

# Seer Self-Consistency: Advance Budget Estimation for Adaptive Test-Time Scaling

Shiyu Ji<sup>†</sup> Yixuan Wang<sup>†</sup> Yijun Liu Qingfu Zhu Wanxiang Che\*

Research Center for Social Computing and Interactive Robotics,

Harbin Institute of Technology, China

{syji, car}@ir.hit.edu.cn

 <https://github.com/noforit/SeerSC>

## Abstract

Test-time scaling improves the inference performance of Large Language Models (LLMs) but also incurs substantial computational costs. Although recent studies have reduced token consumption through dynamic self-consistency, they remain constrained by the high latency of sequential requests. In this paper, we propose SeerSC, a dynamic self-consistency framework that simultaneously improves token efficiency and latency by integrating System 1 and System 2 reasoning. Specifically, we utilize the rapid System 1 to compute the answer entropy for given queries. This score is then used to evaluate the potential of samples for scaling, enabling dynamic self-consistency under System 2. Benefiting from the advance and accurate estimation provided by System 1, the proposed method can reduce token usage while simultaneously achieving a significant decrease in latency through parallel generation. It outperforms existing methods, achieving up to a 47% reduction in token consumption and a 43% reduction in inference latency without significant performance loss.

## 1 Introduction

Large Language Models (LLMs) (Hurst et al., 2024; Yang et al., 2024; Guo et al., 2025) have demonstrated exceptional performance in reasoning tasks, particularly in mathematics. By leveraging methods such as self-consistency (Wang et al., 2022), LLMs can substantially improve their performance during inference through increased computation (Chen et al., 2025; Zhang et al., 2025) without additional training. While improving performance, Test-Time Scaling (TTS) methods also significantly increase the computational cost (Brown et al., 2024; Chen et al., 2024a). Many efforts (Feng et al., 2025; Liu et al., 2025b)

<sup>†</sup>Equal contribution.

\*Corresponding author.

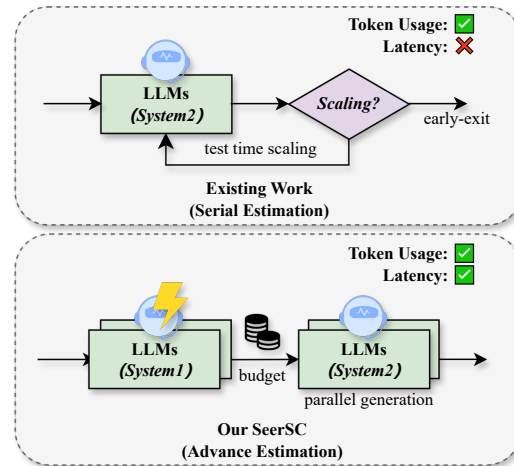


Figure 1: Comparison of the Proposed method with Existing Methods. Unlike sequential estimation methods, SeerSC leverages fast System 1 to guide the parallel generation of System 2.

focus on optimizing the trade-off between performance and efficiency.

For ensemble methods like Self-Consistency (SC), a critical challenge is dynamically determining the optimal number of ensemble samples based on the complexity of the input instance (Wang et al., 2024). Adaptive Consistency (AC) (Aggarwal et al., 2023) generates responses sequentially and decides whether to terminate scaling based on the diversity of the current answers and a specified threshold. Early-Stopping Self-Consistency (ESC) (Li et al., 2024) employs a sliding window to track the diversity of the current responses, enabling dynamic early stopping for self-consistency. By dynamically adjusting the sampling budget, these methods significantly reduce token consumption during the TTS phase.

However, a reduction in token usage does not always translate into a corresponding decrease in latency (Wang et al., 2025b; Huang et al., 2025). Existing dynamic SC methods typically estimate




Figure 2: Observational Experiments on the Math Dataset.

the required computation for each query through **sequential requests**, as shown in Figure 1, which **conflicts with the efficient parallel sampling techniques** used by modern inference engines (Kwon et al., 2023; Zheng et al., 2024). This issue is further exacerbated for recent reasoning models (Guo et al., 2025) as the output length increases. If a computational budget can be allocated to different samples prior to their execution, we can leverage parallel sampling to concurrently optimize both the token consumption and the inference latency, leading to more efficient TTS.

In this paper, we propose *SeerSC*, a dynamic self-consistency framework that integrates intuition (System 1) and reasoning (System 2) (Li et al., 2025) to enable the pre-allocation of computational resources. Observing the consistency in answer diversity between the rapid System 1 and deliberative System 2 generation forms, we propose using the entropy of the fast System 1 answers to guide the computational budget allocation for System 2. Specifically, we first use the System 1 approach to generate multiple responses for the sample in parallel, and subsequently calculate the entropy of the resulting answers. Based on the answer entropy score, we can achieve fine-grained budget allocation, directing limited computational resources to samples with greater scaling potential to maximize performance improvement.

We conduct experiments across multiple model configurations, including reasoning and instruct settings, and utilize mainstream mathematical reasoning benchmarks. By leveraging the advance budget estimation from the rapid System 1, *SeerSC* is able to decrease the token usage during the TTS stage while exploiting parallel generation to mitigate overall latency. Our method demonstrably surpasses current techniques, delivering up to 47% lower token usage and 43% reduced inference la-

tency without substantial performance degradation. Furthermore, compared to sequential methods, our approach can be more effectively integrated with path-pruning techniques (Fu et al., 2025), achieving orthogonal performance gains and expanding its range of application.

Our main contributions can be summarized as follows:

- We propose Seer Self-Consistency (SeerSC), which utilizes the answer entropy of System 1 to guide the dynamic budget allocation of System 2 during TTS stage.
- We propose a fine-grained budget allocation method that ensures computational resources are effectively directed to samples exhibiting a high potential for scaling.
- Experiments conducted across multiple settings demonstrate that the proposed method outperforms existing dynamic SC approaches in both token consumption and latency.

## 2 Observation

As discussed, achieving parallel inference with a dynamic budget requires determining the budget for each sample prior to execution. In search of an effective metric for budget allocation, we initiate an analytical investigation into the gains provided by self-consistency. Our findings resulted in the following two primary observations:

**Answer diversity can reflect the scaling potential.** As shown in Figure 2a, not all samples benefit from Test-Time Scaling (TTS). Thus, an effective metric is required to guide budget allocation. We further perform a statistical analysis of the number of unique answers resulting from Self-Consistency. As shown in Figure 2b, samples with

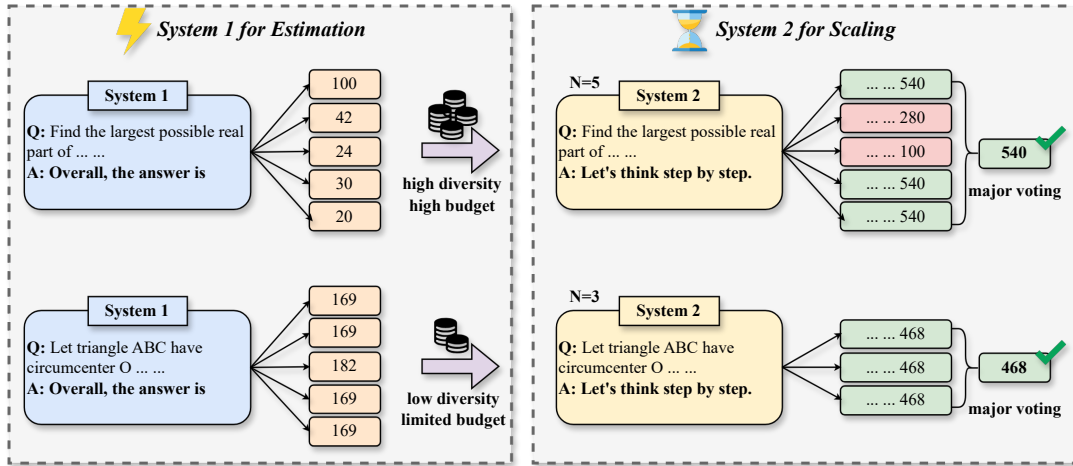


Figure 3: Overall workflow of the proposed SeerSC framework. We first utilize System 1 to rapidly estimate the answer diversity for each sample, and subsequently allocate varying computational resources for the System 2 phase based on the answer entropy.

negligible TTS gain typically yielded a lower number of answers, indicating a higher degree of certainty in the responses. This insight inspires us to use the entropy of the generated answers as a metric to determine the TTS budget for each sample.

**The answer entropy of System 1 and System 2 exhibits consistency.** Nevertheless, acquiring the answer entropy before response generation is unfeasible. While current approaches (Aggarwal et al., 2023; Li et al., 2024) compute the entropy via serial execution, they still suffer from significant latency challenges. In order to compute the certainty of various samples in advance, we investigate the use of low-cost techniques for approximating the answer entropy. As shown in Figure 2c, we illustrate the answer entropy of samples generated via early stopping after a varying number of steps. The leftmost point represents System 1, which generates the answer directly, while the rightmost point represents System 2, which involves full reasoning. As evident from the figure, the answer entropy for the two classes of samples (those that benefit from TTS and those that do not) shows a distinct difference in both System 1 and System 2. This confirms the potential for utilizing the rapid System 1 process to estimate the answer entropy obtained from System 2.

### 3 Methodology

#### 3.1 Overall

Based on these observations, we propose Seer Self-Consistency (SeerSC), a dynamic budget Test-Time Scaling framework that integrates the intuition of

System 1 and the reasoning of System 2. The overall workflow of the framework is illustrated in Figure 3. We first perform System 1-level inference to obtain the answer entropy (§3.2) for each sample. Subsequently, we rank the samples based on their answer entropy to allocate differential amounts of computation (§3.3) according to the pre-defined budget. This allows us to efficiently execute System 2-level TTS within a given budget.

#### 3.2 System 1 for Estimation

As illustrated in Figure 3, given a problem  $X$ , System 1 first performs direct answer generation (see Appendix A.3 for implementation details on reasoning models) with a sampling size of  $M$ , producing an answer set  $\text{Ans} = \{A_1, A_2, \dots, A_M\}$ .

For each answer  $A \in \text{Ans}$ , we define a confidence score  $S(x)$  based on its token-level probabilities. Let  $A = [a_1, a_2, \dots, a_k]$  where  $k = |A|$  and  $a_i$  denotes the  $i$ -th token in the answer. The confidence is defined as:

$$S(A) = \exp\left(\frac{1}{k} \sum_{i=1}^k \log p(a_i)\right) \quad (1)$$

where  $p(a_i)$  denotes the probability of token  $a_i$ . This formulation corresponds to the exponentiated average log-likelihood of the tokens in  $A$ .

We then compute the entropy  $E(X)$  over *answer categories* to guide the load allocation in System 2. Unlike vanilla Shannon entropy that depends on unweighted category frequencies, we integrate category confidence weights into the entropy calculation.

Let  $\mathcal{C} = \{c_1, \dots, c_m\}$  be the set of distinct answer categories and  $\text{Ans} = \{A_1, \dots, A_M\}$  the sampled answers. The (unnormalized) weight of category  $c_j$  is the sum of member-answer confidences:

$$W(c_j) = \sum_{A \in \text{Ans}, A \in c_j} S(A) \quad (2)$$

We normalize to obtain a probability distribution over categories:

$$\tilde{W}(c_j) = \frac{W(c_j)}{\sum_{c \in \mathcal{C}} W(c)} \quad (3)$$

The confidence-weighted entropy of problem  $X$  is then

$$E(X) = - \sum_{c \in \mathcal{C}} \tilde{W}(c) \log \tilde{W}(c) \quad (4)$$

### 3.3 System 2 for Scaling

In System 2, we allocate a computation budget  $B(X)$  for each problem  $X$  based on its estimated entropy  $E(X)$  from System 1. As observed in Section 2, problems with lower entropy require fewer sampled reasoning paths, whereas problems with higher entropy benefit from a larger number of samples. In contrast to Self-Consistency (SC), which assigns a fixed budget of  $N$  samples to every problem, our adaptive allocation follows the rule:

$$B(X) = \begin{cases} 1, & E(X) < \tau_1, \\ N/2, & \tau_1 \leq E(X) < \tau_2, \\ N, & \text{otherwise.} \end{cases} \quad (5)$$

where  $\tau_1$  and  $\tau_2$  are entropy thresholds.

By definition, the entropy  $E(X)$  lies in the range

$$0 \leq E(X) \leq \log M \quad (6)$$

where  $M$  is the number of sampled direct answers in System 1. The lower bound is attained when all direct answers fall into a single category, yielding  $E(X) = 0$ . The upper bound is attained when every sampled direct answer belongs to a distinct category, i.e.,  $m = M$ , where  $m$  is the number of unique categories. In this case, the category distribution is uniform, and the entropy reaches its maximum value  $\log M$ .

Based on this range, we empirically set the thresholds as fractions of  $\log M$ :

$$\tau_1 = \frac{1}{10} \log M, \quad \tau_2 = \frac{1}{3} \log M. \quad (7)$$

This empirical setting provides a practical balance between efficiency and accuracy (see Appendix B for a detailed sensitivity analysis).

## 4 Experiment

### 4.1 Experimental Setup

**Evaluation benchmarks and models.** We conduct evaluations on the following datasets. MATH500 (Lightman et al., 2023) represents a representative subset of problems from the MATH dataset (Hendrycks et al., 2021). AIME2024 (Art of Problem Solving, 2024), AIME2025 (Art of Problem Solving, 2025), and AMC2023 (Art of Problem Solving, 2023) consist of high-difficulty mathematical competition problems. GPQA Diamond (Rein et al., 2024) serves as a benchmark for graduate-level STEM reasoning tasks, providing a challenging evaluation setting. We evaluate on the following models: DeepSeek-R1-Distill-Qwen-7B (Guo et al., 2025), DeepSeek-R1-Distill-Llama-8B (Guo et al., 2025), and Qwen3-4B (Yang et al., 2025). The first two are representative reasoning models, while Qwen3-4B supports both *thinking* and *non-thinking* modes, and we conduct evaluations under both settings. Additionally, we extend our evaluation to the larger-scale QwQ-32B (Team, 2025) model to verify generalizability (see Appendix C.1).

**Selected baselines.** We compare our method against four baseline methods:

- **Chain of Thought (CoT)** (Wei et al., 2022) employs standard CoT prompting, where the model generates a single reasoning trajectory that directly leads to the final answer.
- **Self-Consistency (SC)** (Wang et al., 2022) generates  $N$  independent reasoning paths and determines the final answer by majority voting on the predicted solutions.
- **Adaptive Consistency (AC)** (Aggarwal et al., 2023) sequentially generates reasoning paths in a sequential manner. In each step, it computes the relative frequency of unique answers and terminates sampling once the frequency of a candidate exceeds a predefined threshold or the maximum  $N$  is reached. The final answer is again decided by majority voting.
- **Early-Stopping Self-Consistency (ESC)** (Li et al., 2024) improves standard SC by sequentially generating reasoning paths until the upper limit of  $N$  is reached or the solutions within a sliding window of size  $W$  converge

to the same answer. The final prediction is then determined via majority voting.

**Implementation details.** To ensure a fair comparison, we search for the optimal sampling number  $N$  for the standard SC baseline on each model and dataset, as gains beyond this point are limited due to diminishing returns (Chen et al., 2024a). All baselines and our method are evaluated under this  $N$  setting. For AC baseline, the stopping threshold is set to 0.95. For ESC baseline, we follow the original paper and set the window size  $W = 5$ . For DeepSeek-R1-Distill-Qwen-7B, DeepSeek-R1-Distill-Llama-8B, and Qwen3-4B (thinking mode), we set the maximum output length to 10,240 on Math500 and AMC2023, and to 16,384 on the other datasets. For Qwen3-4B (non-thinking mode), the maximum output length is fixed to 8,192 across all datasets. All experiments are conducted using vLLM on NVIDIA A100-80GB GPUs with CUDA 12.4. The setting of  $N$  and additional implementation details can be found in Appendix A.

## 4.2 Main Results

As shown in Table 1, we report the results of SeerSC and other baseline methods across all evaluated models and datasets. The reported metrics include accuracy, the average number of output tokens and latency per problem, with all experiments repeated multiple times to ensure stability.

**SeerSC matches baselines in accuracy.** Across all models and datasets, SeerSC consistently achieves accuracy comparable to the baseline methods. On medium-difficulty tasks such as MATH500 and AMC2023, its accuracy is essentially on par with SC, with only minor fluctuations (within  $\pm 1$  percentage point). On the more challenging datasets AIME2024/2025 and GPQA-Diamond, SeerSC also sustains stable performance without noticeable degradation. Given the stochasticity introduced by high-temperature decoding, the overall accuracy level remains comparable to that of the baselines.

**SeerSC reduces token usage.** SeerSC consistently surpasses SC in terms of token efficiency. By adaptively determining the number of samples in advance, it avoids a large amount of unnecessary generation. As a result, SeerSC can substantially reduce token usage while maintaining accuracy. On MATH500, AMC2023, and GPQA-D across all

models, it achieves lower token consumption than all other baselines. On AIME2024/2025, although the token usage is slightly higher than that of AC and ESC, it remains considerably lower than SC.

**SeerSC significantly lowers latency.** The core advantage of SeerSC lies in its inference efficiency. Compared to SC, SeerSC substantially reduces latency across the majority of models and datasets. For example, on DeepSeek-R1-Distill-Qwen-7B, SeerSC achieves an average latency of only 5.3s on MATH500, a clear improvement over 8.4s of SC, representing a 37% reduction. On Qwen3-4B (non-thinking), the latency advantage is even more pronounced: SeerSC requires only 3.7s on GPQA-D, compared to 6.5s of SC, yielding a reduction of approximately 43%. By contrast, AC and ESC suffer from severe latency bottlenecks due to their sequential generation design, consistently running slower than SC across models and datasets. On Qwen3-4B (non-thinking), both methods take on average twice as long as SC. On DeepSeek-R1-Distill-Llama-8B and Qwen3-4B (thinking) for AIME2024/2025 and GPQA-D, where long reasoning trajectories dominate computation, the latencies of AC and ESC become comparable to SC. Nevertheless, under these settings, SeerSC still achieves significantly faster inference than all baseline methods.

**Overall.** SeerSC achieves accuracy comparable to SC, AC, and ESC, while significantly outperforming all baselines in terms of token and latency efficiency. It thus provides a more practical and effective trade-off for applications where both low latency and high accuracy are critical.

## 5 Analysis

### 5.1 Latency Analysis




Figure 4: Latency comparison of System 1 and System 2 in SeerSC on AIME-2024, using DeepSeek-R1-Distill-Qwen-7B and Qwen3-4B (non-thinking).

**The latency of System 1 is negligible.** As shown in Figure 4, the latency of System 1 accounts for

Table 1: Main experimental results of SeerSC compared with baseline methods (CoT, SC, AC, and ESC) across diverse models and benchmarks. We report accuracy (%), average tokens ( $\times 10^3$ ), and average latency (s).

Methods	MATH500			AIME2024			AIME2025			AMC2023			GPQA-D		
	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓
<i>DeepSeek-R1-Distill-Qwen-7B</i>															
CoT	90.0	–	–	49.7	–	–	38.8	–	–	83.9	–	–	44.7	–	–
SC	<b>93.9</b>	29.8	8.4	<b>74.4</b>	344.4	158.6	58.6	368.6	174.7	<b>91.6</b>	42.5	13.2	48.1	128.5	51.8
AC	93.8	19.4	12.4	73.7	<b>223.7</b>	214.1	59.3	<b>286.9</b>	233.7	91.3	<b>32.8</b>	24.8	47.7	98.3	73.1
ESC	93.6	22.2	13.9	74.0	270.4	261.2	58.3	319.5	266.8	91.1	34.9	26.7	48.0	105.9	82.0
SeerSC	93.0	<b>17.7</b>	<b>5.3</b>	74.1	315.6	<b>151.8</b>	<b>60.2</b>	356.0	<b>173.8</b>	91.5	33.5	<b>9.4</b>	<b>48.6</b>	<b>95.3</b>	<b>37.8</b>
<i>DeepSeek-R1-Distill-Llama-8B</i>															
CoT	84.5	–	–	39.3	–	–	27.8	–	–	78.6	–	–	46.2	–	–
SC	91.5	32.8	16.9	<b>68.8</b>	379.7	376.0	44.3	387.0	395.1	88.6	45.6	28.1	50.7	74.6	62.5
AC	91.3	24.3	21.8	67.7	290.0	361.4	44.8	312.2	392.3	88.6	34.5	39.5	49.9	62.7	71.6
ESC	<b>91.7</b>	26.7	23.3	67.9	337.3	436.6	<b>44.9</b>	339.1	424.3	<b>88.9</b>	37.5	40.9	50.4	66.2	75.1
SeerSC	90.2	<b>21.2</b>	<b>10.6</b>	66.3	<b>288.6</b>	<b>294.0</b>	43.8	<b>308.3</b>	<b>300.6</b>	87.3	<b>31.5</b>	<b>16.6</b>	<b>50.8</b>	<b>53.0</b>	<b>43.5</b>
<i>Qwen3-4B (thinking)</i>															
CoT	88.5	–	–	60.9	–	–	51.0	–	–	78.8	–	–	59.2	–	–
SC	<b>91.7</b>	36.5	17.7	<b>71.4</b>	191.3	176.4	<b>65.3</b>	209.9	205.2	<b>87.0</b>	53.3	61.7	63.0	70.8	54.8
AC	91.6	23.8	19.0	71.2	<b>126.8</b>	165.2	64.9	<b>151.5</b>	197.9	86.4	39.5	54.5	62.5	53.2	56.0
ESC	<b>91.7</b>	27.0	21.5	71.3	138.4	186.6	65.2	168.5	222.6	86.9	43.1	62.5	<b>63.4</b>	58.8	62.1
SeerSC	90.6	<b>19.1</b>	<b>9.5</b>	70.7	156.8	<b>142.2</b>	64.3	183.7	<b>178.2</b>	84.8	<b>32.8</b>	<b>35.7</b>	62.5	<b>37.6</b>	<b>29.8</b>
<i>Qwen3-4B (non-thinking)</i>															
CoT	83.3	–	–	23.0	–	–	18.7	–	–	67.2	–	–	48.7	–	–
SC	<b>89.1</b>	7.2	1.8	39.1	98.2	36.0	22.6	49.2	17.3	76.8	13.3	4.0	<b>57.6</b>	24.4	6.5
AC	89.0	5.5	3.1	39.2	<b>81.3</b>	59.9	22.2	<b>44.5</b>	32.8	<b>76.9</b>	11.5	7.8	57.3	16.2	11.3
ESC	89.0	6.0	3.6	<b>39.8</b>	85.6	68.4	22.4	46.7	34.7	76.6	12.0	8.0	<b>57.6</b>	18.2	11.8
SeerSC	88.1	<b>4.4</b>	<b>1.5</b>	38.9	93.6	<b>33.6</b>	<b>22.9</b>	47.8	<b>16.1</b>	76.3	<b>9.5</b>	<b>3.3</b>	57.1	<b>13.0</b>	<b>3.7</b>

less than 1% of the total latency of SeerSC. In practical applications, the overhead of System 1 would be even smaller, since the KV cache produced during its prefilling stage can be reused in System 2. As a result, the additional delay corresponds only to decoding a few tokens and is therefore negligible (see Appendix C.2 for a evaluation of Time to First Token).

**SeerSC exhibits better scaling on the latency–accuracy trade-off** We conduct latency scaling experiments on Qwen3-4B (non-thinking) and DeepSeek-R1-Distill-Qwen-7B using AIME-2024 and GPQA-D, evaluating SeerSC against baseline methods. As illustrated in Figure 5, the CoT method serves as the starting point of extended performance, where allocating more computation time to the model enables us to examine the scaling behavior of each method.

The results show that on AIME2024, a challenging mathematical reasoning benchmark, SeerSC exhibits scaling performance comparable to SC, while achieving slightly higher accuracy under low-latency conditions. Moreover, both SeerSC and SC reduce latency by about 50% compared to AC and ESC on this task. On GPQA-D, a graduate-level scientific reasoning task, SeerSC demonstrates superior scaling performance. Specif-

ically, on DeepSeek-R1-Distill-Qwen-7B, SeerSC reaches the convergence accuracy of SC with only 50% of the latency. Compared to AC and ESC, SeerSC achieves the same accuracy with a 75% reduction in latency.

In summary, across different models and datasets, SeerSC consistently exhibits favorable latency scaling properties, significantly outperforming the sequential methods AC and ESC.

## 5.2 Ablation Study

**Temperature in System 1 can balance accuracy and latency.** The sampling temperature used in System1 primarily affects the distribution of anticipatory entropy, which in turn influences the load allocation of System 2. As illustrated in Figure 6(left), a low temperature shifts the entropy distribution toward lower values, reducing latency but also lowering accuracy. In contrast, a high temperature shifts the distribution toward higher values, increasing latency while improving accuracy. Therefore, selecting an appropriate temperature is crucial for achieving a favorable trade-off.

**Larger sampling size in System 1 leads to more accurate entropy estimation.** According to the Monte Carlo sampling principle, increasing the number of samples allows the estimated entropy




Figure 5: Latency scaling results of SeerSC compared with baseline methods on AIME-2024 and GPQA-D using Qwen3-4B (non-thinking) and DeepSeek-R1-Distill-Qwen-7B. SeerSC matches or exceeds SC under low-latency conditions, reduces latency by about 50% compared to AC and ESC on AIME-2024, and achieves the same accuracy with up to 75% lower latency on GPQA-D.




Figure 6: Ablation study on System 1 using DeepSeek-R1-Distill-Qwen-7B on MATH500. (Left) effect of sampling temperature; (Right) effect of sampling size.

Table 2: Effect of confidence-weighted entropy estimation in System 1 using DeepSeek-R1-Distill-Qwen-7B on MATH500 and AIME-2024.

Method	MATH500		AIME2024	
	Acc.↑	Lat.↓	Acc.↑	Lat.↓
w. Confidence Weighted	<b>93.0</b>	<b>5.3</b>	<b>74.11</b>	<b>151.8</b>
w.o. Confidence Weighted	92.2	5.6	73.56	156.5

distribution to approach its expected value, thereby improving stability. As shown in Figure 6(right), with more samples, accuracy gradually increases while latency decreases. However, once an appropriate sampling size is reached, further increasing it provides little benefit and instead adds extra latency to System 1.

**Incorporating confidence-weighted entropy estimation improves performance.** The Shannon entropy distribution in System 1 is constrained by the number of answer categories, causing entropy values to accumulate around discrete points, which is suboptimal for load allocation in System 2. By integrating confidence weighting, the entropy distribution becomes smoother and more continu-

ous, while also increasing the contribution of high-confidence answers in the entropy calculation. As shown in Table 2, downstream task results indicate that confidence-weighted entropy estimation not only improves performance but also moderately reduces latency.

### 5.3 Orthogonality

It is worth noting that SeerSC is orthogonal to several existing parallel test-time scaling optimization techniques, such as weighted voting and path pruning. We conducted experiments to validate this claim. For extended experiments regarding orthogonality, please refer to Appendix C.3 and C.4.

**SeerSC is orthogonal to weighted vote.** Existing weighted voting assigns confidence-based weights to paths, allowing high-confidence samples to exert greater influence and improve accuracy. (Wan et al., 2025; Zhou et al., 2025). Building on standard SeerSC, we incorporated the weighted voting strategy from (Fu et al., 2025), which leverages tail confidence—the confidence scores of the last window in a path—to determine its weight in voting. As shown in Table 3, both DeepSeek-R1-Distill-Qwen-7B and Qwen3-4B (non-thinking) achieved accuracy improvements across datasets. Notably, on AIME-2024, Qwen3-4B (non-thinking) outperformed vanilla SeerSC by +3 percentage points. Moreover, combining the two methods does not affect token consumption or latency efficiency.

**SeerSC is orthogonal to path pruning.** Existing path pruning approaches generally prune low-quality reasoning paths during parallel generation in order to reduce token usage and improve voting

Table 3: Orthogonality study of SeerSC with weighted vote and path pruning on DeepSeek-R1-Distill-Qwen-7B and Qwen3-4B (non-thinking). We report accuracy (%), average tokens ( $\times 10^3$ ), and average latency (s) on five datasets. Results show that SeerSC can be effectively combined with these strategies, improving accuracy or efficiency without introducing additional latency.

Methods	MATH500			AIME2024			AIME2025			AMC2023			GPQA-D		
	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓
<i>DeepSeek-R1-Distill-Qwen-7B</i>															
SeerSC	93.0	17.7	<b>5.3</b>	74.1	315.6	151.8	60.2	356.0	<b>173.8</b>	91.5	33.5	<b>9.4</b>	48.6	95.3	<b>37.8</b>
+ Weighted Vote	<b>93.8</b>	17.8	<b>5.3</b>	<b>75.7</b>	315.8	152.0	<b>60.5</b>	357.1	174.0	<b>92.2</b>	<b>33.4</b>	<b>9.4</b>	<b>49.0</b>	95.7	38.1
+ Path Pruning	93.4	<b>17.5</b>	<b>5.3</b>	74.2	<b>306.1</b>	<b>151.7</b>	58.3	<b>338.9</b>	<b>173.8</b>	91.5	<b>33.4</b>	<b>9.4</b>	48.2	<b>91.1</b>	<b>37.8</b>
<i>Qwen3-4B (non-thinking)</i>															
SeerSC	88.1	<b>4.4</b>	1.5	38.9	93.6	33.6	22.9	47.8	<b>16.1</b>	76.3	9.5	<b>3.3</b>	57.1	13.0	<b>3.7</b>
+ Weighted Vote	<b>88.8</b>	4.5	<b>1.4</b>	<b>41.9</b>	93.7	<b>33.4</b>	<b>23.9</b>	47.8	16.2	<b>77.9</b>	9.6	<b>3.3</b>	<b>57.5</b>	13.2	3.8
+ Path Pruning	88.2	<b>4.4</b>	1.5	38.9	<b>87.1</b>	33.9	22.7	<b>44.8</b>	<b>16.1</b>	76.8	<b>9.4</b>	<b>3.3</b>	56.5	<b>11.8</b>	3.9

accuracy (Wang et al., 2025a; Hong et al., 2025). On top of standard SeerSC, we adopted the pruning strategy from (Fu et al., 2025), which measures the minimum confidence along a path. If this value falls below a fixed threshold, the path is pruned. Specifically, the pruning threshold is set to 7 for DeepSeek-R1-Distill-Qwen-7B and 10.5 for Qwen3-4B (non-thinking). As shown in Table 3, both models achieved improved token efficiency on AIME-2024/2025 and GPQA-D, and in most tasks, a slight improvement in voting accuracy was also observed.

## 6 Related Work

### 6.1 Test-Time Scaling

Test-Time Scaling (TTS) allocates extra computation at inference to improve performance (Snell et al., 2024), without parameter updates, and can be broadly categorized into *sequential* and *parallel* methods (Zhang et al., 2025).

Sequential TTS refines a single candidate solution step by step. Chain-of-Thought (Wei et al., 2022) is the most representative approach, where intermediate reasoning steps are explicitly generated to improve accuracy and interpretability. Extensions such as Self-Refine (Madaan et al., 2023) further iteratively improve answers, though excessive reasoning can sometimes hurt performance (Chen et al., 2024b).

Parallel TTS generates multiple candidate solutions in parallel and aggregates them to form the final answer. Self-Consistency (Wang et al., 2022) selects the majority outcome among sampled paths, while Best-of- $N$  (BoN) (Lightman et al., 2023) uses a reward model to select the best candidate. These methods improve robustness but often face diminishing returns as the sample size

increases (Snell et al., 2024).

### 6.2 Efficient Optimization of Parallel TTS.

Recent work has focused on improving the efficiency of TTS (Zeyu et al., 2025; Xu et al., 2025). Existing approaches to optimizing parallel TTS can be broadly grouped into three categories:

(1) *Adaptive sampling*. These methods dynamically adjust the number of samples  $N$ . For instance, AC (Aggarwal et al., 2023) and ESC (Li et al., 2024) stop sampling early once a consensus emerges among recent answers; DSC (Wang et al., 2024) uses prompting to decide the sample size; RASC (Wan et al., 2025) trains a classifier to decide early stopping. While these sequential sampling methods hinders GPU parallelization and introduces additional latency (Hong et al., 2025).

(2) *Path pruning*. Another line of work prunes low-quality reasoning paths during generation to reduce tokens and improve aggregation accuracy. For example, Fast BoN (Sun et al., 2024) employs a reward model to retain promising trajectories, ST-BoN (Wang et al., 2025a) measures pairwise distances between sampled paths to prune inconsistent ones, and Slim-SC (Hong et al., 2025) prunes similar or redundant trajectories.

(3) *Hybrid methods*. Some methods combine adaptive sampling and pruning. For instance, RPC (Zhou et al., 2025) uses perplexity-based pruning with dynamic sampling; DeepThink (Fu et al., 2025) incorporates multiple confidence metrics to prune low-confidence paths, adaptively decides whether to stop sampling, and applies weighted voting to improve accuracy.

## 7 Conclusion

We proposed **SeerSC**, a dynamic self-consistency framework. Compared with sequential adaptive methods, SeerSC achieves lower latency without sacrificing accuracy, and it remains orthogonal and complementary to confidence-weighted voting and path pruning, enabling further performance gains.

## Limitations

Although the proposed method demonstrates advantages in both token consumption and latency, there remains scope for further improvement. In domains with verifiable outcomes (e.g., mathematics), the answer entropy from System 1 is directly accessible. However, in broader domains like code generation, existing entropy-based dynamic TTS methods do not provide reliable estimation. Our future work will focus on extending the application of dynamic test-time scaling to a wider range of scenarios by leveraging techniques such as clustering and semantic vectors.

## Acknowledge

We gratefully acknowledge the support of the National Natural Science Foundation of China (NSFC) via grant 62236004 and 62476073.

## References

- Pranjal Aggarwal, Aman Madaan, Yiming Yang, and 1 others. 2023. Let’s sample step by step: Adaptive-consistency for efficient reasoning and coding with llms. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12375–12396.
- Art of Problem Solving. 2023. 2023 AMC. [https://artofproblemsolving.com/wiki/index.php/2023\\_AMC\\_12A\\_Problems](https://artofproblemsolving.com/wiki/index.php/2023_AMC_12A_Problems).
- Art of Problem Solving. 2024. 2024 AIME. [https://artofproblemsolving.com/wiki/index.php/AIME\\_Problems\\_and\\_Solutions](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions).
- Art of Problem Solving. 2025. 2025 AIME. [https://artofproblemsolving.com/wiki/index.php/AIME\\_Problems\\_and\\_Solutions](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions).
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. 2024. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787*.
- Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Ion Stoica, Matei Zaharia, and James Zou. 2024a. Are more llm calls all you need? towards scaling laws of compound inference systems. *arXiv preprint arXiv:2403.02419*.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. 2025. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, and 1 others. 2024b. Do not think that much for  $2+3=?$  on the overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*.
- Sicheng Feng, Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025. Efficient reasoning models: A survey. *arXiv preprint arXiv:2504.10903*.
- Yichao Fu, Xuwei Wang, Yuandong Tian, and Jiawei Zhao. 2025. Deep think with confidence. *arXiv preprint arXiv:2508.15260*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Colin Hong, Xu Guo, Anand Chanaan Singh, Esha Choukse, and Dmitrii Ustiugov. 2025. Slim-sc: Thought pruning for efficient scaling with self-consistency. *arXiv preprint arXiv:2509.13990*.
- Jenny Y Huang, Mehul Damani, Yousef El-Kurdi, Ramon Astudillo, and Wei Sun. 2025. Latency and token-aware test-time compute. *arXiv preprint arXiv:2509.09864*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Xinglin Wang, Bin Sun, Heda Wang, and Kan Li. 2024. Escape sky-high cost: Early-stopping self-consistency for multi-step reasoning. In *The Twelfth*

- International Conference on Learning Representations*.
- Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, and 1 others. 2025. From system 1 to system 2: A survey of reasoning large language models. *arXiv preprint arXiv:2502.17419*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiakai Liu, Chaojie Wang, Rui Yan, Wei Shen, Fuxiang Zhang, Jiacheng Xu, Yang Liu, and Yahui Zhou. 2025a. Skywork-reward-v2: Scaling preference data curation via human-ai synergy. *arXiv preprint arXiv:2507.01352*.
- Yue Liu, Jiaying Wu, Yufei He, Ruihan Gong, Jun Xia, Liang Li, Hongcheng Gao, Hongyu Chen, Baolong Bi, Jiaheng Zhang, and 1 others. 2025b. Efficient inference for large reasoning models: A survey. *arXiv preprint arXiv:2503.23077*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, and 1 others. 2023. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- Hanshi Sun, Momin Haider, Ruiqi Zhang, Huitao Yang, Jiahao Qiu, Ming Yin, Mengdi Wang, Peter Bartlett, and Andrea Zanette. 2024. Fast best-of-n decoding via speculative rejection. *Advances in Neural Information Processing Systems*, 37:32630–32652.
- Qwen Team. 2025. [Qwq-32b: Embracing the power of reinforcement learning](#).
- Guangya Wan, Yuqi Wu, Jie Chen, and Sheng Li. 2025. Reasoning aware self-consistency: Leveraging reasoning paths for efficient llm sampling. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3613–3635.
- Xinglin Wang, Shaoxiong Feng, Yiwei Li, Peiwen Yuan, Yueqi Zhang, Chuyi Tan, Boyuan Pan, Yao Hu, and Kan Li. 2024. Make every penny count: Difficulty-adaptive self-consistency for cost-efficient reasoning. *arXiv preprint arXiv:2408.13457*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Yiming Wang, Pei Zhang, Siyuan Huang, Baosong Yang, Zhuosheng Zhang, Fei Huang, and Rui Wang. 2025a. Sampling-efficient test-time scaling: Self-estimating the best-of-n sampling in early decoding. *arXiv preprint arXiv:2503.01422*.
- Zili Wang, Tianyu Zhang, Lei Zhu, Haoli Bai, Lu Hou, Shiming Xiang, Xianzhi Yu, and Wulong Liu. 2025b. Faster and better llms via latency-aware test-time scaling. *arXiv preprint arXiv:2505.19634*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. 2025. Tokenskip: Controllable chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*.
- Zenan Xu, Zexuan Qiu, Guanhua Huang, Kun Li, Siheng Li, Chenchen Zhang, Kejiao Li, Qi Yi, Yuhao Jiang, Bo Zhou, and 1 others. 2025. Adaptive termination for multi-round parallel reasoning: An universal semantic entropy-guided framework. *arXiv preprint arXiv:2507.06829*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, and 1 others. 2024. Qwen2. 5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*.
- XING Zeyu, Xing Li, Huiling Zhen, Xianzhi Yu, Mingxuan Yuan, and Sinno Jialin Pan. 2025. Large reasoning models know how to think efficiently. In *ES-FoMo III: 3rd Workshop on Efficient Systems for Foundation Models*.
- Qiyuan Zhang, Fuyuan Lyu, Zexu Sun, Lei Wang, Weixu Zhang, Wenye Hua, Haolun Wu, Zhihan Guo, Yufei Wang, Niklas Muennighoff, and 1 others. 2025. A survey on test-time scaling in large language models: What, how, where, and how well? *arXiv preprint arXiv:2503.24235*.

Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Livia Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, and 1 others. 2024. Sglang: Efficient execution of structured language model programs. *Advances in neural information processing systems*, 37:62557–62583.

Zhi Zhou, Tan Yuhao, Zenan Li, Yuan Yao, Lan-Zhe Guo, Xiaoxing Ma, and Yu-Feng Li. 2025. Bridging internal probability and self-consistency for effective and efficient llm reasoning. *arXiv preprint arXiv:2502.00511*.

## A Implementation Details

### A.1 Optimal $N$ for SC

Table 4 summarizes the optimal sampling number  $N$  for the SC across models and datasets.

### A.2 Implementation details of SeerSC

For DeepSeek-R1-Distill-Qwen-7B and DeepSeek-R1-Distill-Llama-8B, System 1 is run with a generation temperature of 0.5, while for Qwen3-4B (thinking) and Qwen3-4B (non-thinking), System 1 is run with the temperature of 1.0. For all models and datasets, the number of direct answers in System 1 is fixed at 64.

### A.3 System 1 Implementation for Reasoning Models

Reasoning models (e.g., DeepSeek-R1 series or QwQ) typically generate an extended reasoning process enclosed within `<think>` and `</think>` tags before producing the final answer. However, System 1 requires rapid, direct answer generation to estimate entropy efficiently.

To enable this, we employ a forced decoding strategy. Specifically, we append a sequence to the model input that simulates the completion of the thinking process and prompts the immediate start of the final answer. The appended sequence is:

```
\n</think>So the final answer is \
```

By forcing the model to process this sequence immediately after the user query, we bypass the generation of the reasoning trace. The model proceeds as if it has already completed its reasoning, directly generating the content within `\boxed{}`.

**Example.** Given the input query:

```
<|User|>How many positive... within
\boxed{ }. <|Assistant|><think>
```

A standard generation would produce:

```
(thinking...)</think>So the final answer
is \boxed{1}.
```

With our suffix injection, the effective input becomes:

```
<|User|>...<|Assistant|><think>\n</think>So
the final answer is \
```

And the model output becomes:

```
\boxed{1}.
```

This technique ensures that System 1 inference remains computationally lightweight and fast.

## B Hyperparameter Sensitivity Analysis

The entropy thresholds  $\tau_1$  and  $\tau_2$  in Equation 5 are critical hyperparameters that determine the trade-off between accuracy and latency. We empirically determined these values based on performance on the MATH training dataset using the DeepSeek-R1-Distill-Qwen-7B model.

**Impact of the First Threshold ( $\tau_1$ ).** We first investigated the impact of a single threshold  $\tau_1$  (effectively setting  $\tau_2 = \tau_1$ , resulting in a binary budget allocation of either 1 or  $N$ ). As shown in Table 5a, increasing  $\tau_1$  leads to more samples being assigned a budget of 1. Consequently, the method transitions from the behavior of SC (high accuracy, high latency) to that of CoT (low accuracy, low latency). A threshold of  $1/10$  provides a strong initial filtering, significantly reducing latency while preserving most of the accuracy gains.

**Impact of the Second Threshold ( $\tau_2$ ).** To achieve a finer-grained budget allocation, we introduced the second threshold  $\tau_2$  (assigning a budget of  $N/2$  for samples with entropy between  $\tau_1$  and  $\tau_2$ ). We fixed  $\tau_1 = 1/10$  and varied  $\tau_2$ . As presented in Table 5b, introducing  $\tau_2$  further optimizes the efficiency. The setting ( $\tau_1 = 1/10, \tau_2 = 1/3$ ) achieves a significant reduction in latency (from 7.4s to 5.3s) compared to using a single threshold, with only a negligible drop in accuracy (0.3%).

**Summary.** These empirical results justify our choice of  $\tau_1 = \frac{1}{10} \log M$  and  $\tau_2 = \frac{1}{3} \log M$ . While this configuration may not be the global optimum for every specific scenario, our extensive experiments across diverse datasets and model scales (ranging from 4B to 32B) demonstrate its consistent effectiveness and robustness. We adopt these fixed values to avoid the additional computational overhead associated with per-task hyperparameter

Table 4: Optimal  $N$  for SC on different models and datasets.

Model	MATH500	AIME2024	AIME2025	AMC2023	GPQA-D
DeepSeek-R1-Distill-Qwen-7B	8	32	32	8	16
DeepSeek-R1-Distill-Llama-8B	8	32	32	8	8
Qwen3-4B (thinking)	8	16	16	8	8
Qwen3-4B (non-thinking)	8	32	16	8	16

Table 5: Hyperparameter Sensitivity Analysis on the MATH training dataset using DeepSeek-R1-Distill-Qwen-7B. We identify the optimal configuration step-by-step: first determining  $\tau_1$ , then tuning  $\tau_2$ .

(a) Impact of the First Threshold  $\tau_1$  (single threshold setting).  $\tau_1 = 1/10$  offers a favorable balance between SC and CoT behaviors.

Metric	(0, -)	(1/20, -)	(1/10, -)	(1/5, -)	(1/3, -)	(1, -)
Acc. (%)	93.9	93.8	<b>93.3</b>	92.8	92.0	90.0
Lat. (s)	8.4	7.8	<b>7.4</b>	5.8	4.2	1.1

(b) Impact of the Second Threshold  $\tau_2$  (with  $\tau_1$  fixed at  $1/10$ ). The configuration (1/10, 1/3) further reduces latency with negligible accuracy loss.

Metric	(1/10, -)	(1/10, 1/6)	(1/10, 1/5)	(1/10, 1/3)	(1/10, 1/2)	(1/10, 1)
Acc. (%)	93.3	93.2	93.0	<b>93.0</b>	92.4	92.0
Lat. (s)	7.4	6.6	6.1	<b>5.3</b>	4.3	3.6

searching on validation sets. The exploration of more fine-grained or adaptive thresholding mechanisms remains a promising avenue for our future work.

## C Additional Experiments

### C.1 Generalizability to Larger Models

To assess generalizability, we extended our evaluation to the larger-scale **QwQ-32B** model with a sampling budget of  $N = 8$ . As shown in Table 6, the core conclusions observed in smaller models remain valid: SeerSC maintains accuracy parity with the SC baseline while significantly reducing computational costs. This confirms that SeerSC’s adaptive budget allocation remains robust and effective at larger model scales.

### C.2 Impact on Time to First Token (TTFT)

We further investigated the impact of the System 1 pre-estimation stage on the Time to First Token (TTFT), a critical metric for streaming applications. We measured the TTFT on a single AIME2024 query using *DeepSeek-R1-Distill-Qwen-7B* served by vLLM.

As shown in Table 7, although SeerSC requires a pre-estimation step (generating 64 direct answers), the increase in TTFT is minimal. Specifically, the TTFT increases from 0.022s (SC) to 0.234s (SeerSC), adding an overhead of only  $\approx 0.2$ s.

Given that the average completion time for a complex reasoning task (e.g., AIME) is typically around 30 seconds, this slight delay is negligible relative to the total inference latency and does not perceptibly affect the user experience.

### C.3 Impact of Weighted Voting on Baselines

We extend the confidence-based weighted voting (Fu et al., 2025) to baselines (SC, AC, ESC) on DeepSeek-R1-Distill-Qwen-7B to evaluate potential gains. As shown in Table 8, while weighted voting yields marginal accuracy improvements for baselines, it fails to address their inherent latency bottlenecks. In contrast, **SeerSC maintains a significant efficiency advantage**—e.g., on GPQA-D, it is  $1.4\times$  faster than SC and  $\approx 2\times$  faster than AC/ESC—confirming that our efficiency gains stem from proactive budget allocation rather than the voting mechanism.

### C.4 Compatibility with Other TTS Paradigms

In Section 5.3, we demonstrated that SeerSC is orthogonal to parallel optimization techniques such as weighted voting and path pruning. Beyond these strategies, SeerSC serves as a flexible budget allocation framework that can also be integrated with broader Test-Time Scaling (TTS) paradigms. To demonstrate this extensibility, we conducted experiments combining SeerSC with Reward Model-based selection (Best-of-N) and Length-Controlled

Table 6: Experimental results on **QwQ-32B** model. SeerSC demonstrates strong generalizability, achieving significant reductions in token usage and latency on diverse datasets, while maintaining accuracy comparable to baselines. We report accuracy (%), average tokens ( $\times 10^3$ ), and average latency (s).

Methods	MATH500			AIME2024			AIME2025			AMC2023			GPQA-D		
	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓	Acc.↑	Token↓	Lat.↓
<i>QwQ-32B</i>															
CoT	91.1	–	–	65.1	–	–	51.8	–	–	81.1	–	–	63.5	–	–
SC	<b>93.6</b>	33.3	16.2	74.0	93.4	95.3	66.0	101.3	87.6	86.0	51.7	29.6	<b>66.4</b>	69.6	58.0
AC	93.5	21.0	19.1	73.3	<b>73.6</b>	135.2	65.3	<b>91.4</b>	123.7	86.0	36.9	41.2	66.1	<b>53.9</b>	70.5
ESC	93.4	23.9	22.2	73.6	78.5	127.1	65.1	93.8	115.5	86.5	40.4	45.8	66.1	58.6	77.9
SeerSC	93.0	<b>19.5</b>	<b>9.6</b>	<b>74.4</b>	88.3	<b>90.6</b>	<b>66.2</b>	95.4	<b>83.2</b>	85.5	<b>36.7</b>	<b>19.1</b>	66.3	59.7	<b>43.4</b>

Table 7: Comparison of Time to First Token (TTFT) between SC and SeerSC on a single AIME query. The overhead introduced by System 1 is negligible ( $\approx 0.2$ s).

Method	TTFT (s)	Overhead (s)
SC	0.022	–
SeerSC	0.234	+0.212

scaling (TokenSkip).

**Combining with Best-of-N.** We integrated SeerSC with Best-of-N (BoN) using DeepSeek-R1-Distill-Qwen-7B as the generator and Skywork-Reward-V2-Qwen3-8B (Liu et al., 2025a) as the reward model. Instead of majority voting, we used the reward model to rank the candidates generated within the SeerSC budget.

As shown in Table 9(a), this combination yields consistent accuracy improvements, particularly on the challenging GPQA-D benchmark (+2.4%). However, the introduction of a reward model incurs additional latency for scoring, which is a trade-off users must consider.

**Combining with Length-Controlled Scaling.** We also explored combining SeerSC with TokenSkip (Xia et al., 2025), a sequential scaling method designed to reduce token usage for reasoning models. We used a Qwen2.5-7B-Instruct model fine-tuned with TokenSkip (length parameter 0.7).

As shown in Table 9(b), combining SeerSC with TokenSkip significantly reduces token consumption (up to 16% on GSM8K). This comes with a slight accuracy drop, which is primarily attributable to the performance degradation inherent in the TokenSkip fine-tuning process.

Table 8: Impact of Weighted Voting on baselines (DeepSeek-R1-Distill-Qwen-7B). Values in parentheses denote changes relative to vanilla versions. Despite accuracy gains in baselines, SeerSC retains superior latency efficiency.

Methods	MATH500		AMC2023		GPQA-D	
	Acc. (%)↑	Lat. (s)↓	Acc. (%)↑	Lat. (s)↓	Acc. (%)↑	Lat. (s)↓
SC	93.9	8.4	91.6	13.2	48.1	51.8
+ Weighted Vote	94.1 (+0.2)	8.4	<b>92.5</b> (+0.9)	13.2	48.7 (+0.6)	52.1
AC	93.8	12.4	91.3	24.8	47.7	73.1
+ Weighted Vote	<b>94.0</b> (+0.2)	12.5	92.3 (+1.0)	25.2	47.7 (+0.0)	73.5
ESC	93.6	13.9	91.1	26.7	48.0	82.0
+ Weighted Vote	93.4 (-0.2)	13.8	92.1 (+1.0)	26.7	48.6 (+0.6)	82.3
SeerSC	93.0	<b>5.3</b>	91.5	<b>9.4</b>	48.6	<b>37.8</b>
+ Weighted Vote	93.8 (+0.8)	<b>5.3</b>	92.2 (+0.7)	<b>9.4</b>	<b>49.0</b> (+0.4)	38.1

Table 9: Experimental results of combining SeerSC with other TTS paradigms. The results demonstrate that SeerSC is compatible with both reward model-based and length-controlled methods.

(a) Combination with **Best-of-N** (Generator: DeepSeek-R1-Distill-Qwen-7B; RM: Skywork-Reward-V2-Qwen3-8B). BoN improves accuracy, at the cost of scoring latency.

Method	MATH500			GPQA-D		
	Acc. (%)	Token ( $10^3$ )	Lat. (s)	Acc. (%)	Token ( $10^3$ )	Lat. (s)
SeerSC	93.0	17.7	<b>5.3</b>	48.6	95.3	<b>37.8</b>
+ Best-of-N	<b>93.4</b> (+0.4)	17.7	7.3	<b>51.0</b> (+2.4)	95.3	48.7

(b) Combination with **TokenSkip** (Model: Qwen2.5-7B-Instruct fine-tuned with TokenSkip). TokenSkip reduces token usage significantly, demonstrating orthogonality with SeerSC.

Method	GSM8K			MATH500		
	Acc. (%)	Token	Lat. (s)	Acc. (%)	Token	Lat. (s)
SeerSC	92.9	1437	0.8	<b>78.6</b>	2726	0.5
+ TokenSkip	91.0 (-1.9)	<b>1205</b> (-16%)	0.8	76.0 (-2.6)	<b>2379</b> (-13%)	0.5