper, we tackle the challenge of efficient test-time scaling using a fixed reasoning model without additional training, applied to complex reasoning tasks such as mathematical problem solving. Specifically, inspired by a novel observation on the positional distribution difference of linguistic markers that appear in correct and incorrect responses, we focus on manually inserting self-reflection triggers like "Wait" during the generation process to improve the pass rate of a fixed model. To achieve this, we propose Iterative Deepening Sampling (ID-Sampling), a novel algorithmic framework that iteratively increases the sampling budget following a geometric progression, while incorporating self-reflection mechanisms at each expansion step. We theoretically demonstrate that ID-Sampling effectively balances computational efficiency and response quality, ensuring that the budget is not excessively wasted while still improving model performance. We evaluate ID-Sampling on a few challenging reasoning benchmarks, and demonstrate its effectiveness in Best-of-N sampling and majority voting settings on various reasoning models. Additionally, we provide an ablation study analyzing how the rate of budget increase per iteration impacts both pass rate and inference time. Our results highlight the potential of ID-Sampling as a scalable approach for improving LLM reasoning performance through adaptive self-reflection mechanisms.

2 RELATED WORKS

There are two primary directions for test-time scaling, namely scaling on multiple responses, and scaling the reasoning steps in a single response.

To efficiently scale to multiple responses, researchers have proposed both aggregating independent samples and employing structured methods such as tree-search-based architectures. While there have been many works on tree-search strategies early on Zhang et al. (2023); Liu et al. (2023); Hao et al. (2023); Chen et al. (2024a); Zhou et al. (2023); Zhang et al. (2024a;b), researchers are putting more attention on aggregating independent responses given that it is hard to create a good process reward model to accurately estimate the quality of a partial solution DeepSeek-AI et al. (2025). Specifically in this direction, researchers mostly rely on two main classes of aggregating strategies, namely self-consistency Wang et al. (2022) and Best-of-N Snell et al. (2024); Wu et al. (2025). There are many works that focus on this setting to make it more efficient and effective. For example, Sun et al. Sun et al. (2024) proposed to use speculative rejection in BoN to reject bad responses through early scores. Chen et al. Chen et al. (2024e) proposed to use extremely high temperature on the first token to greatly improve the Best-of-N performance on math and coding tasks.

To scale within a single response, researchers have first introduced an intermediate step by the so-called chain-of-thoughts Wang & Zhou (2024). Then, starting with Self-Refine Madaan et al. (2023), there are many works that studied how to effectively use an LLM to give themselves their

108 own evaluation and reflection Yao et al. (2023); Shinn et al. (2023), and study how this reflection-
 109 based mechanism can be adapted to different applications Chen et al. (2023); Gou et al. (2023);
 110 Chen et al. (2024c). While there have previously been a lot of discussions on whether LLM can
 111 actually self-evaluate and self-correct themselves Huang et al. (2023); Chen et al. (2024f); Verma
 112 et al. (2024); Wang et al. (2025b), recent studies have shown that by training with high-quality data
 113 that involve such a process, LLM can actually achieve a useable level of capability and help self-
 114 correct in the process Huang et al. (2024); Zeng et al. (2024). Inspired by the recent success of
 115 Deepseek-R1 DeepSeek-AI et al. (2025), people have realized that scaling within a single-response
 116 can be efficiently achieved through reinforcement learning, and many follow-up works have also
 117 studied how to learn such scaling capability more efficiently Wang et al. (2025a). In this paper, we
 118 manually inject *trigger sentence* in the middle of generation, which is similar to the recent work s1
 119 Muennighoff et al. (2025). However, our work focuses on injecting at an early stage in the thinking
 120 process, while they focus on injecting when the current response is completely finished for budget
 121 forcing. As we later show in our experiments, this design difference leads to a substantial divergence
 122 in results: their approach requires further training, whereas ours provides benefits directly with
 123 existing models.

124 While the research in the two search strategies is mostly separate, their methods are orthogonal and
 125 can be used together to make the sampling process more efficient Snell et al. (2024), and some re-
 126 cent research has also studied how to balance the two, and has shown that scaling on the number
 127 of responses can be more efficient than scaling within a single response Wang et al. (2025a); Mar-
 128 janović et al. (2025); Sadhukhan et al. (2025). In this paper, we focus on the intersection of both
 129 scaling strategies and study how to strategically trigger the self-correction capability of LLMs more
 130 efficiently. We take into account the additional cost of scaling within one response by measuring the
 131 total wall-clock time used, and calculate the equivalent N when comparing the answers as needed.

132 3 PRELIMINARIES

134 3.1 BEST-OF-N SAMPLING

136 For reasoning-intensive tasks such as mathematical problem-solving and coding, Best-of-N sam-
 137 pling is one of the most widely used strategies for data generation. This approach involves sampling
 138 N outputs from the same model using predefined sampling parameters — typically with a higher
 139 temperature than single-sample settings — followed by a selection process to determine the best re-
 140 sponse. The selection criteria depend on the intended use of the samples. During training, responses
 141 are typically evaluated using rule-based checkers for mathematical problems or online judges for
 142 coding tasks to identify correct answers within the sampled set. At test time, a reward model is often
 143 employed to score the generated responses, with the highest-scoring output selected as the final an-
 144 swer. This methodology effectively balances exploration and optimization, making it a fundamental
 145 component in enhancing the performance of LLMs on reasoning tasks.

146 The BoN sampling is a simple yet effective method that can be fully parallelized to enhance per-
 147 formance. Increasing N guarantees improved results during training when a ground-truth checker is
 148 available and generally leads to better performance at inference time, provided that the reward model
 149 is sufficiently accurate. However, if paired with a reward model that is not good enough, BoN might
 150 fail to scale efficiently and might even decrease its performance when more samples are included for
 151 aggregation. Therefore, getting a good reward model is necessary for good performance for BoN.

152 3.2 MAJORITY VOTING

154 Similar to BoN sampling, majority voting, also known as self-consistency (cons@n), provides an
 155 alternative approach for aggregating N different responses Wang et al. (2022). As the name sug-
 156 gests, this method involves generating N responses, counting the frequency of each unique answer,
 157 and selecting the most frequently occurring response as the final output. This approach leverages
 158 the inherent redundancy in multiple generations to improve robustness and reliability, making it
 159 particularly useful for tasks requiring high confidence in correctness.

160 Majority voting offers the advantage of aggregating responses without relying on a reward model.
 161 Additionally, although empirically less effective, it can be extended to a weighted version, where
 162 weights are assigned based on PRM scores or confidence estimations of the generated answers Wang

162
163
164
165
166
167
168
169
170
171

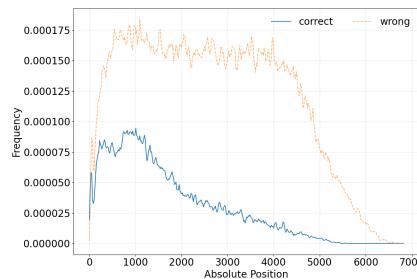


Figure 3: Positional frequency of thinking-related linguistic markers in Deepseek-R1 responses on the AIME-2024 dataset. A higher frequency indicates the likelihood of a given position containing such markers within a response. The plot is smoothed with a moving average for better visualization.

et al. (2022). While majority voting benefits from not requiring a highly accurate reward model, it faces challenges in identifying equivalent answers. For instance, it can be very hard to identify two code samples to be the same, and even in mathematical problems, expressions such as $\frac{1}{\sqrt{3}}$ and $\frac{\sqrt{3}}{3}$ are equivalent but must be recognized as such to ensure correct vote counting. A common solution in mathematical domains involves using symbolic-based checkers to compare answer pairs and identify equivalences. However, this process can be computationally expensive, requiring up to $O(N^2)$ comparisons.

4 ITERATIVE DEEPENING SAMPLING

4.1 OBSERVATION: FREQUENCIES OF LINGUISTIC MARKERS IN REASONING MODELS

Recent studies have highlighted that overthinking can negatively impact the performance of reasoning models Wang et al. (2025a); Marjanović et al. (2025). A key observation is that correct answers tend to be shorter and contain fewer linguistic markers associated with "thinking" behaviors.

Here, we present a novel perspective by analyzing not only the frequency but also the positional distribution of thinking-related markers within model responses. A complete list of these markers is provided in the appendix. As illustrated in Fig. 3, correct answers tend to exhibit fewer think markers, with a clear concentration in the early stages of the response. Their frequency drops sharply after the first 1000 tokens. In contrast, incorrect answers often contain self-corrections distributed throughout the response, with a noticeable decline only near the end. While this distinct distribution pattern is less informative than response length for classifying correctness due to its complexity, it is nonetheless a characteristic that we can mimic during generation.

4.2 METHODS

In this paper, we focus on efficiently scaling test-time computation by replacing the standard sampling algorithm with a more effective, custom-designed alternative. Given a fixed overall computational budget, an important challenge is determining how much additional budget should be allocated to refining a given prefix x_0 that the model has already sampled. Efficient budget allocation is crucial, as any saved resources can be redirected to increasing the number of N in Best-of-N sampling or deepening tree search, ultimately improving overall performance. Understanding this trade-off is key to optimizing both search efficiency and model output quality.

More specifically, the self-evaluation and self-correction process of the LLM can be manually triggered by introducing a predefined *trigger sentence*, such as "Wait! Maybe I made some mistakes! I need to rethink from scratch." or simply "Wait". In most cases, LLMs respond to this trigger by restarting their reasoning process and making self-corrections. Repeatedly inserting this sentence increases the overall length of the reasoning trajectory, potentially improving problem-solving accuracy by facilitating iterative refinement. In this paper, we propose a method for strategically placing a fixed *trigger sentence* at increasing intervals within longer contextual windows, aiming to balance computational efficiency and the effectiveness of self-correction.

216 **Algorithm 1** Iterative Deepening Sampling

217 1: **function** GENERATE_WITH_ID-SAMPLING(LLM, Dataset, B_0 , γ , B):
218 2: *Prefixes* \leftarrow *Questions from Dataset*
219 3: *Budget* $\leftarrow B_0$
220 4: **while** *Prefixes* is not empty **and** *Budget* $\leq B$ **do**
221 5: *NewPrefixes* $\leftarrow \{\}$
222 6: *Outputs* \leftarrow LLM.generate(*input* = *Prefixes*, *max_tokens* = *Budget*)
223 7: **for** each *Output* in *Outputs* **do**
224 8: **if** *Output* is finished **then**
225 9: LOGANSWER(*Output*)
226 10: **else if** *Output* finished thinking **then**
227 11: *Output* \leftarrow LLM.generate(*input* = *Output*)
228 12: LOGANSWER(*Output*)
229 13: **else**
230 14: *NewPrefix* \leftarrow PADTRIGGERSentence(*Output*)
231 15: *NewPrefixes*.append(*NewPrefix*)
232 16: **end if**
233 17: **end for**
234 18: *Budget* \leftarrow *Budget* $\times \gamma$
235 19: *Prefixes* \leftarrow *NewPrefixes*
236 20: **end while**
237 21: **end function**

238
239
240
241
242

243 Motivated by the analysis in Section 4.1, we propose to introduce more *trigger sentences* at earlier
244 stages and gradually reduce their occurrence as the response length increases. Suppose we have
245 already used a budget of b to generate a prefix x_0 and are now considering whether to immediately
246 introduce a *trigger sentence*. Iterative Deepening (ID) sampling allocates an additional budget of $\gamma \times$
247 b before inserting the *trigger sentence* the next time, where $\gamma > 1$ is a tunable hyperparameter. This
248 iterative process continues until reaching a maximum budget B , beyond which no further *trigger*
249 sentences are introduced. If the reasoning process reaches a natural stopping point—i.e., a complete
250 answer is generated within the allocated budget—the process terminates early. This is because
251 generating a full response from scratch generally leads to more reliable outputs than attempting to
252 refine an already complete solution. The complete procedure is outlined in Algorithm 1, where B_0
253 represents the initial budget. The function LLM.generate conducts the generation within a given
254 budget and is adaptable to different tree-search strategies. The function PadTriggerSentence handles
255 the insertion of *trigger sentences* while ensuring redundancy is minimized if necessary.

256 The definition and allocation of computational budget depend on the specific test-time scaling al-
257 gorithm employed, leading to variations in implementation strategies. In this paper, we focus on
258 the setting where N responses are sampled independently, and we have provided an illustration of
259 ID-sampling in Fig. 2. Since response generation is independent and typically performed in par-
260 allel, the computational cost primarily depends on the total number of generated responses N and
261 their respective lengths. To manage computational efficiency, we define the budget as the maximum
262 number of tokens generated in a given round, which corresponds to the *max_token* parameter in
263 LLM serving engines such as vLLM Kwon et al. (2023). Additionally, to avoid inserting the *trigger*
264 *sentence* mid-sentence, we extend generation until the completion of a reasoning step. Here, a step
265 is identified by token splits such as ‘\n’ or ‘\n\n’. This ensures that *trigger sentence* placement
266 aligns with the basic logical structure of the generated response, preserving coherence and stability
267 in the reasoning process. On the other hand, this also helps introduce some randomness in inserting
268 *trigger sentence*, avoiding inserting them on exactly the same token at each run.

269 In implementations, the maximum generation token limit is a common parameter provided by mod-
270 ern serving engines such as vLLM or OpenAI’s API services. Consequently, ID-sampling can still
271 efficiently leverage the speedups offered by these engines. The pseudocode is provided in Alg. 1.

270 4.3 THEORETICAL ANALYSIS
271

272 A key challenge in ID-sampling is that budget control occurs before each manually triggered self-
273 evaluation and self-correction step without explicitly analyzing the actual generated responses. This
274 can lead to unnecessary iterations, potentially increasing computational costs. However, due to the
275 design of our ID-sampling method, we establish important theoretical guarantees. In particular,
276 we provide a bound on the total number of tokens generated before reaching the final answer, as
277 formalized in the following theorem.

278 **Theorem 4.1.** *Suppose the final answer obtained through ID-Sampling needs a budget of L in
279 normal sampling without manual injection in the middle. Then the total number of additional budget
280 used is no more than $\frac{\gamma*L}{\gamma-1}$.*

282 **Proof Sketch** Note that the budget for each generation iteration follows a geometric sequence
283 with common ratio γ . The theorem follows directly from a summation of this geometric series. A
284 complete proof is provided in the appendix.

285 The theorem does not guarantee the quality of the generated responses or the alignment of the output
286 distribution. Rather, it establishes that our method incurs no substantial additional computational
287 overhead when generating a single response. Our intuition, however, is that by introducing *trigger
288 sentences* at an early stage, ID-sampling can yield empirical gains in pass rate by biasing the model’s
289 output distribution toward higher-quality responses.

290 It is important to note that Iterative Deepening (ID) sampling does not impose any assumptions or
291 constraints on the model’s inherent self-correction or self-evaluation capabilities. In some cases,
292 the model may naturally generate a response that already includes a *trigger sentence* before an
293 explicit manual insertion. We observe that the built-in reasoning capabilities of recent state-of-
294 the-art models, such as OpenAI-o1 Jaech et al. (2024) and DeepSeek-R1 DeepSeek-AI et al. (2025),
295 significantly impact the effectiveness of ID-sampling. To better understand these effects, we conduct
296 a comprehensive study using DeepSeek-R1, which we present in the experimental section.

297 5 EXPERIMENTS
298300 5.1 EXPERIMENT SETUP
301

302 **Models and Datasets** In our experiments, we primarily target improvements in reasoning models
303 where test-time scaling actually holds, while also reporting results for non-reasoning models as
304 a point of reference. We use Llama-3.1-1B-Instruct Dubey et al. (2024) and Phi-4 Abdin et al.
305 (2024) as examples for non-reasoning models, and DeepSeek-R1-Distill-Qwen-7BDeepSeek-AI
306 et al. (2025), Qwen3-8B and Qwen3-32B Yang et al. (2025) for reasoning models. For reward models
307 for BoN, we employ Qwen-2.5-Math-PRM-7B Zhang et al. (2025). We evaluate ID-sampling
308 on three benchmark datasets: MATH-500 Lightman et al. (2023), AIME, and GPQA-diamond Rein
309 et al. (2024). The first two are math-focused datasets, while GPQA-diamond is a science QA bench-
310 mark. Due to space constraints, results on GPQA-diamond are presented in the appendix.

311 **Baselines** We compare our approach against two settings: (i) vanilla sampling without any manual
312 intervention, and (ii) S1 Muennighoff et al. (2025) without the training mentioned, which is simply
313 appends word “wait” after the original reasoning process is completed twice.

315 **Evaluation** We report three key pass rate metrics across the datasets:

- 317 1. Best-of-N (BoN) – The accuracy when a reward model selects the best response from N
318 generated samples.
- 319 2. Pass@N - The pass rate, which measures whether or not at least one of the N total responses
320 is correct. We have used the unbiased estimation version Chen et al. (2021) with 40 total
321 samples to address the statistical significance problem.
- 322 3. Majority Voting (cons@N) – The accuracy when responses are aggregated via unweighted
323 majority voting, selecting the most frequent answer.

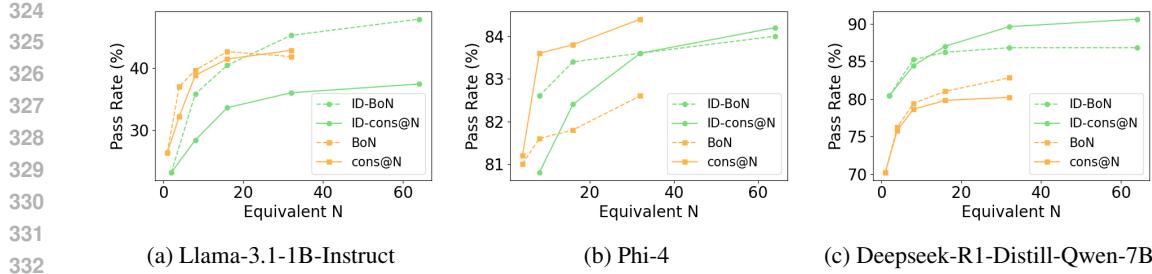


Figure 4: Math-500 dataset: Pass rate results for different models. The x-axis is the equivalent N after considering the extra time used by ID-sampling.

For both BoN and majority voting, a single aggregated answer is compared against the ground-truth solution to measure accuracy. We report BoN on Math500 datasets and for non-reasoning models on AIME datasets, given that its response length is within the context length limit of the reward models, and report Pass@n for reasoning models on AIME datasets because the normal response length is $> 12K$ and beyond the context length of popular reward models.¹

For both math benchmarks, we employ symbolic checkers to ensure that all mathematically equivalent answers are accepted as correct.

Parameters For *trigger sentence*, we choose to use a whole sentence "Wait! Maybe I made some mistakes! I need to rethink from scratch." for non-reasoning models and a single word "wait" for reasoning models. [For the word "wait", the natural generation behavior without ID sampling follows the same pattern shown in Fig. 3.](#) We will provide a discussion about this choice later. By default, we use $\gamma = 2.0$ for ID-sampling, and we will later provide ablation study results on this. For ID-sampling with BoN, we set the initial budget B_0 as 256 tokens. For reasoning models, we set the thinking budget as 4096 for any model on Math-500 datasets, and 16384 for AIME datasets. If the model fully used the thinking budget, an end-of-think token ($</\text{think}>$) is padded, and an additional 1000 tokens are provided for the final answer to be provided. For non-thinking models, we set a fixed max sequence length of 4096. Due to the page limit, we leave other hyperparameters in the appendix.

5.2 RESULTS

5.2.1 MATH-500

Before evaluating the pass rate, we first analyze the runtime overhead of ID-sampling on Math-500 datasets. We observe that ID-sampling incurs approximately 1.6–1.9 \times the total wall-clock time compared to the baseline of vanilla sampling for non-reasoning models, and 1.1–1.3 \times for reasoning models. Given that this additional computational cost could instead be allocated to generating more responses and selecting the best one, we present our results in terms of an equivalent N. Specifically, if the original results correspond to Bo8, we report them as equivalent $N = 16$, as ID-sampling consistently completes within twice the runtime of the original method.

We present our results for different models on Math-500 in Fig. 4. We observe a notable difference in the effectiveness of ID-sampling for non-reasoning models based on the aggregation method. While ID-sampling yields performance gains with increasing sample size (N) under a Best-of-N (BoN) selection strategy, it consistently performs poorly when results are combined via majority voting. This suggests that for non-reasoning models lacking self-reflection capabilities, the patterns of correct versus incorrect responses differ in ways that ID-sampling fails to exploit.

By contrast, our primary focus is on reasoning models with test-time scaling capabilities. With the DeepSeek-R1-Distill-Qwen-7B, a strong reasoning model with built-in self-evaluation and self-correction, ID-sampling consistently outperforms vanilla sampling. This is because while stronger

¹We have tried to use techniques like Yarn to extend the context length limit to a longer context. However, the model will lead to a consistent decrease rather than an increase when we increase the number of responses, and thus, we decided to just report Pass@n on longer context answers.

N	Vanilla		ID-sampling	
	BoN	cons@N	BoN	cons@N
1	0.00	0.00	3.45	3.45
4	0.00	0.00	0.00	3.45
8	0.00	0.00	0.00	3.45
16	3.45	0.00	0.00	3.45
32	3.45	0.00	10.34	6.90

(a) Llama-3.1-1B-Instruct

(b) Phi-4

Table 1: AIME-24 dataset: Pass rate results for different models with different number of samples. The best results in each setting are highlighted in bold. Since performance saturates quickly on these weaker LLM models, meaning that even the highest pass rate with $N = 32$ does not exceed the $N = 1$ pass rate for ID-sampling, we omit the use of equivalent N in this setting.

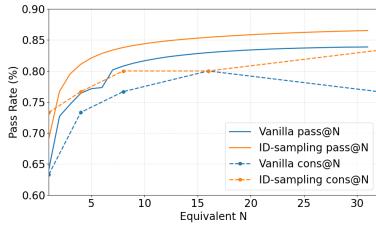


Figure 5: Qwen3-8B results on AIME-24.

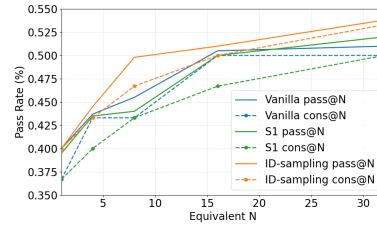


Figure 6: Deepseek-R1-distill-Qwen-7B results on AIME-25.

models excel at self-correction, they remain suboptimal at determining when to initiate the self-correction process. Compared to earlier models, each *trigger sentence* has a more pronounced effect, allowing ID-sampling to correct errors that might otherwise persist without explicit intervention.

5.2.2 AIME

The AIME datasets (AIME-24 and AIME-25) are highly challenging benchmarks that have become widely adopted for evaluating reasoning models DeepSeek-AI et al. (2025); Huang et al. (2024). In this setting, reward models exhibit substantially lower accuracy. Accordingly, rather than reporting BoN, we adopt Pass@N and Cons@N as our evaluation metrics.

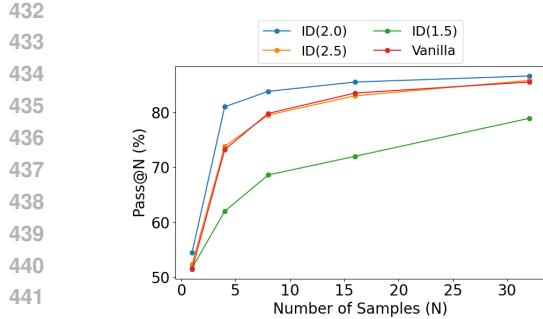
We first present the results for non-reasoning models in Table 1. We observe that ID-sampling can help improve the pass rate, and because of the fast saturation of pass rate, ID-sampling can even help the model to surpass its previous ceiling, i.e., $N = 1$ with ID-sampling is better than $N = 32$ with vanilla sampling. More importantly, for reasoning models, the wall-clock-time difference between vanilla sampling and ID-sampling becomes negligible. Specifically, for Qwen3 models, the time ratio remains within 1 ± 0.02 , and for R1-distill models, it is within 1 ± 0.1 with $\gamma = 2.0$. As a result, the use of equivalent N has little practical effect in this context. We present results on reasoning models in Fig. 5, 1, 6, 7, and Table 2. Across multiple models on the two most recent AIME

n	1	4	8	16
Vanilla Pass@n	67.8	81.6	85.9	88.4
S1 Pass @n	67.6	82.0	86.1	88.4
ID-sampling Pass@n	69.5	82.8	86.1	88.9
Vanilla Cons@n	56.7	76.7	80.0	83.3
S1 Cons@n	56.7	76.7	76.7	83.3
ID-sampling Cons@n	63.3	80.0	80.0	83.3

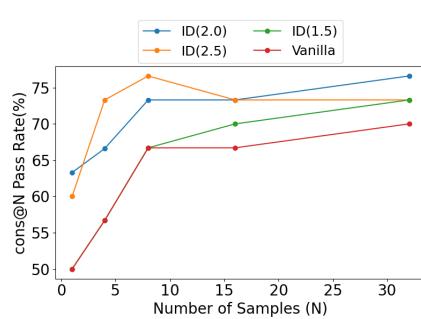
Table 2: Comparison between Vanilla sampling, S1, and ID-sampling on Qwen3-32B on AIME-25.

	Rel. Time
Vanilla	1.00
ID(2.5)	1.07
ID(2.0)	1.09
ID(1.5)	1.39

Table 3: Relative wall-clock time cost of ID-sampling for varying γ values using DeepSeek-R1-Distill-Qwen-7B on the AIME-24 dataset.



(a) Pass@N



(b) cons@N

Figure 7: AIME-24 Dataset: Pass rate results for DeepSeek-R1-Distill-Qwen-7B with different γ for ID-sampling, compared to vanilla sampling. The numbers in brackets are the γ used in ID-sampling. As different values of γ result in significantly varying runtimes, as shown in Table 3, we again omit the use of equivalent N in this analysis.

datasets, ID-sampling consistently outperforms both vanilla sampling and S1. This demonstrates the general applicability of our method, without requiring any assumptions about the specifics of post-training. By contrast, S1 only barely matches the performance of vanilla sampling on Deepseek-R1-distill-Qwen-7B, reinforcing the need for training as originally emphasized in the S1 paper. This is because S1 does not substantially alter the overall response pattern, particularly when measured by the positional frequency of linguistic markers, and therefore derives no benefit from it. For Qwen3-8B, which typically exhausts the full thinking budget, S1 has no opportunity to be applied, and its performance largely overlaps with vanilla sampling on both AIME-24 and AIME-25. For clarity, we therefore omit the S1 curve for Qwen3-8B models.

Ablation Study Given the stable performance gains of ID-sampling on AIME, here we conducted an ablation study to analyze the impact of the scaling factor γ on ID-sampling. We present the pass rate results for DeepSeek-R1-Distill-Qwen-7B on AIME-24 in Figures 7a and 7b, and report the relative inference time for each setting in Table 3.

We find that adjusting γ significantly impacts both performance and computational cost. In terms of runtime, $\gamma = 1.5$ yields the highest cost, whereas the other two settings remain within $1.1 \times$ the wall-clock time of vanilla sampling. Regarding performance, $\gamma = 2.0$ consistently achieves the best results in terms of Pass@N, while $\gamma = 2.5$ occasionally outperforms in terms of Cons@N, though with a small margin. Overall, ID-sampling proves to be a more effective sampling strategy than vanilla sampling, provided that γ is not too small, which would undermine our goal of emphasizing early-stage *trigger sentence* injection.

We observed a non-convex relationship between γ and performance, even without accounting for time, with one of the best settings at $\gamma = 2.0$. We attribute this to the fact that ID-sampling aims to match the patterns of correct responses, but can mismatch when the *trigger sentence* is inserted too frequently, i.e. γ is small. Overall, as most reasoning models right now share a similar reasoning pattern, we recommend $\gamma = 2.0$ as it strikes a favorable balance between performance and efficiency, and we have demonstrated its effectiveness through the extensive experiments presented above.

6 DISCUSSIONS

Choice of Trigger Sentence In our experiments, we use different *trigger sentences* for non-reasoning models and reasoning models. The difference is caused by the nature of the models. For non-reasoning models, a single word "wait" cannot trigger the self-correction process and can only introduce noise to the generation. For reasoning models, we found that for an unknown reason, adding the whole sentence will trigger the generation of an end-of-think token for a certain probability, and lead to an early stop in the generation sequence. While this does not always harm the performance, this is not something we want, as we want to use the *trigger sentence* to trigger

486 self-correction. Overall, as test-time scaling is highly correlated to reasoning models, a single word,
 487 “wait,” will be sufficient as shown in our experiments.
 488

489 **Limitations** Our proposed method also has several clear limitations. The most significant is that
 490 the performance of our method could significantly depend on the underlying model. While we have
 491 used our experiments to show that ID-sampling works on popular models, it is possible for some
 492 models to be incompatible with the algorithm, especially if the models are trained with losses that in-
 493 incorporate more than just the accuracy of the final answer. Moreover, incorporating *trigger sentences*
 494 into the generation process requires ID-sampling to invoke multiple generation steps. While this the-
 495oretically incurs no additional cost on the KV-cache, in practice, it can lead to increased inference
 496 time. To address this concern, we use equivalent N in our experiments, which accounts for total
 497 wall-clock time, to enable a fair comparison across different sampling methods. Additionally, due
 498 to the need for multiple sampling steps, ID-sampling is currently tested only on open-source models.
 499 While the same idea can also be applied to black-box models through multi-round generations, this
 500 will introduce extra assistant tokens during generation, which may cause a slight distribution shift
 501 when applied to black-box models.
 502

503 7 CONCLUSIONS

504 In this paper, we introduced Iterative Deepening (ID) Sampling, a simple yet effective algorithm for
 505 scaling test-time compute more efficiently than standard sampling. We showed that ID-sampling
 506 improves test-time performance on challenging reasoning tasks like math across diverse reasoning
 507 models. Our findings suggest that while current models possess strong self-correction capabilities,
 508 they remain limited in autonomously deciding when to invoke such mechanisms. Furthermore, our
 509 results highlight that guiding generation to mimic the patterns of correct responses can provide a
 510 benefit at test-time.
 511

512 REFERENCES

514 Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar,
 515 Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. Phi-4 techni-
 516 cal report. *arXiv preprint arXiv:2412.08905*, 2024.

517 Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. Step-level value preference optimiza-
 518 tion for mathematical reasoning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen
 519 (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 7889–
 520 7903, Miami, Florida, USA, November 2024a. Association for Computational Linguistics.
 521 doi: 10.18653/v1/2024.findings-emnlp.463. URL [https://aclanthology.org/2024.](https://aclanthology.org/2024.findings-emnlp.463)
 522 [findings-emnlp.463](https://aclanthology.org/2024.findings-emnlp.463).
 523

524 Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. Alphamath almost zero: process supervision
 525 without process, 2024b. URL <https://arxiv.org/abs/2405.03553>.

526 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
 527 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
 528 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

529 Weizhe Chen, Sven Koenig, and Bistra Dilkina. Reprompt: Planning by automatic prompt engineer-
 530 ing for large language models agents. *arXiv preprint arXiv:2406.11132*, 2024c.

531 Weizhe Chen, Sven Koenig, and Bistra Dilkina. Why solving multi-agent path finding with large
 532 language model has not succeeded yet. *arXiv preprint arXiv:2401.03630*, 2024d.

533 Weizhe Chen, Zhicheng Zhang, Guanlin Liu, Renjie Zheng, Wenlei Shi, Chen Dun, Zheng Wu,
 534 Xing Jin, and Lin Yan. Flaming-hot initiation with regular execution sampling for large language
 535 models. *arXiv preprint arXiv:2410.21236*, 2024e.

536 Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models
 537 to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.

540 Ziru Chen, Michael White, Raymond Mooney, Ali Payani, Yu Su, and Huan Sun. When is tree
 541 search useful for llm planning? it depends on the discriminator. *arXiv preprint arXiv:2402.10890*,
 542 2024f.

543 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
 544 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
 545 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
 546 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
 547 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
 548 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
 549 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
 550 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai
 551 Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang,
 552 Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang,
 553 Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang,
 554 Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang,
 555 R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shengfeng
 556 Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing
 557 Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen
 558 Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong
 559 Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu,
 560 Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xi-
 561 aoshua Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia
 562 Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng
 563 Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong
 564 Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong,
 565 Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou,
 566 Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying
 567 Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda
 568 Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu,
 569 Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu
 570 Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforce-
 ment learning, 2025. URL <https://arxiv.org/abs/2501.12948>.

571 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
 572 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
 573 *arXiv preprint arXiv:2407.21783*, 2024.

574 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen.
 575 Critic: Large language models can self-correct with tool-interactive critiquing. *arXiv preprint*
 576 *arXiv:2305.11738*, 2023.

577 Xinyu Guan, Li Lyra Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, and Mao Yang.
 578 rstar-math: Small llms can master math reasoning with self-evolved deep thinking. *arXiv preprint*
 579 *arXiv:2501.04519*, 2025.

580 Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu.
 581 Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*,
 582 2023.

583 Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song,
 584 and Denny Zhou. Large language models cannot self-correct reasoning yet. *arXiv preprint*
 585 *arXiv:2310.01798*, 2023.

586 Zhen Huang, Haoyang Zou, Xuefeng Li, Yixiu Liu, Yuxiang Zheng, Ethan Chern, Shijie Xia, Yi-
 587 wei Qin, Weizhe Yuan, and Pengfei Liu. O1 replication journey—part 2: Surpassing o1-preview
 588 through simple distillation, big progress or bitter lesson? *arXiv preprint arXiv:2411.16489*, 2024.

589 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 590 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv*
 591 *preprint arXiv:2412.16720*, 2024.

594 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.
 595 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 596 serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating
 597 Systems Principles*, 2023.

598 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 599 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *arXiv preprint
 600 arXiv:2305.20050*, 2023.

602 Jiacheng Liu, Andrew Cohen, Ramakanth Pasunuru, Yejin Choi, Hannaneh Hajishirzi, and Asli
 603 Celikyilmaz. Making ppo even better: Value-guided monte-carlo tree search decoding. *arXiv
 604 preprint arXiv:2309.15028*, 2023.

606 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri
 607 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement
 608 with self-feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

609 Sara Vera Marjanović, Arkil Patel, Vaibhav Adlakha, Milad Aghajohari, Parishad BehnamGhader,
 610 Mehar Bhatia, Aditi Khandelwal, Austin Kraft, Benno Krojer, Xing Han Lù, et al. Deepseek-r1
 611 thoughtology: Let's think about llm reasoning. *arXiv preprint arXiv:2504.07128*, 2025.

613 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
 614 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
 615 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

616 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Di-
 617 rani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a bench-
 618 mark. In *First Conference on Language Modeling*, 2024.

619 Ranajoy Sadhukhan, Zhuoming Chen, Haizhong Zheng, Yang Zhou, Emma Strubell, and Beidi
 620 Chen. Kinetics: Rethinking test-time scaling laws. *arXiv preprint arXiv:2506.05333*, 2025.

622 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
 623 language agents with verbal reinforcement learning. In *Advances in Neural Information Process-
 624 ing Systems (NeurIPS)*, 2023. URL https://papers.nips.cc/paper_files/paper/2023/file/1b44b878bb782e6954cd888628510e90-Paper-Conference.pdf.

626 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 627 can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.

629 Hanshi Sun, Momin Haider, Ruiqi Zhang, Huitao Yang, Jiahao Qiu, Ming Yin, Mengdi Wang, Peter
 630 Bartlett, and Andrea Zanette. Fast best-of-n decoding via speculative rejection. *arXiv preprint
 631 arXiv:2410.20290*, 2024.

633 Mudit Verma, Siddhant Bhambri, and Subbarao Kambhampati. On the brittle foundations of react
 634 prompting for agentic large language models. *arXiv preprint arXiv:2405.13966*, 2024.

635 Junlin Wang, Shang Zhu, Jon Saad-Falcon, Ben Athiwaratkun, Qingyang Wu, Jue Wang, Shuai-
 636 wen Leon Song, Ce Zhang, Bhuwan Dhingra, and James Zou. Think deep, think fast: Investigat-
 637 ing efficiency of verifier-free inference-time-scaling methods. *arXiv preprint arXiv:2504.14047*,
 638 2025a.

640 Xinglin Wang, Yiwei Li, Shaoxiong Feng, Peiwen Yuan, Yueqi Zhang, Jiayi Shi, Chuyi Tan, Boyuan
 641 Pan, Yao Hu, and Kan Li. Every rollout counts: Optimal resource allocation for efficient test-time
 642 scaling. *arXiv preprint arXiv:2506.15707*, 2025b.

643 Xuezhi Wang and Denny Zhou. Chain-of-thought reasoning without prompting. *arXiv preprint
 644 arXiv:2402.10200*, 2024.

646 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 647 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
arXiv preprint arXiv:2203.11171, 2022.

648 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws:
 649 An empirical analysis of compute-optimal inference for llm problem-solving. In *The Thirteenth*
 650 *International Conference on Learning Representations*, 2025.

651

652 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 653 Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
 654 *arXiv:2505.09388*, 2025.

655

656 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao.
 657 ReAct: Synergizing reasoning and acting in language models. In *International Conference on*
 658 *Learning Representations (ICLR)*, 2023. URL https://openreview.net/pdf?id=WE_v1uYUL-X.

659

660 Zhiyuan Zeng, Qinyuan Cheng, Zhangyue Yin, Bo Wang, Shimin Li, Yunhua Zhou, Qipeng Guo,
 661 Xuanjing Huang, and Xipeng Qiu. Scaling of search and learning: A roadmap to reproduce o1
 662 from reinforcement learning perspective. *arXiv preprint arXiv:2412.14135*, 2024.

663

664 Di Zhang, Xiaoshui Huang, Dongzhan Zhou, Yuqiang Li, and Wanli Ouyang. Accessing gpt-4 level
 665 mathematical olympiad solutions via monte carlo tree self-refine with llama-3 8b. *arXiv preprint*
 666 *arXiv:2406.07394*, 2024a.

667

668 Di Zhang, Jianbo Wu, Jingdi Lei, Tong Che, Jiatong Li, Tong Xie, Xiaoshui Huang, Shufei Zhang,
 669 Marco Pavone, Yuqiang Li, et al. Llama-berry: Pairwise optimization for o1-like olympiad-level
 670 mathematical reasoning. *arXiv preprint arXiv:2410.02884*, 2024b.

671

672 Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B Tenenbaum, and Chuang Gan.
 Planning with large language models for code generation. *arXiv preprint arXiv:2303.05510*,
 673 2023.

674

675 Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu,
 676 Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical
 677 reasoning. *arXiv preprint arXiv:2501.07301*, 2025.

678

679 Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Lan-
 680 guage agent tree search unifies reasoning acting and planning in language models. *arXiv preprint*
 681 *arXiv:2310.04406*, 2023.

682

683 A DISCLAIMER OF LLM USAGE

685 Apart from being the subject of study, large language models were used solely to polish the writing
 686 in this paper. All substantive contributions are attributable to the authors.

689 B COMPLETE PROOF FOR THEOREM 4.1

691 Let L_{total} represent the total generated length resulting from an iterative process. Let l denote the
 692 length generated in the final iteration. We naturally have $l \leq L$, otherwise the generation will be
 693 finished earlier.

694 We aim to establish an upper bound for L_{total} . We analyze the contributions to the total length
 695 by considering the sequence of generated lengths in reverse chronological order. Assume that this
 696 sequence of lengths can be bounded above by a geometric sequence whose first term is l and whose
 697 common ratio is $r = 1/\gamma$, where γ is a constant greater than 1 ($\gamma > 1$).

698 Under this assumption, the total length L_{total} is bounded by the sum of this infinite geometric series:

$$700 \quad L_{total} \leq \sum_{k=0}^{\infty} l \left(\frac{1}{\gamma} \right)^k$$

702 Since it is assumed that $\gamma > 1$, the common ratio $r = 1/\gamma$ satisfies $0 < r < 1$. Therefore, the
 703 geometric series converges, and its sum is:
 704

$$705 \sum_{k=0}^{\infty} l \left(\frac{1}{\gamma} \right)^k = \frac{l}{1 - \frac{1}{\gamma}}$$

708 We can simplify the denominator of this expression:
 709

$$710 \frac{l}{1 - \frac{1}{\gamma}} = \frac{l}{\frac{\gamma-1}{\gamma}} = \frac{\gamma l}{\gamma - 1}$$

713 Thus, we have the bound:
 714

$$715 L_{total} \leq \frac{\gamma l}{\gamma - 1}$$

716 Finally, using the given condition that $l \leq L$, we substitute L to obtain the final upper bound for the
 717 total length:
 718

$$719 L_{total} \leq \frac{\gamma l}{\gamma - 1} \leq \frac{\gamma L}{\gamma - 1}$$

721 Therefore, the total generated length L_{total} is bounded above by $\frac{\gamma L}{\gamma - 1}$.
 722

723 C HYPERPARAMETERS

726 All experiments are conducted on an NVIDIA A100 GPU (80GB memory) using the vLLM serving
 727 engine Kwon et al. (2023). By default, we use a carefully selected set of hyperparameters and
 728 configurations to balance efficiency and performance:

- 730 • Based on the typical sampling budget used in training (10–40), we set the maximum number of N to 32 for sampling.
 731
- 732 • All numbers are based on our implementations and the numbers are slightly different to the
 733 ones reported on technical reports due to the prompt used and the early cut-off based on the
 734 budget.
 735
- 736 • By default, we use $temperature = 0.6$, $top_p = 0.95$. For Qwen3, we additionally use
 737 $top_k = 20$ as recommended by the official technical report.
 738

739 D ADDITIONAL DISCUSSIONS

741 **Potential Extension to Search-based Test-time Scaling** In this paper, our study primarily focuses on ID-sampling without integrating it into tree-search algorithms. However, we emphasize
 742 that our proposed framework can be directly extended to methods such as beam search and Monte
 743 Carlo Tree Search (MCTS). In these cases, the budget in Algorithm 1 can be directly interpreted as
 744 the computational budget in MCTS or the number of iterations in beam search. One challenge in this
 745 extension is that self-correction mechanisms reduce the reliance on early model outputs being correct,
 746 distinguishing the need for a new generation of PRMs from previously released ones. However,
 747 future researchers who develop or have access to more accurate PRMs that explicitly model the self-
 748 correction process can leverage our ID-sampling framework to enhance the reasoning capabilities
 749 of fixed models.
 750

751 **Choice of Datasets** In our experiments, we focus primarily on mathematical reasoning datasets,
 752 with results on a science QA dataset provided in the appendix. While test-time scaling techniques
 753 could, in principle, be extended to more general reasoning-related domains like coding, such exten-
 754 sions are not directly applicable in our setting. Methodologically, the injection of *trigger sentences*
 755 is ill-suited for code generation tasks, where inserting additional tokens may introduce syntactic or
 semantic errors. We therefore leave the integration of ID-sampling to future work.

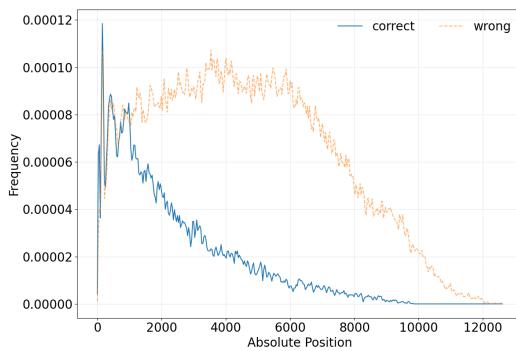


Figure 8: Positional frequency of thinking-related linguistic markers in Deepseek-R1-distill-Qwen-7B responses using ID-sampling on the AIME-2024 dataset.

E LINGUISTIC MARKERS

In the paper, we have presented the occurrence frequencies of the linguistic markers as our motivation. We provide the full list of such markers here, which is the same as Wang et al. (2025a): "however", "wait", "alternatively", "hmm".

F ADDITIONAL EXPERIMENT RESULTS

F.1 GPQA RESULTS

Model	Sampling	Cons@1	Cons@4	Cons@8	Cons@16	Cons@32
Deepseek-R1-7B	Vanilla Sampling	48.9	59.7	65.2	67.3	70.1
	ID Sampling	50.0	60.8	66.8	67.9	71.1
Qwen3-8B	Vanilla Sampling	55.4	65.2	71.1	73.9	74.4
	ID Sampling	55.4	66.8	72.2	73.9	76.6

Table 4: ID-sampling performance on difference models on GPQA-diamond dataset. Specifically, Deepseek-R1-7B denotes Deepseek-R1-distill-Qwen-7B.

In GPQA, the wall-clock time of ID-sampling and vanilla sampling differs by only about $1.10\times$, a negligible overhead compared to the inherent runtime randomness of LLMs. Accordingly, we do not report equivalent- N results in this setting. As shown in Table 4, we observe that on both models, ID-sampling stably outperform vanilla sampling when combined with majority voting (Cons@ n). We choose not to report Pass@ N , as for multiple-choice questions such as those in GPQA, as this metric primarily incentivizes diversity in responses rather than better answer quality.

G FREQUENCY OF LINGUISTIC MARKER

In Fig. 8, we show the frequency of linguistic marker after using ID-sampling. Specifically, because of the smooth applied to the curve, there are only a few spikes that are noticeable. Aside from these spikes, we do not observe significant changes in the overall shape of the frequency curves for either correct or incorrect trajectories when comparing vanilla sampling to ID sampling. Nonetheless, the average frequency for correct trajectories appears slightly higher, while the frequency for incorrect trajectories is roughly the same.

We believe the lack of major change in incorrect trajectories is due to the limited number of "wait" prompts added. And these added patterns convert more incorrect to correct trajectories, without significantly disrupting the overall generation patterns. Similarly, the slight increase in average frequency for correct trajectories is likely due to a subset of previously incorrect paths being corrected.

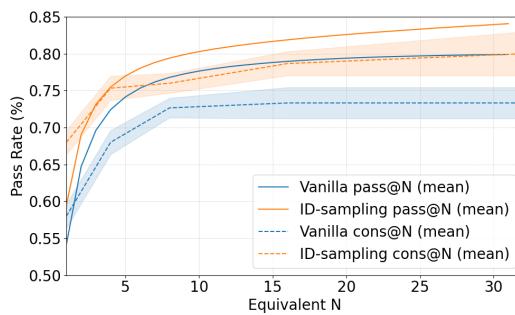


Figure 9: Performance of ID-sampling on Qwen3-8B evaluated on AIME-25. Results are averaged over five random seeds, with shaded regions indicating standard deviation.

This increases the overall average, as these corrected trajectories tend to exhibit more frequent occurrences of reasoning-related linguistic markers in their earlier segments.

H STATISTICAL SIGNIFICANCE

Due to the constraints of computational resources and the breadth of domains and models required to demonstrate the generalizability of our approach, we report results from a single run in the main paper. Here, we supplement those findings with an illustrative statistical evaluation using Qwen3-8B on AIME25.

As shown in Fig. 9, ID sampling consistently outperforms vanilla sampling. The standard deviation of Pass@n is extremely small and almost imperceptible in the plot because we use an unbiased estimator of Pass@n. Although the variance of Cons@n is larger, our proposed method remains consistently superior across runs.