

# $\phi$ -Decoding: Adaptive Foresight Sampling for Balanced Inference-Time Exploration and Exploitation

Fangzhi Xu<sup>1,2,6\*</sup> Hang Yan<sup>1,7\*</sup> Chang Ma<sup>3</sup> Haiteng Zhao<sup>4</sup>

Jun Liu<sup>1,6†</sup> Qika Lin<sup>5†</sup> Zhiyong Wu<sup>2†</sup>

<sup>1</sup>School of Computer Science and Technology, Xi'an Jiaotong University

<sup>2</sup>Shanghai AI Lab <sup>3</sup>The University of Hong Kong

<sup>4</sup>Peking University <sup>5</sup>National University of Singapore

<sup>6</sup>Ministry of Education Key Laboratory of Intelligent Networks and Network Security

<sup>7</sup>Shaanxi Province Key Laboratory of Big Data Knowledge Engineering  
{fangzhixu98, whucs2013wzy}@gmail.com liukeen@xjtu.edu.cn

## Abstract

Inference-time optimization scales computation to derive deliberate reasoning steps for effective performance. While previous search-based strategies address the short-sightedness of auto-regressive generation, the vast search space leads to excessive *exploration* and insufficient *exploitation*. To strike an efficient balance to derive the optimal step, we frame the decoding strategy as *foresight sampling*, leveraging simulated future steps to obtain globally optimal step estimation. Built on it, we propose a novel decoding strategy, named  $\phi$ -Decoding. To provide a precise and expressive estimation of step value,  $\phi$ -Decoding approximates two distributions via foresight and clustering. Sampling from the joint distribution, the optimal steps can be selected for exploitation. To support adaptive computation allocation, we propose in-width and in-depth pruning strategies, featuring a light-weight solution to achieve inference efficiency. Extensive experiments across seven benchmarks show  $\phi$ -Decoding outperforms strong baselines in both performance and efficiency. Additional analysis demonstrates its generalization across various LLMs and scalability across a wide range of computing budgets.<sup>1</sup>

## 1 Introduction

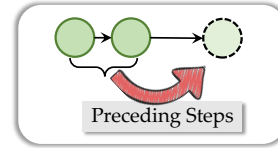
Large language models (LLMs) (Achiam et al., 2023; Team et al., 2023) present remarkable performances in solving reasoning-intensive tasks through step-by-step thoughts (Wei et al., 2022). Recent advancements (Team, 2024; Guo et al., 2025) have significantly boosted LLM reasoning by large-scale post-training on well-curated datasets.

\*means equal contribution. Work done during Fangzhi's internship at Shanghai AI Lab.

†denotes corresponding author.

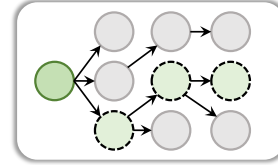
<sup>1</sup>The code is released at <https://github.com/xufangzhi/phi-Decoding>.

(a) Auto-Regressive



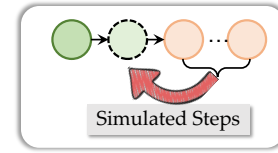
- ✓ Efficiency
- ✗ Performance
- ✗ Global Optima

(b) Search-Based



- ✗ Efficiency
- ✓ Performance
- ✓ Global Optima

(c) Foresight Sampling



- ✓ Efficiency
- ✓ Performance
- ✓ Global Optima

Figure 1: Comparisons of different decoding paradigms. (a) is auto-regressive decoding, which has high efficiency but lacks global awareness. (b) represents search-based methods, which requires huge search space with extensive time cost. (c) is the foresight sampling strategy. It leverages the simulated future steps to estimate the step value, which can strike a balanced inference-time exploration and exploitation.

Nevertheless, the cost associated with the post-training procedure hinders its reproducibility. This naturally motivates us to explore the inference-time strategy for optimizing the LLM reasoning chains.

Inference-time optimization involves employing more reasoning tokens that encode thinking steps to perform effective reasoning. However, the natural shortsightedness of auto-regressive generation, which predicts the next step only with preceding steps, makes most inference algorithms unable to achieve global optima (Ma et al., 2024) (Fig. 1(a)). Most previous work solves this by deliberately optimizing each step using search-based methods (Yao

et al., 2024; Hao et al., 2023; Xie et al., 2024; Wu et al., 2024), the expanding and backtracking of tree search algorithms enable LLMs to find global-optimal reasoning paths (Fig. 1(b)). However, the vast search space results in excessive *exploration* and insufficient *exploitation*. Conversely, if we could derive a precise estimation of globally-aware step values, an efficient balance between inference-time exploration and exploitation could be achieved.

Based on this, we frame the decoding strategy as *foresight sampling*, as depicted in Fig. 1(c). It relies on the future simulation to obtain the globally optimal estimation of the current step. Central to the foresight sampling is the critical task of *how to estimate step value with the foresight steps*. Intuitively, the step estimation with foresight can be derived either by incorporating the process reward model (PRM) (Snell et al., 2024) or through model uncertainty (Ma et al., 2024). However, PRMs are not widely available for all reasoning scenarios, which poses challenges for scalability. Delegating the step assessment to model uncertainty risks the issue of local optima, potentially resulting in suboptimal performance.

Another issue in stepwise exploration and exploitation is *whether every step requires deliberation for decision-making*. Naturally, more computational resources should be allocated to challenging steps, while conserving compute for simpler steps. Previous inference-time optimization methods widely overlook this principle. In addition, such concept of *over-thinking* has been widely observed in the o1-like attempts (Chen et al., 2024; Manvi et al., 2024). Therefore, it is both intriguing and promising to develop a light-weight solution that can adaptively balances computational workload without extra training.

In this paper, we propose a novel inference-time optimization algorithm named  *$\phi$ -Decoding*, which introduces an adaptive foresight sampling strategy to achieve efficient exploration and exploitation during inference. To give the reliable and expressive step value estimation,  *$\phi$ -Decoding* capitalizes on foresight paths to derive two distributions: one from the derived step *advantage* values, capturing uncertainty discrepancies between successive steps, and another from *alignment* of these foresight paths via clustering. Sampling from the joint distribution,  *$\phi$ -Decoding* selects optimal steps for exploitation. To efficiently allocate the computa-

tions,  *$\phi$ -Decoding* introduces both the in-width and in-depth pruning strategies, which provides adaptive inference-time scaling.

On diverse reasoning benchmarks,  *$\phi$ -Decoding* improves the average performance of LLaMA3.1-Instruct-8B by >14% over auto-regressive CoT. Inference-time scaling across diverse computing budgets shows the consistent superiority of  *$\phi$ -Decoding* over other baselines, offering a balance between performance (Accuracy) and computational efficiency (#FLOPS). Further analysis of the generalization across various backbone LLMs and scalability to the competition-level task highlights the superiority of  *$\phi$ -Decoding*.

The major contributions of our work are:

- (1) **An adaptive inference-time optimization algorithm  *$\phi$ -Decoding* without external auxiliary.**  *$\phi$ -Decoding* estimates the step value based on the joint distribution derived from foresight paths. In-width and in-depth pruning strategies are introduced to alleviate the overthinking issue.
- (2) **Extensive experiments with state-of-the-art performances.**  *$\phi$ -Decoding* improves the average reasoning of LLaMA3.1-8B-Instruct by over 14% across various reasoning benchmarks, presenting a great trade-off between effectiveness and efficiency compared with baselines.
- (3) **Comprehensive analysis and insightful findings.**  *$\phi$ -Decoding* proves its generalization ability to various LLMs, ranging from the 70B-sized model to R1-distilled LLM. Additionally, the inference-time scaling across a wide range of computing budgets reveals the consistent advantages, where  *$\phi$ -Decoding* matches the performance of the suboptimal baseline with  $6\times$  efficiency.

## 2 Methodology

### 2.1 Preliminary

In the context of auto-regressive language generation, the selection of the current step  $\hat{a}_t$  is based on the following probability distribution:

$$\hat{a}_t \sim p_\theta(a_t|x, \mathbf{a}_{<t}) \quad (1)$$

where  $x$  is the instruction or the input query, and  $\mathbf{a}_{<t}$  represents the preceding steps.  $\theta$  denotes the LLM parameters, where  $p_\theta$  is derived from the distribution of language modeling.

To overcome the short-sighted limitation of auto-regressive generation and achieve efficient exploration, *foresight sampling* conditions the generation process not only on the preceding steps  $\mathbf{a}_{<t}$

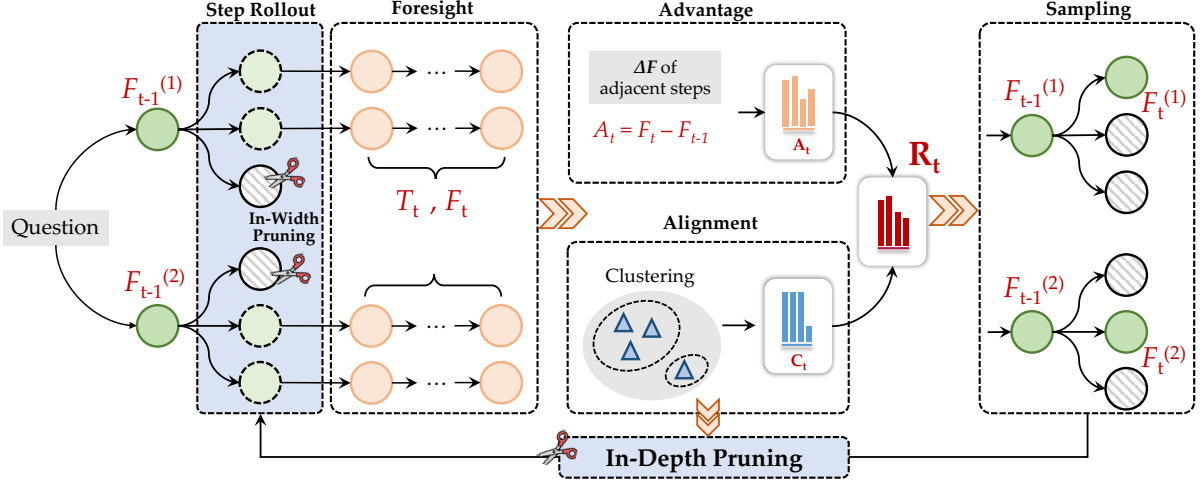


Figure 2: The overall framework of  $\phi$ -Decoding. We visualize the decoding process at the timestamp  $t$ . For simplicity, we set step beam size  $M$  as 2, the number of rollouts  $N$  as 3, and the number of clusters  $K$  as 2.

but also on an estimation of future outcomes  $\mathbf{a}_{>t}$ . We use the Boltzmann distribution to model the probabilities of different outcomes during the decoding process, incorporating both the influence of preceding steps and an estimation of future states, such as:

$$\hat{a}_t \sim p_\theta(a_t|x, \mathbf{a}_{<t})\mathbb{E}_{\mathbf{a}_{>t}}p_\theta(\mathbf{a}_{>t}|x, a_t, \mathbf{a}_{<t}) \quad (2)$$

It is non-trivial to have a precise calculation of  $\mathbb{E}_{\mathbf{a}_{>t}}p_\theta(\mathbf{a}_{>t}|x, a_t, \mathbf{a}_{<t})$ . Therefore, we try to derive an estimation of this future simulation quality.

$$\hat{a}_t \sim p_\theta(a_t|x, \mathbf{a}_{<t})\exp[R(x, \mathbf{a}_{\leq t}, \mathbf{a}_{>t})/\tau] \quad (3)$$

where  $R$  denotes the optimized function for step value estimation based on the foresight steps.  $\tau$  represents the temperature hyper-parameter, which controls the diversity of generation.

Therefore, the ultimate objective of  $\phi$ -Decoding is to design the step value estimation function  $R(x, \mathbf{a}_{\leq t}, \mathbf{a}_{>t})$ . We include the key techniques of  $\phi$ -Decoding in Fig. 2, which depicts the decoding process at the timestamp  $t$ . The complete algorithm as well as the overall decoding pipeline are presented in Appendix B.

## 2.2 Step Value Estimation

To thoroughly optimize the formulation, we propose to evaluate the foresight paths from *advantage* (absolute value) and *alignment* (relative value).

**Dynamic Advantage Estimation.** We follow the beam search strategy. At the timestamp  $t$ , we rollout  $N$  candidate steps from each beam. Based on

the idea of *foresight sampling*, the probability  $F_t$  of the foresight path can be derived:

$$F_t = p_\theta(\mathbf{a}_{>t}|x, a_t, \mathbf{a}_{<t}), \quad (4)$$

where the index of the candidate step is omitted for simplicity.

To measure the advantage brought by the candidate step  $a_t$ , we define the calculation of *Advantage*  $A_t$  as:

$$\begin{aligned} A_t &= p_\theta(\mathbf{a}_{>t}|x, a_t, \mathbf{a}_{<t}) - p_\theta(\mathbf{a}_{>t-1}|x, a_{t-1}, \mathbf{a}_{<t-1}) \\ &= F_t - F_{t-1} \end{aligned} \quad (5)$$

It is represented as the  $\Delta$  of the foresight probability  $F$  between the adjacent steps. Notably, we implement the calculation of  $p_\theta$  with the averaged log probability of the sequence, which alleviates the influence from the foresight length.

Since the calculation of *Advantage* for each candidate step is independent, it estimates the absolute value of the step. For better illustration, we define  $R_1(x, \mathbf{a}_{\leq t}, \mathbf{a}_{>t}) = \exp(A_t/\tau_1)$ .

**Alignment Assessment by Clustering.** One potential risk of the uncertainty-based estimation is the issue of local optima. That is, LLMs may be trapped in the incorrect step with exceptionally high confidence.

To address this limitation, we introduce the definition of *alignment* to provide the relative preference among the foresight paths. This is achieved by employing a clustering strategy following the foresight sampling process. Specifically, the foresight paths at each timestamp are grouped into clusters.

The number of clusters is defined as  $K$ . The *alignment* value of  $a_t$  is determined based on the size of the cluster to which it belongs:

$$C_t = \frac{|\text{Cluster}(a_t)|}{\#\text{Foresight Paths}} \quad (6)$$

where  $|\text{Cluster}(a_t)|$  denotes the size of the cluster  $a_t$  belongs to.

*Alignment* actually provides the relative estimation of the step value, which reflects the consistency among the foresight paths. The more closely the expected outcome aligns with those of other candidates, the greater the step value would be. Similarly, we define  $R_2(x, \mathbf{a}_{\leq t}, \mathbf{a}_{> t}) = \exp(C_t/\tau_2)$ .

**Sampling From Joint Distribution** Combining the rewarding from  $R_1$  and  $R_2$ , we can derive the definition of  $R$  function, which is in the form of:

$$\begin{aligned} R(x, \mathbf{a}_{\leq t}, \mathbf{a}_{> t}) &= \text{Norm}(A_t) + \text{Norm}(C_t) \\ &= \frac{\exp(A_t/\tau_1)}{\sum_{a_t} \exp(A_t/\tau_1)} + \frac{\exp(C_t/\tau_2)}{\sum_{a_t} \exp(C_t/\tau_2)} \end{aligned} \quad (7)$$

Replacing this formulation of  $R$  into Eq. 3, the objective becomes the sampling on the joint distribution of *Advantage* and *Alignment*.

In the implementation, we set  $\tau_1 = \tau_2 = 0.6$  and combine  $R_1$  and  $R_2$  with equal weighting for simplicity. We leave the discussion of the weighted version in future work.

### 2.3 Dynamic Pruning Strategy

To optimize the computation allocation and alleviate the over-thinking issue, we introduce an efficient and effective pruning strategy. It is designed from two dimensions: in-width and in-depth. Figure 2 visualizes the function of the pruning strategies.

**In-Width Pruning.** Although foresight sampling addresses the short-sightedness of language models, it inevitably introduces additional computational cost. Intuitively, some steps with obvious errors can be filtered out directly, without needing to simulate future steps. To achieve this, we assess the generation confidence of each step  $a_t$  based on its probability:

$$s_t = p_\theta(a_t|x, \mathbf{a}_{< t}). \quad (8)$$

There are in total  $M * N$  candidate steps at time-step  $t$ . We then calculate the mean and variance

of these step confidence:

$$\mu_t = \frac{1}{M * N} \sum_i s_t^{(i)}, \quad \sigma_t^2 = \frac{1}{M * N} \sum_i (s_t^{(i)} - \mu)^2 \quad (9)$$

where  $\mu$  and  $\sigma^2$  denote the mean and variance values respectively.  $M * N$  is the number of candidates under the setting of step beam search as defined in Sec. 2.2.

Based on this calculation, we exclude any steps with generation confidence that is exceptionally low, i.e., those with  $s_t^{(i)} < \mu - \sigma$ . The remaining steps are kept for foresight:

$$\mathcal{S}_t = \{a_t^{(i)} | \mu - \sigma \leq s_t^{(i)}\} \quad (10)$$

Adhering to this principle enables the attainment of in-width pruning. Notably, the extent of pruning can be regulated by adjusting the threshold using  $\mu - k\sigma$ , where  $k \in \mathbb{Z}^+$ .

**In-Depth Pruning.** Foresight sampling enables the deliberate thinking of each step. Previous work (Wang and Zhou, 2024) uncovers that the early steps are much more important, necessitating increased computational resources for optimization. As the final answer approaches, LLMs exhibit greater determination in their reasoning paths. Motivated by it, we can save some computational costs with the strategy of early stopping.

To avoid extra computing and make the solution as simple as possible, we employ the clustering result introduced in Sec. 2.2. In detail, we derive the size of the largest cluster, written as  $|\text{Cluster}_{\max}|$ . The condition of early-stopping is controlled by the threshold:

$$\frac{|\text{Cluster}_{\max}|}{\#\text{Foresight Paths}} \geq \delta \quad (11)$$

Then, the LLM completes the remaining reasoning steps under the auto-regressive setting. For convenience, we set  $\delta = 0.7$  for all experiments.

## 3 Experiments

### 3.1 Evaluation Benchmarks and Metrics

**Benchmarks** To comprehensively evaluate the LLM performances on downstream tasks, we mainly include the following 6 representative reasoning benchmarks: GSM8K (Cobbe et al., 2021), MATH-500 (Hendrycks et al., 2021), GPQA (Rein et al., 2023), ReClor (Yu et al.,

2020), LogiQA (Liu et al., 2021), and ARC-Challenge (Clark et al., 2018). Furthermore, we incorporate the competition-level benchmark AIME (Jia, 2024) to highlight the scalability of  $\phi$ -Decoding to address more challenging scenarios.

**Metrics** We report the Pass@1 accuracy (Acc.) for each benchmark. To better illustrate the trade-off between efficiency and performance, the FLOPS metric is also computed, following the definition of (Kaplan et al., 2020). Please refer to Appendix A.1 for more evaluation details.

### 3.2 Baselines and Backbone LLMs

In the experiments, we compare  $\phi$ -Decoding with the following 5 baseline methods.

**Auto-Regressive (CoT).** It produces the chain-of-thought reasoning through the auto-regressive language generation.

**Tree-of-Thought (ToT) (Yao et al., 2024).** It builds a tree structure for a given problem, where each node represents a reasoning step. We use the BFS version as the implementation.

**Monte Carlo Tree Search (MCTS).** It constructs a search tree and dynamically updates the step value via expanding and backtracking. We follow Reasoning as Planning (RaP) (Hao et al., 2023) for implementation.

**Guided Decoding (Xie et al., 2024).** It utilizes self-evaluation at each step to perform a stochastic beam search.

**Predictive Decoding (Ma et al., 2024).** It proposes the look-ahead strategy and leverages Model Predictive Control to reweigh LLM distributions, producing non-myopic language modeling.

For the 6 reasoning benchmarks in the main experiments, all the baseline methods are evaluated on two backbone LLMs: LLaMA3.1-8B-Instruct (Dubey et al., 2024) and Mistral-v0.3-7B-Instruct (Jiang et al., 2023). To assess generalization and scalability, we further evaluate the Qwen2.5-3B (Yang et al., 2024) and LLaMA3.1-70B (Dubey et al., 2024) LLMs, while also boosting Deepseek R1-series LLM (i.e., R1-Distill-LLaMA-8B) (Guo et al., 2025) for competition-level tasks.

All the experiments are implemented on A100 of 80GB VRAM GPUs. The inference process is accelerated by the vLLM engine (Kwon et al., 2023).

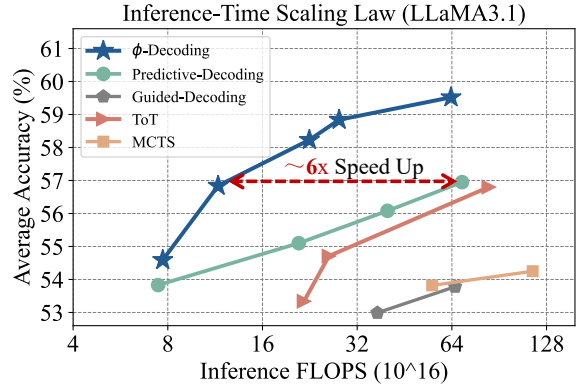


Figure 3: Inference-time scaling law on LLaMA3.1-8B-Instruct. The horizontal axis denotes the inference-time computational cost, while the vertical axis represents the average performances on 6 benchmarks.

The generation temperature is set to 0.6. Please refer to Appendix A.2 for more implementation details.

### 3.3 Main Results

Table 1 presents the results on 6 reasoning benchmarks across 2 representative open-source LLMs.

**$\phi$ -Decoding significantly enhances the average performances of backbone LLMs.** Compared with the standard CoT strategy,  $\phi$ -Decoding can achieve the inference-time optimization without extra training. Specifically, notable average improvements of 14.62% and 6.92% are observed in LLaMA3.1-Instruct and Mistral-v0.3-Instruct models respectively.

**$\phi$ -Decoding strikes a superior balance between effectiveness and efficiency over strong baselines.** In general,  $\phi$ -Decoding outperforms the four strong baselines by a large margin, with consistent lower computational cost. Compared with the recent promising MCTS-style method,  $\phi$ -Decoding showcases a notable average improvement of 3.25-5.70%, achieved with one-third of the cost. When contrasted with the recent SOTA baseline *Predictive Decoding*,  $\phi$ -Decoding shows remarkable superiority particularly in its adeptness at generalizing across various backbone LLMs.

### 3.4 On the Inference-Time Scaling

Figure 3 presents the inference-time scaling law on LLaMA3.1-8B-Instruct. From the scaling curves,  $\phi$ -Decoding presents the consistent superiority on each computational budget, ranging from  $8 \times 10^{16}$

Models	GSM8K	Math-500	GPQA	ReClor	LogiQA	ARC-c	Avg.	FLOPS
LLaMA3.1-8B-Instruct								
Auto-Regressive (CoT)	70.28	31.00	26.56	49.40	33.33	58.91	44.91	$1.34 \times 10^{16}$
Tree-of-Thoughts	75.74	31.60	<u>31.25</u>	59.00	45.93	80.72	54.04	$7.03 \times 10^{17}$
MCTS	80.44	<u>34.40</u>	24.11	61.40	42.70	79.95	53.83	$17.90 \times 10^{17}$
Guided Decoding	75.51	31.20	30.58	60.20	43.47	81.74	53.78	$6.54 \times 10^{17}$
Predictive Decoding	<u>81.43</u>	34.00	31.03	<b>64.00</b>	<u>46.70</u>	<u>84.56</u>	<u>56.95</u>	$6.89 \times 10^{17}$
$\phi$ -Decoding	<b>86.58</b>	<b>38.20</b>	<b>34.60</b>	<b>64.00</b>	<b>48.39</b>	<b>85.41</b>	<b>59.53</b>	$6.43 \times 10^{17}$
Mistral-v0.3-7B-Instruct								
Auto-Regressive (CoT)	49.05	<u>12.20</u>	23.88	52.20	37.02	69.54	40.65	$0.81 \times 10^{16}$
Tree-of-Thoughts	53.90	10.80	26.34	55.60	<u>41.63</u>	73.63	43.65	$4.99 \times 10^{17}$
MCTS	<u>60.12</u>	10.80	22.77	<u>56.80</u>	40.71	<u>74.74</u>	<u>44.32</u>	$9.33 \times 10^{17}$
Guided Decoding	53.90	10.80	<u>27.46</u>	53.20	36.71	73.55	42.60	$7.03 \times 10^{17}$
Predictive Decoding	58.00	11.00	22.10	54.20	39.78	73.55	43.11	$4.73 \times 10^{17}$
$\phi$ -Decoding	<b>60.42</b>	<b>16.40</b>	<b>29.24</b>	<b>58.20</b>	<b>43.01</b>	<b>78.16</b>	<b>47.57</b>	$3.55 \times 10^{17}$

Table 1: Main results. The optimal results are highlighted in bold, whereas suboptimal results are underlined. The Avg. column indicates the averaged results across the six benchmarks. *FLOPS* denotes the calculated computational cost, with lower values indicating lower costs.

to  $64 \times 10^{16}$  FLOPS. Furthermore, when considering similar performance levels (e.g., an average performance of  $\sim 57\%$ ),  $\phi$ -Decoding demonstrates over  $6\times$  efficiency compared to even suboptimal methods. Meanwhile, it is observed that *Predictive Decoding* and *ToT* also exhibit the stable improvement trend with the inference cost increasing.

## 4 Analysis

### 4.1 Ablation Studies

Some key components of  $\phi$ -Decoding are ablated to verify their contributions to the overall performances in Table 2. *w/o foresight sampling* indicates that the look-ahead process is ablated, relying solely on step uncertainty for sampling. *w/o cluster* denotes that we simply sample on the foresight uncertainty distribution without considering the cluster distribution. *w/o dynamic pruning* means the breadth and depth pruning strategies are ablated. We have the following findings.

#### Foresight sampling mitigates auto-regressive generation limitations with extra inference cost.

As the basis of our sampling strategy, simulating the future steps brings remarkable performance gains (2.98%-6.09%). It proves the finding that the short-sightedness of the standard auto-regressive language generation can be reduced by increasing the inference-time computation.

#### Cluster distribution is beneficial to the overall performances.

As one of the contributions, we

incorporate the cluster of foresight steps to mitigate the unreliability of the accumulated generation probability. The results demonstrate that the cluster can calibrate the sampling distribution, leading to 0.95%-1.97% average gains.

#### Dynamic pruning largely reduces the computational costs.

It is observed that the dynamic pruning strategy provides obvious efficiency improvement from the metric of *FLOPS*. Also, the dynamic pruning strategy surprisingly enhances model performance by eliminating distractions from negative rollouts during sampling.

### 4.2 Generalization and Scalability

Next, we analyze the generalization and scalability of  $\phi$ -Decoding to (i) larger backbone LLM; and (ii) competition-level benchmarks.

#### $\phi$ -Decoding still works when scaling to 70B model size.

Figure 3 shows the results on LLaMA3.1-70B-Instruct across four benchmarks. The model performance is further enhanced with the proposed algorithm. It uncovers the superior generalization capability of  $\phi$ -Decoding. Limited by space, we leave the discussion of smaller backbone LLM (i.e., Qwen2.5-3B-Inst.) for Appendix C. The experiments on the 3B-sized model also reflect the obvious advantages brought by  $\phi$ -Decoding. Across the 6 reasoning benchmarks,  $\phi$ -Decoding improves the backbone LLM by 3.80% in average. Combining all these generalization experiments, it is concluded that  $\phi$ -Decoding works

Models	GSM8K	Math-500	GPQA	ReClor	LogiQA	ARC-c	Avg.	FLOPS
LLaMA3.1-8B-Instruct								
$\phi$ -Decoding	<b>86.58</b>	<b>38.20</b>	<b>34.60</b>	<b>64.00</b>	<b>48.39</b>	<b>85.41</b>	<b>59.53</b>	$6.43 \times 10^{17}$
w/o foresight sampling	81.80	35.00	30.58	60.60	46.39	84.90	56.55	$1.27 \times 10^{17}$
w/o cluster	85.60	37.40	30.58	61.00	45.47	85.32	57.56	$6.37 \times 10^{17}$
w/o dynamic pruning	86.35	38.20	29.46	61.00	46.39	85.67	57.85	$8.00 \times 10^{17}$
Mistral-v0.3-7B-Instruct								
$\phi$ -Decoding	<b>60.42</b>	<b>16.40</b>	<b>29.24</b>	<b>58.20</b>	<b>43.01</b>	<b>78.16</b>	<b>47.57</b>	$3.55 \times 10^{17}$
w/o foresight sampling	57.54	11.40	25.22	42.40	36.70	75.60	41.48	$1.19 \times 10^{17}$
w/o cluster	60.19	15.00	29.24	56.60	42.24	76.45	46.62	$3.55 \times 10^{17}$
w/o dynamic pruning	59.97	15.20	26.56	53.20	36.41	75.77	44.52	$6.41 \times 10^{17}$

Table 2: Ablation Studies on LLaMA3.1-8B-Instruct and Mistral-v0.3-7B-Instruct models. *w/o foresight sampling* ablates the simulation of future steps. *w/o cluster* ablates the calculation of *Alignment* value. *w/o dynamic pruning* ablates both of the pruning strategies.

Tasks	AR (CoT)	$\phi$ -Decoding	$\Delta$
GSM8K	92.27	94.31	<b>+2.04</b>
MATH-500	41.40	44.80	<b>+3.40</b>
ReClor	67.60	84.80	<b>+17.20</b>
LogiQA	51.00	56.37	<b>+5.37</b>

Table 3: Generalization experiments on LLaMA3.1-70B-Instruct. The improvements over Auto-Regressive (CoT) are reported in the last column.

Methods	AIME2024	$\Delta$
LLaMA3.1-8B-Instruct	9.17	-
+ Predictive Decoding	13.33	<b>+4.16</b>
+ $\phi$ -Decoding	16.67	<b>+7.50</b>
R1-Distill-LLaMA-8B	37.81	-
+ Predictive Decoding	20.00	-17.81
+ $\phi$ -Decoding	46.67	<b>+8.86</b>

Table 4: Results on AIME 2024. We compare  $\phi$ -Decoding with Predictive-Decoding based on two backbone LLMs: LLaMA3.1-8B-Instruct and R1-Distill-LLaMA-8B.

well with a wide size range of LLMs, showcasing the superiority.

**Our inference-time optimization can scale to improve performances on the competition-level task even with the strongest reasoning LLM.** Table 4 shows the results on AIME 2024 benchmark. In addition to LLaMA3.1-8B-Instruct. and Mistral-v0.3-7B-Instruct., we also incorporate the DeepSeek-R1 model, utilizing the R1-Distill-LLaMA-8B variant due to resource constraints. Even based on a well-trained deep thinking model,  $\phi$ -Decoding can still help push the upper boundary on the competition-level task. Such findings

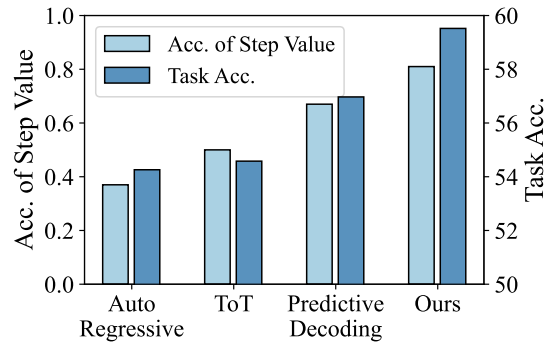


Figure 4: Analysis on the accuracy of step value estimation. The bar in light blue represents the accuracy of the step values, while the bar in dark blue denotes the averaged task performances.

are exciting and insightful to implement inference-time optimization aimed at addressing challenging problems with LLM.

### 4.3 Accuracy of Step Value Estimation

The core of these decoding approaches is to estimate the precise step value through self-rewarding. To measure how the estimated step value matches the actual rewards, we employ the calculation of the *Accuracy of Step Value* via distribution match. Please refer to Appendix D for details. Based on the calculation, we visualize the results in Figure 4, revealing the following finding.

**The estimation of step value is positively correlated with the correctness of the final answer.** Of the four inference-time decoding approaches illustrated in Figure 4, a more accurate estimation of the step value results in improved task performance. Among them,  $\phi$ -Decoding achieves the optimal estimation of step values as well as the final accuracy

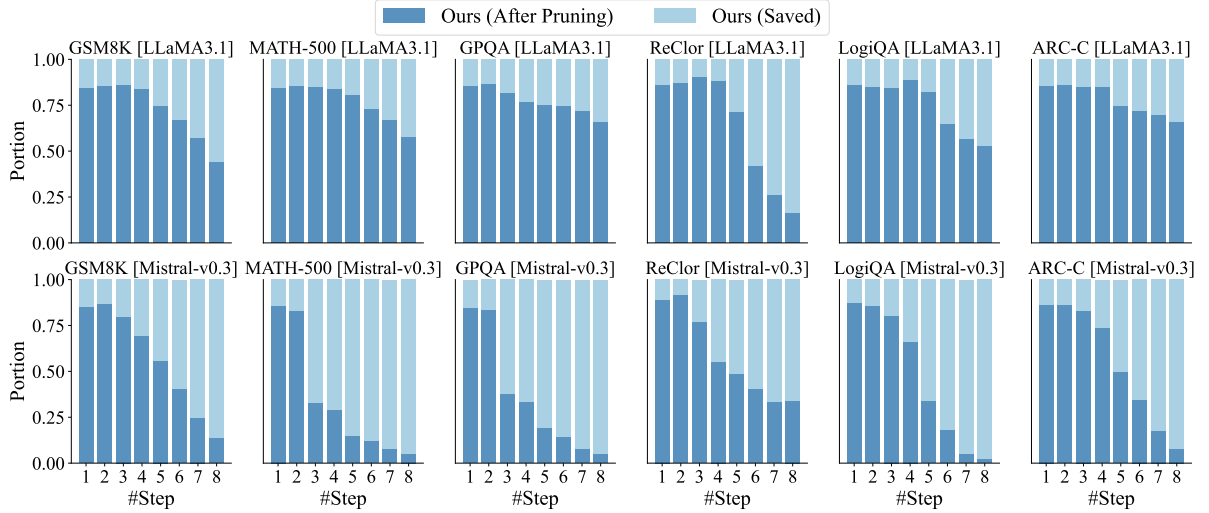


Figure 5: Visualization of step-wise effects with alleviated overthinking. The first row displays the results for each independent benchmark using the LLaMA backbone, whereas the second row reflects the results with the Mistral backbone.

with obvious advantages.

#### 4.4 Analysis on Step-wise Overthinking

Beyond simply reporting the *FLOPS* metric, a detailed analysis of the effects of pruning strategies is presented in Figure 5. It is observed that early steps are more critical, which involves relatively more computational costs. At these early steps, it mainly relies on breadth pruning strategy to avoid redundant step exploration, reducing  $\sim 20\%$  of the costs. With the steps growing, depth pruning takes over to alleviate overthinking through early stopping. This finding inspires us to allocate more inference-time computational resources to the early steps, which are proved to be critical for the reasoning tasks.

### 5 Related Works

**Inference-Time Optimization.** To alleviate the post-training workload (Zelikman et al., 2024; Liu et al., 2024; Team, 2024; Guo et al., 2025; Sun et al., 2024b), inference-time optimization methods arouse wide concerns, showcasing a notable performance boost in reasoning scenarios (Snell et al., 2024; Sun et al., 2023; Zhao et al., 2024; Zhang et al., 2025). Mainstream methods can be categorized into searching-based (Yao et al., 2024; Hao et al., 2023; Xie et al., 2024; Wu et al., 2024) and sampling-based (Ma et al., 2024; Chen et al., 2023; Zhang et al., 2024). Although these works achieve the globally-optimal inference, they either induce large computation costs or yield inadequate step value estimation. Other classical methods, such as Best-of-N, usually involve delegating the

step selection to the external reward model (Wang et al., 2024; Guan et al., 2025), and self-reflection strategies (Cheng et al., 2024; Xu et al., 2024) usually involve extra training.  $\phi$ -Decoding stands out as an optimal and efficient decoding choice without reliance on external auxiliary.

**Adaptive Inference-time Scaling.** Though scaling of inference-time computations has proved to be effective (Snell et al., 2024), the issue of over-thinking is widely observed and remains to be addressed (Chen et al., 2024). One line of works (Team et al., 2025; Han et al., 2024) stress on the control of the generation length, while another line of methods (Manvi et al., 2024; Sun et al., 2024a) leverage the idea of early-stopping. In contrast, the adaptive scaling technique presented in our work is training-free and independent of external models. Based on the self-evaluation of step-wise value,  $\phi$ -Decoding introduces the comprehensive pruning strategy from the dimensions of width and depth. It stands out as a light-weight solution to alleviate the inference-time over-thinking.

### 6 Conclusion

This work focuses on inference-time optimization for LLMs, leveraging computational scaling to enhance performance. Building on stepwise reasoning and foresight sampling, we address two key research questions: (1) *How can we achieve superior step value estimation?* and (2) *Is deliberative planning necessary for every step?* We introduce a novel decoding strategy,  $\phi$ -Decoding,

that efficiently balances exploration and exploitation during inference. Extensive evaluations across seven diverse LLM benchmarks demonstrate  $\phi$ -Decoding’ state-of-the-art performance and efficiency. Furthermore, its ability to generalize to a wide range of LLMs (3B, 7B, 8B, and 70B) and scale across various computational budgets underscores the superiority of  $\phi$ -Decoding in inference-time optimization.

## Acknowledgement

This work was supported by National Key Research and Development Program of China (2022YFC3303600), National Natural Science Foundation of China (No. 62137002, 62293550, 62293553, 62293554, 62437002, 62477036, 62176209, 62176207), "LENOVO-XJTU" Intelligent Industry Joint Laboratory Project, Shaanxi Undergraduate and Higher Education Teaching Reform Research Program (No. 23BY195), and Xi’an Jiaotong University City College Research Project (No. 2024Y01), Project of China Knowledge Centre for Engineering Science and Technology, the Youth AI Talents Fund of China Association of Automation (Grant No.HBRC-JKYZD-2024-311).

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. 2024. [Do not think that much for  \$2+3=?\$  on the overthinking of o1-like llms](#). *arXiv preprint arXiv:2412.21187*.
- Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. 2023. [Universal self-consistency for large language model generation](#). *arXiv preprint arXiv:2311.17311*.
- Kanzhi Cheng, Yantao Li, Fangzhi Xu, Jianbing Zhang, Hao Zhou, and Yang Liu. 2024. [Vision-language models can self-improve reasoning via reflection](#). *arXiv preprint arXiv:2411.00855*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. [Training verifiers to solve math word problems](#). *arXiv preprint arXiv:2110.14168*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- Xinyu Guan, Li Lina Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, and Mao Yang. 2025. [rstar-math: Small llms can master math reasoning with self-evolved deep thinking](#). *arXiv preprint arXiv:2501.04519*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*.
- Tingxu Han, Chunrong Fang, Shiyu Zhao, Shiqing Ma, Zhenyu Chen, and Zhenting Wang. 2024. [Token-budget-aware llm reasoning](#). *arXiv preprint arXiv:2412.18547*.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. [Reasoning with language model is planning with world model](#). *arXiv preprint arXiv:2305.14992*.
- 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](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Maxwell Jia. 2024. [Aime 2024](#).
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *arXiv preprint arXiv:2001.08361*.
- 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*.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2021. [Logiqa: a](#)

- challenge dataset for machine reading comprehension with logical reasoning. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3622–3628.
- Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. 2024. [Acemath: Advancing frontier math reasoning with post-training and reward modeling](#). *arXiv preprint arXiv:2412.15084*.
- Chang Ma, Haiteng Zhao, Junlei Zhang, Junxian He, and Lingpeng Kong. 2024. [Non-myopic generation of language models for reasoning and planning](#). *arXiv preprint arXiv:2410.17195*.
- Rohin Manvi, Anikait Singh, and Stefano Ermon. 2024. [Adaptive inference-time compute: LLMs can predict if they can do better, even mid-generation](#). *arXiv preprint arXiv:2410.02725*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *arXiv preprint arXiv:2311.12022*.
- 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. 2024a. [Fast best-of-n decoding via speculative rejection](#). *arXiv preprint arXiv:2410.20290*.
- Qiushi Sun, Zhirui Chen, Fangzhi Xu, Kanzhi Cheng, Chang Ma, Zhangyue Yin, Jianing Wang, Chengcheng Han, Renyu Zhu, Shuai Yuan, et al. 2024b. [A survey of neural code intelligence: Paradigms, advances and beyond](#). *arXiv preprint arXiv:2403.14734*.
- Qiushi Sun, Zhangyue Yin, Xiang Li, Zhiyong Wu, Xipeng Qiu, and Lingpeng Kong. 2023. [Corex: Pushing the boundaries of complex reasoning through multi-model collaboration](#). *arXiv preprint arXiv:2310.00280*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. [Gemini: a family of highly capable multimodal models](#). *arXiv preprint arXiv:2312.11805*.
- Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. 2025. [Kimi k1. 5: Scaling reinforcement learning with llms](#). *arXiv preprint arXiv:2501.12599*.
- Qwen Team. 2024. [Qwq: Reflect deeply on the boundaries of the unknown](#).
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. 2024. [Math-shepherd: Verify and reinforce llms step-by-step without human annotations](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers), pages 9426–9439.
- Xuezhi Wang and Denny Zhou. 2024. [Chain-of-thought reasoning without prompting](#). *arXiv preprint arXiv:2402.10200*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in neural information processing systems*, 35:24824–24837.
- Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. 2024. [Inference scaling laws: An empirical analysis of compute-optimal inference for problem-solving with language models](#). *arXiv preprint arXiv:2408.00724*.
- Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, and Michael Xie. 2024. [Self-evaluation guided beam search for reasoning](#). *Advances in Neural Information Processing Systems*, 36.
- Fangzhi Xu, Qiushi Sun, Kanzhi Cheng, Jun Liu, Yu Qiao, and Zhiyong Wu. 2024. [Interactive evolution: A neural-symbolic self-training framework for large language models](#). *arXiv preprint arXiv:2406.11736*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. [Qwen2. 5 technical report](#). *arXiv preprint arXiv:2412.15115*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. [Tree of thoughts: Deliberate problem solving with large language models](#). *Advances in Neural Information Processing Systems*, 36.
- Weihao Yu, Zihang Jiang, Yanfei Dong, and Jiashi Feng. 2020. [Reclor: A reading comprehension dataset requiring logical reasoning](#). In *International Conference on Learning Representations*.
- Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D Goodman. 2024. [Quiet-star: Language models can teach themselves to think before speaking](#). *arXiv preprint arXiv:2403.09629*.

Xiaoying Zhang, Baolin Peng, Ye Tian, Jingyan Zhou, Lifeng Jin, Linfeng Song, Haitao Mi, and Helen Meng. 2024. [Self-alignment for factuality: Mitigating hallucinations in llms via self-evaluation](#). *arXiv preprint arXiv:2402.09267*.

Xinyu Zhang, Yuxuan Dong, Yanrui Wu, Jiaxing Huang, Chengyou Jia, Basura Fernando, Mike Zheng Shou, Lingling Zhang, and Jun Liu. 2025. [Physreason: A comprehensive benchmark towards physics-based reasoning](#). *arXiv preprint arXiv:2502.12054*.

Haiteng Zhao, Chang Ma, Guoyin Wang, Jing Su, Lingpeng Kong, Jingjing Xu, Zhi-Hong Deng, and Hongxia Yang. 2024. [Empowering large language model agents through action learning](#). *arXiv preprint arXiv:2402.15809*.

## A Implementation Details

### A.1 Calculation of FLOPS

Following (Kaplan et al., 2020), we calculate the inference-time FLOPS to measure the computational efficiency. The definition of the metric FLOPS is:

$$\text{FLOPS} \approx 6nP \quad (12)$$

where  $n$  represents the total number of the output tokens, and  $P$  is the number of parameters of the LLM. In the tables above, we report the average results of FLOPS across the benchmarks.

### A.2 Inference Setup

We provide the detailed implementation setup in Table 5. Considering the huge cost ahead, the hyperparameters are merely searched within a very small range. We leave it for future works to derive the optimal experimental configuration.

Task	Hyper-Parameter Setup					
LLaMA3.1-8B-Instruct						
GSM8K	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
MATH-500	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
GPQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(1,8)$	$K=3$	$\delta=0.7$	
ReClor	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
LogiQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
ARC-C	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
AIME2024	$M=3$	$N=2$	$(T_{\min}, T_{\max})=(32,64)$	$K=3$	$\delta=0.7$	
Mistralv0.3-7B-Instruct						
GSM8K	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(2,8)$	$K=3$	$\delta=0.7$	
MATH-500	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(1,8)$	$K=3$	$\delta=0.7$	
GPQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(1,8)$	$K=3$	$\delta=0.7$	
ReClor	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(2,8)$	$K=3$	$\delta=0.7$	
LogiQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(2,8)$	$K=3$	$\delta=0.7$	
ARC-C	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(2,8)$	$K=3$	$\delta=0.7$	
Qwen2.5-3B-Instruct						
GSM8K	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
MATH-500	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
GPQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(3,8)$	$K=3$	$\delta=0.7$	
ReClor	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
LogiQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
ARC-C	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(4,8)$	$K=3$	$\delta=0.7$	
LLaMA3.1-70B-Instruct						
GSM8K	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(7,8)$	$K=3$	$\delta=0.7$	
MATH-500	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(3,8)$	$K=3$	$\delta=0.7$	
ReClor	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(2,8)$	$K=3$	$\delta=0.7$	
LogiQA	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(6,8)$	$K=3$	$\delta=0.7$	
DeepSeek R1-Distill-LLaMA-8B						
AIME2024	$M=4$	$N=4$	$(T_{\min}, T_{\max})=(16,32)$	$K=3$	$\delta=0.7$	

Table 5: Experimental setup of  $\phi$ -Decoding.  $M$  denotes the step beam size.  $N$  is the number of rollouts for each step beam.  $T_{\min}$  and  $T_{\max}$  represent the least and the most foresight step number respectively.  $K$  is the number of clusters while  $\delta$  means the early-stopping threshold using clustering.

## B Algorithm of $\phi$ -Decoding

The pseudo code of  $\phi$ -Decoding is presented in Algorithm 1. To make a high-level overview of  $\phi$ -Decoding, we also provide the pipeline in Figure 6.

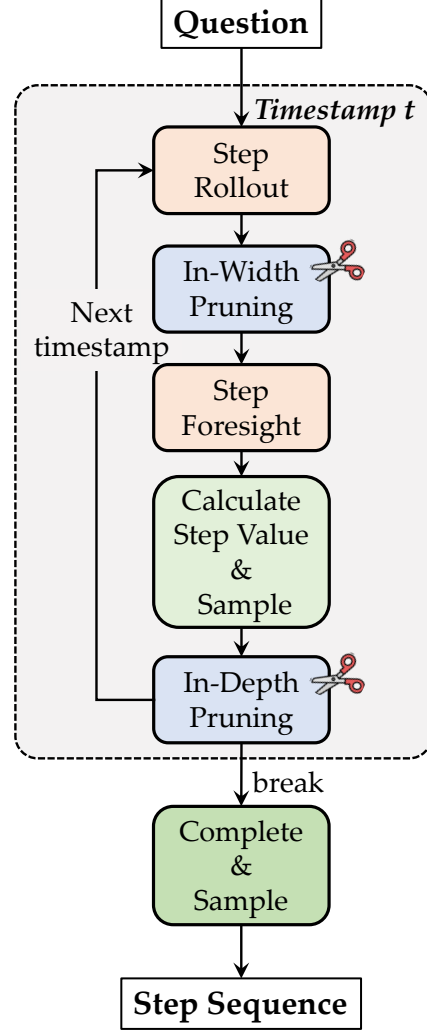


Figure 6: Overall pipeline of  $\phi$ -Decoding.

## C Generalization to Smaller LLMs

Besides the generalization to 70B-sized backbone, we also supplement the evaluations on 3B-sized model. Table 8 presents the performances on Qwen2.5-3B-Instruct model.

Compared with the auto-regressive chain-of-thought baseline,  $\phi$ -Decoding provides obvious performance gains across 6 reasoning benchmarks, improving the average performance by 3.80%.

## D Accuracy of Step Value Estimation

To measure whether the estimated step value aligns with the actual rewards, we conduct the analysis

---

**Algorithm 1**  $\phi$ -Decoding

---

**Input:** Input query  $x$ , LLM  $\pi_\theta$ , step beam size  $M$ , number of rollouts on each beam  $N$ , minimum and maximum number of step foresight  $T_{\min}$  and  $T_{\max}$ , number of clusters  $K$ , early-stopping threshold  $\delta$ .

**Output:** Step sequence.

```
for  $t = 1, 2, \dots, T_{\max}$  do
  ▷ Step Rollout (In Parallel)
  for  $m = 1, \dots, M$  do
    for  $n = 1, 2, \dots, N$  do
      Sample single step  $a_t^{(m,n)}, s_t^{(m,n)} \sim p_\theta(\cdot | x, \mathbf{a}_{<t}^{(m)})$ 
    end for
  end for
  ▷ In-Width Pruning (filter erroneous candidates)
  Derive mean  $\mu_t$  and variance  $\sigma_t^2$  from these step confidence  $s_t$ 
  Prune steps and keep the remaining ones for foresight:  $\mathcal{S}_t \leftarrow \{a_t^{(m,n)} | \mu_t - \sigma_t \leq s_t^{(m,n)}\}$ 
  ▷ Step Foresight (In Parallel)
  for each  $a_t^{(m,n)}$  in  $\mathcal{S}_t$  do
    Derive foresight steps and foresight scores:  $\mathbf{a}_{>t}^{(m,n)}, F_t^{(m,n)} \sim p_\theta(\cdot | x, \mathbf{a}_{\leq t}^{(m,n)})$ 
  end for
  ▷ Step Value Estimation (In Parallel)
  for  $i = 1, 2, \dots, |\mathcal{S}_t|$  do
     $m, n \leftarrow$  the superscript of  $i^{th}$  candidate in  $\mathcal{S}_t$ 
    Derive Advantage via  $\Delta F$  of adjacent steps:  $A_t^{(m,n)} \leftarrow F_t^{(m,n)} - F_{t-1}^{(m)}$ 
    Derive Alignment via clustering:  $C_t^{(m,n)} \leftarrow \text{Cluster}(\{\mathbf{a}_{>t}^{(m,n)}\})$ 
    Combine Advantage and Alignment:  $R(x, \mathbf{a}_{\leq t}^{(m,n)}, \mathbf{a}_{>t}^{(m,n)}) \leftarrow \text{Norm}(A_t^{(m,n)}) + \text{Norm}(C_t^{(m,n)})$ 
     $w_i \leftarrow \exp \left[ R(x, \mathbf{a}_{\leq t}^{(m,n)}, \mathbf{a}_{>t}^{(m,n)}) / \tau \right]$ 
  end for
  ▷ Sample  $M$  Steps
  for  $m = 1, 2, \dots, M$  do
    Sample without replace:  $i \sim \text{Categorical}(\{\frac{w_i}{\sum_j w_j}\}_{i=1}^{|\mathcal{S}_t|})$ 
    Sampled step:  $a_t^{(m)} \leftarrow \mathcal{S}_t[i]$ 
  end for
  ▷ In-Depth Pruning (early-stop)
  break if  $t \geq T_{\max}$  and  $\text{EarlyStop}(\delta)$  is True;
end for
Complete all candidates at the last foresight step and sample only one based on the  $R$  function.
Return Step sequence.
```

---

Cluster Methods	GSM8K	Math-500	GPQA	ReClor	LogiQA	ARC-c	Avg.	FLOPS
$\phi$ -Decoding (LLaMA3.1-8B-Instruct)								
TF-IDF	<b>86.58</b>	38.20	<b>34.60</b>	<b>64.00</b>	<b>48.39</b>	<b>85.41</b>	<b>59.53</b>	$6.43 \times 10^{17}$
SBERT (109M)	86.43	<b>39.20</b>	33.26	63.20	47.48	<b>85.41</b>	59.16	$6.52 \times 10^{17}$
SBERT (22.7M)	86.05	36.80	33.26	62.40	45.47	<b>85.41</b>	58.23	$6.61 \times 10^{17}$

Table 6: Variants of cluster strategies.

$K$	$\sigma$	GSM8K	Math-500	GPQA	ReClor	LogiQA	ARC-c	Avg.	FLOPS
$\phi$ -Decoding (LLaMA3.1-8B-Instruct)									
3	0.7	<b>86.58</b>	38.20	<b>34.60</b>	64.00	<b>48.39</b>	<b>85.41</b>	<b>59.53</b>	$6.43 \times 10^{17}$
2	0.8	85.52	<b>39.40</b>	33.04	<b>64.20</b>	46.85	<b>85.41</b>	59.07	$6.26 \times 10^{17}$
4	0.5	83.93	38.20	32.37	64.00	43.78	84.81	57.85	$6.15 \times 10^{17}$

Table 7: Various setups of cluster.

Tasks	AR(CoT)	$\phi$ -Decoding	$\Delta$
Qwen2.5-3B-Instruct			
GSM8K	78.62	85.60	+6.98
MATH-500	41.00	45.20	+4.20
GPQA	28.57	28.79	+0.22
ReClor	53.60	59.40	+5.80
LogiQA	42.70	46.08	+3.38
ARC-C	77.47	79.69	+2.22
<b>Avg.</b>	53.66	57.46	+3.80

Table 8: Generalization to smaller backbone Qwen2.5-3B-Instruct.

in Sec. 4.3. At each timestamp  $t$ , we can derive the value estimation of the candidate steps via the decoding strategy. These step values can approximate a distribution  $P_1$ . Meanwhile, we can derive the explicit outcome of each candidate step using the foresight paths. Comparing the outcome with ground-truth, the outcome accuracy for these candidate steps can also form a distribution  $P_2$ , where  $|P_1| = |P_2|$ . We derive the distribution matching as the accuracy of step value estimation:

$$\text{Accuracy} = 1 - \frac{\sum_{i=1}^{|P_1|} (P_1(i) - P_2(i))^2}{|P_1|} \quad (13)$$

where  $P_1(i) \in P_1$ ,  $P_2(i) \in P_2$ . In the implementation of  $P_1$ , we use the model estimated step values for sampling-based methods ( $\phi$ -Decoding and Predictive Decoding). For auto-regressive and ToT methods, we allocate binary rewards for the selected steps (rewarded as 1) and other candidates (rewarded as 0). The final accuracy score is calculated by averaging the results on each timestamp.

## E In-depth Analysis of Cluster Strategies

### E.1 Variants of Cluster

In the main experiments, we implement the cluster strategy with TF-IDF, which is from the syntax perspective. It can also be replaced with sentence-BERT (SBERT) (Reimers and Gurevych, 2019) to obtain the sentence embedding for clustering.

Table 6 presents the comparisons between different cluster strategies. SBERT (109M) employs the pretrained sentence embedding model of multi-qa-mpnet-base-dot-v1, while SBERT (22.7M) utilizes the model of all-MiniLM-L6-v2.

From the results, clustering with the external embedding model can also lead to similar competitive performances, slightly lower than the TF-IDF strategy. Also, it is observed that increasing the size of the sentence embedding models can bring improvements in the average performances.

### E.2 Hyperparameters of Cluster

Table 7 offers the analysis on different hyperparameters. We keep the other configuration fixed for fair comparison, where  $M=4$  and  $N=4$ . Under this setting, the maximum number of foresight paths for clustering is 16. Based on the results, the cluster size  $K=2$  or 3 would be good choices. With  $K$  increasing, it may bring much uncertainty.